U.S. DEPARTMENT OF ENERGY | Energy Efficiency & Renewable Energy

FEDERAL ENERGY MANAGEMENT PROGRAM

Best Practices Guide for Energy-Efficient Data Center Design

February 2010

FEMP
Federal Energy Management Program

Prepared by the National Renewable Energy Laboratory (NREL), a national laboratory of the U.S. Department of Energy, Office of Energy Efficiency and Renewable Energy; NREL is operated by the Alliance for Sustainable Energy, LLC.

Acknowledgements

This report was prepared by John Bruschi, Peter Rumsey, Robin Anliker, Larry Chu, and Stuart Gregson of Rumsey Engineers under contract to the National Renewable Energy Laboratory. The work was supported by the Federal Energy Management Program led by Will Lintner.

Contacts

William Lintner
U.S. Department of Energy FEMP
william.lintner.ee.doe.gov
202-586-3120

Bill Tschudi
Lawrence Berkeley National Laboratory
wftschudi@lbl.gov
510-495-2417

Otto VanGeet
National Renewable Energy Laboratory
Otto.VanGeet@nrel.gov
303-384-7369

On the Cover

Center for Disease Control and Prevention's Arlen Specter Headquarters and Operations Center reached LEED© Silver rating through sustainable design and operations that decrease energy consumption by 20% and water consumption by 36% beyond standard codes. PIX 16419.

Table of Contents

Figures

Table

Summary

This guide provides an overview of best practices for energy-efficient data center design which spans the categories of Information Technology (IT) systems and their environmental conditions, data center air management, cooling and electrical systems, on-site generation, and heat recovery. IT system energy efficiency and environmental conditions are presented first because measures taken in these areas have a cascading effect of secondary energy savings for the mechanical and electrical systems. This guide concludes with a section on metrics and benchmarking values by which a data center and its systems energy efficiency can be evaluated. No design guide can offer 'the most energy-efficient' data center design but the guidelines that follow offer suggestions that provide efficiency benefits for a wide variety of data center scenarios.

Background

Data center spaces can consume up to 100 to 200 times as much electricity as standard office spaces. With such large power consumption, they are prime targets for energy-efficient design measures that can save money and reduce electricity use. However, the critical nature of data center loads elevates many design criteria—chiefly reliability and high power density capacity—far above energy efficiency. Short design cycles often leave little time to fully assess efficient design opportunities or consider first cost versus life-cycle-cost issues. This can lead to designs that are simply scaled up versions of standard office space approaches or that reuse strategies and specifications that worked "good enough" in the past without regard for energy performance. This Best Practices Guide has been created to provide viable alternatives to inefficient data center building practices.

Information Technology (IT) Systems

In a typical data center with a highly efficient cooling system, IT equipment loads can account for over half of the entire facility's energy use. Use of efficient IT equipment will significantly reduce these loads within the data center, which consequently will downsize the equipment needed to cool them. Purchasing servers equipped with energy-efficient processors, fans, and power supplies, high-efficient network equipment, consolidating storage devices, consolidating power supplies, and implementing virtualization are the most advantageous ways to reduce IT equipment loads within a data center.

Efficient Servers

Rack servers tend to be the main perpetrators of wasting energy and represent the largest portion of the IT energy load in a typical data center. Servers take up most of the space and drive the entire operation. The majority of servers run at or below 20% utilization most of the time, yet still draw full power during the process. Recently vast improvements in the internal cooling systems and processor devices of servers have been made to minimize this wasted energy.

When purchasing new servers it is recommended to look for products that include variable speed fans as opposed to a standard constant speed fan for the internal cooling component. With variable speed fans it is possible to deliver sufficient cooling while running slower, thus consuming less energy. The Energy Star program aids consumers by recognizing high-efficiency servers. Servers that meet Energy Star efficiency requirements will, on average, be 30% more efficient than standard servers.

Additionally, a throttle-down drive is a device that reduces energy consumption on idle processors, so that when a server is running at its typical 20% utilization it is not drawing full power. This is also sometimes referred to as "power management." Many IT departments fear that throttling down servers or putting idle servers to sleep will negatively impact server reliability; however, hardware itself is designed to handle tens of thousands of on-off cycles. Server power draw can also be modulated by installing "power cycler" software in servers. During low demand, the software can direct individual devices on the rack to power down. Potential power management risks include slower performance and possibly system failure; which should be weighed against the potential energy savings.

Multi-core processor chips allow simultaneous processing of multiple tasks, which leads to higher efficiency in two ways. First, they offer improved performance within the same power and cooling load as compared to single-core processors. Second, they consolidate shared devices over a single processor core. Not all applications are capable of taking advantage of multi-core processors. Graphics-intensive programs and high performance computing still require the higher clock-speed single-core designs.

Further energy savings can be achieved by consolidating IT system redundancies. Consider one power supply per server rack instead of providing power supplies for each server. For a given redundancy level, integrated rack mounted power supplies will operate at a higher load factor (potentially 70%) compared to individual server power supplies (20% to 25%). This increase in power supply load factor vastly improves the power supply efficiency (see Figure 1) in following section on power supplies). Sharing other IT resources such as Central Processing Units (CPU), disk drives, and memory optimizes electrical usage as well. Short term load shifting combined with throttling resources up and down as demand dictates is another strategy for improving long term hardware energy efficiency.

Storage Devices

Power consumption is roughly linear to the number of storage modules used. Storage redundancy needs to be rationalized and right-sized to avoid rapid scale up in size and power consumption.

Consolidating storage drives into a Network Attached Storage or Storage Area Network are two options that take the data that does not need to be readily accessed and transports it offline. Taking superfluous data offline reduces the amount of data in the production environment, as well as all the copies. Consequently, less storage and CPU requirements on the servers are needed, which directly corresponds to lower cooling and power needs in the data center.

For data that cannot be taken offline, it is recommended to upgrade from traditional storage methods to thin provisioning. In traditional storage systems an application is allotted a fixed amount of anticipated storage capacity, which often results in poor utilization rates and wasted energy. Thin provisioning technology, in contrast, is a method of maximizing storage capacity utilization by drawing from a common pool of purchased shared storage on an as-need basis, under the assumption that not all users of the storage pool will need the entire space simultaneously. This also allows for extra physical capacity to be installed at a later date as the data approaches the capacity threshold.

Network Equipment

As newer generations of network equipment pack more throughput per unit of power, there are active energy management measures that can also be applied to reduce energy usage as network demand varies. Such measures include idle state logic, gate count optimization, memory access algorithms and Input/Output buffer reduction.

As peak data transmission rates continue to increase, requiring dramatically more power, increasing energy is required to transmit small amounts of data over time. Ethernet network energy efficiency can be substantially improved by quickly switching the speed of the network links to the amount of data that is currently transmitted.

Power Supplies

Most data center equipment uses internal or rack mounted alternating current/direct current (AC-DC) power supplies. Historically, a typical rack server's power supply converted AC power to DC power at efficiencies of around 60% to 70%. Today, through the use of higher-quality components and advanced engineering, it is possible to find power supplies with efficiencies up to 95%. Using higher efficiency power supplies will directly lower a data center’s power bills and indirectly reduce cooling system cost and rack overheating issues. At $0.12/kWh, savings of $2,000 to $6,000 per year per rack (10 kW to 25 kW, respectively) are possible just from improving the power supply efficiency from 75% to 85%. These savings estimates include estimated secondary savings due to lower uninterruptible power supply (UPS) and cooling system loads.

The impact of real operating loads should also be considered to select power supplies that offer the best efficiency at the load level at which they are expected to most frequently operate. The optimal power supply load level is typically in the mid-range of its performance curve: around 40% to 60%, as shown in Figure 1.

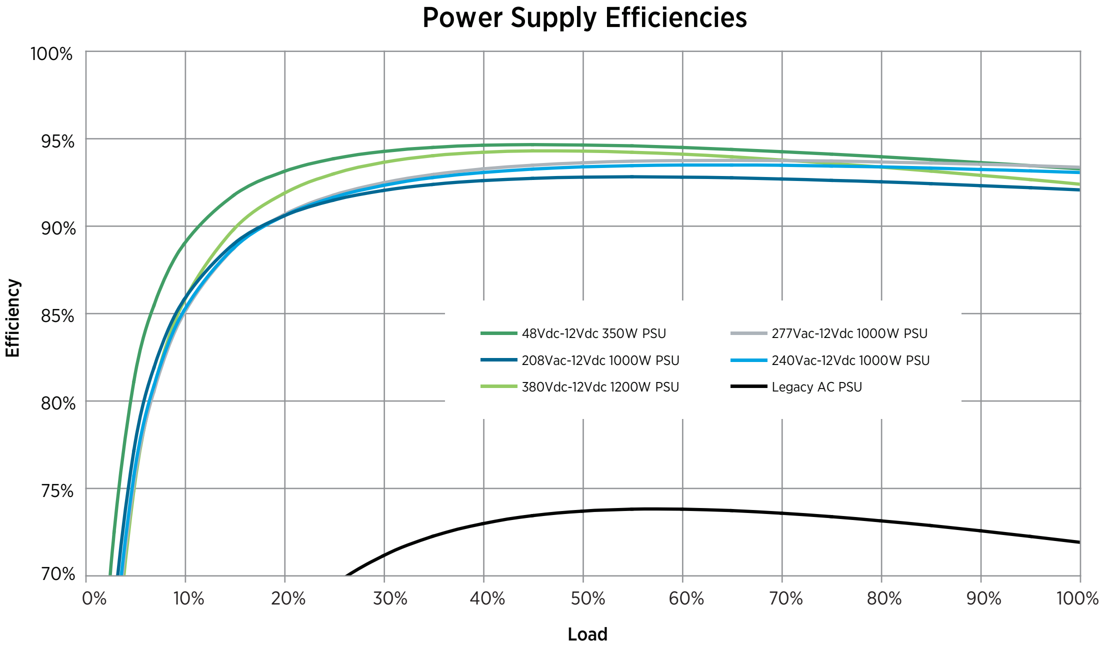

Figure 1. Efficiencies at varying load levels for typical power supplies

(Source: Quantitative Efficiency Analysis of Power Distribution Configurations for Data Centers, The Green Grid)

Efficient power supplies usually have a minimal incremental cost at the server level. Power supplies that meet the recommended efficiency guidelines of the Server System Infrastructure (SSI) Initiative should be selected. There are also several certification programs currently in place that have standardized the efficiencies of power supplies in order for vendors to market their product. For example, the 80 PLUS program offers certifications for power supplies with efficiencies of 80% or greater at 20%, 50%, and 100% of their rated loads with true power factors of 0.9 or greater.

Consolidation

Hardware Location

Lower data center supply fan power and more efficient cooling system performance can be achieved when equipment with similar heat load densities and temperature requirements are grouped together. Isolating equipment by environmental requirements of temperature and humidity allow cooling systems to be controlled to the least energy-intensive set points for each location.

This concept can be expanded to data facilities in general. Consolidating underutilized data center spaces to a centralized location can ease the utilization of data center efficiency measures by condensing the implementation to one location, rather than several.

Virtualization

Virtualization is a method of running multiple independent virtual operating systems on a single physical computer. It is a way of allowing the same amount of processing to occur on fewer servers by increasing server utilization. Instead of operating many servers at low CPU utilization, virtualization combines the processing power onto fewer servers that operate at higher utilization. Virtualization will drastically reduce the number of servers in a data center, reducing required server power and consequently the size of the necessary cooling equipment. Some overhead is required to implement virtualization, but this is minimal compared to the savings that can be achieved.

Environmental Conditions

2009 ASHRAE Guidelines and IT-Reliability

The first step in designing the cooling and air management systems in a data center is to look at the standardized operating environments for equipment set forth by the American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) or Network Equipment Building System (NEBS). In 2008, ASHRAE in collaboration with IT equipment manufacturers expanded their recommended environmental envelope for inlet air entering IT equipment. The revision of this envelope allows greater flexibility in facility operations, and contributes to reducing the overall energy consumption. The expanded recommended and allowable envelopes for Class 1 and 2 data centers are shown in Figure 2 and tabulated in Table 1 (for more details on data center type, different levels of altitude, etc., refer to the referenced ASHRAE publication, *Thermal Guidelines for Data Processing Environments*, 2nd Edition).

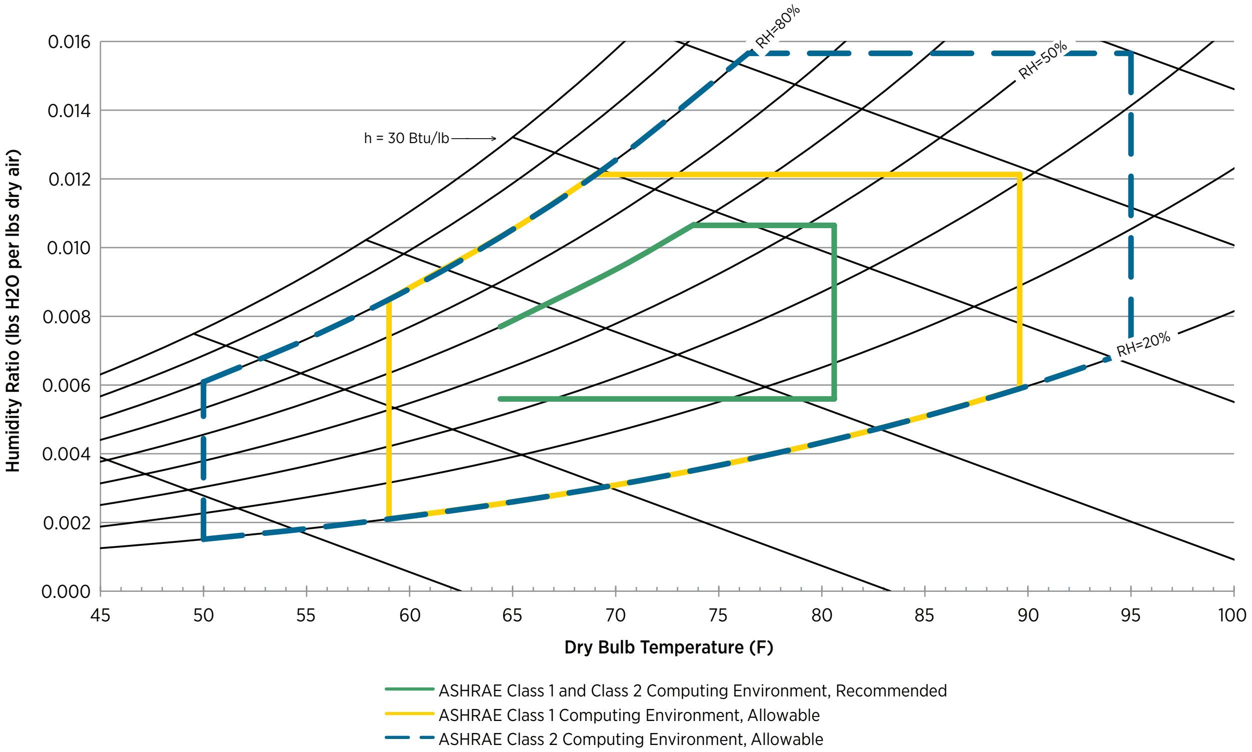

Figure 2. 2009 ASHRAE environmental envelope for IT equipment air intake conditions (Source: Rumsey Engineers)

Table 1. ASHRAE Recommended and Allowable Inlet Air Conditions for Class 1 and 2 Data Centers (Source: Rumsey Engineers)

	Class 1 and Class 2 Recommended Range	Class 1 Allowable Range	Class 2 Allowable Range
Low Temperature Limit	64.4°F DB	59°F DB	50°F DB
High Temperature Limit	80.6°F DB	89.6°F DB	95°F DB
Low Moisture Limit	41.9°F DP	20% RH	20% RH
High Moisture Limit	60% RH & 59°F DP	80% RH & 62.6°F DP	80% RH & 69.8°F DP

It is important to recognize the difference between the recommended and allowable envelopes presented in the ASHRAE guidelines. The recommended environmental envelope is intended to guide operators of data centers on the energy-efficient operation of data centers while maintaining high reliability. The allowable envelope outlines the environmental boundaries tested by equipment manufacturers for equipment functionality, not reliability.

Another important factor to consider regarding the optimal server inlet air temperature is that variable speed fans in the servers are usually controlled to the internal server temperature. Operating the data center at server inlet air conditions above the recommended range may cause these internal fans to operate at higher speeds and consume more power. For example, a data sheet for a Dell PowerEdge blade server indicates a 30% increase in server fan speed with an increase in inlet air temperature from 77°F to 91°F. This increase in inlet air temperature results in more than doubling the server fan power by applying the fan affinity law where fan power increases with the cube of fan speed. Thus, the effect of increasing server inlet air temperature on server fan power should be carefully weighed against the potential data center cooling system energy savings.

Air Management

Air management for data centers entails all the design and configuration details that go into minimizing or eliminating mixing between the cooling air supplied to equipment and the hot air rejected from the equipment. Effective air management implementation minimizes the bypass of cooling air around rack intakes and the recirculation of heat exhaust back into rack intakes. When designed correctly, an air management system can reduce operating costs, reduce first cost equipment investment, increase the data center's power density (Watts/square foot), and reduce heat related processing interruptions or failures. A few key design issues include the configuration of equipment's air intake and heat exhaust ports, the location of supply and returns, the large scale airflow patterns in the room, and the temperature set points of the airflow.

Implement Cable Management

Under-floor and over-head obstructions often interfere with the distribution of cooling air. Such interferences can significantly reduce the air handlers' airflow as well as negatively affect the air distribution. Cable congestion in raised-floor plenums can sharply reduce the total airflow as well as degrade the airflow distribution through the perforated floor tiles. Both effects promote the development of hot spots.

A minimum effective (clear) height of 24 inches should be provided for raised floor installations. Greater under-floor clearance can help achieve a more uniform pressure distribution in some cases.

A data center should have a cable management strategy to minimize air flow obstructions caused by cables and wiring. This strategy should target the entire cooling air flow path, including the rack-level IT equipment air intake and discharge areas as well as under-floor areas.

Persistent cable management is a key component of maintaining effective air management. Instituting a cable mining program (i.e. a program to remove abandoned or inoperable cables) as part of an ongoing cable management plan will help optimize the air delivery performance of data center cooling systems.

Aisle Separation and Containment

A basic hot aisle/cold aisle configuration is created when the equipment racks and the cooling system's air supply and return are designed to prevent mixing of the hot rack exhaust air and the cool supply air drawn into the racks. As the name implies, the data center equipment is laid out in rows of racks with alternating cold (rack air intake side) and hot (rack air heat exhaust side) aisles between them. Strict hot aisle/cold aisle configurations can significantly increase the air-side cooling capacity of a data center's cooling system.

All equipment is installed into the racks to achieve a front-to-back airflow pattern that draws conditioned air in from cold aisles, located in front of the equipment, and rejects heat out through the hot aisles behind the racks. Equipment with non-standard exhaust directions must be addressed in some way (shrouds, ducts, etc.) to achieve a front-to-back airflow. The rows of racks are placed back-to-back, and holes through the rack (vacant equipment slots) are blocked off on the intake side to create barriers that reduce recirculation, as shown in Figure 3 below. Additionally, cable openings in raised floors and ceilings should be sealed as tightly as possible. With proper isolation, the temperature of the hot aisle no longer impacts the temperature of the racks or the reliable operation of the data center; the hot aisle becomes a heat exhaust. The air-side cooling system is configured to supply cold air exclusively to the cold aisles and pull return air only from the hot aisles.

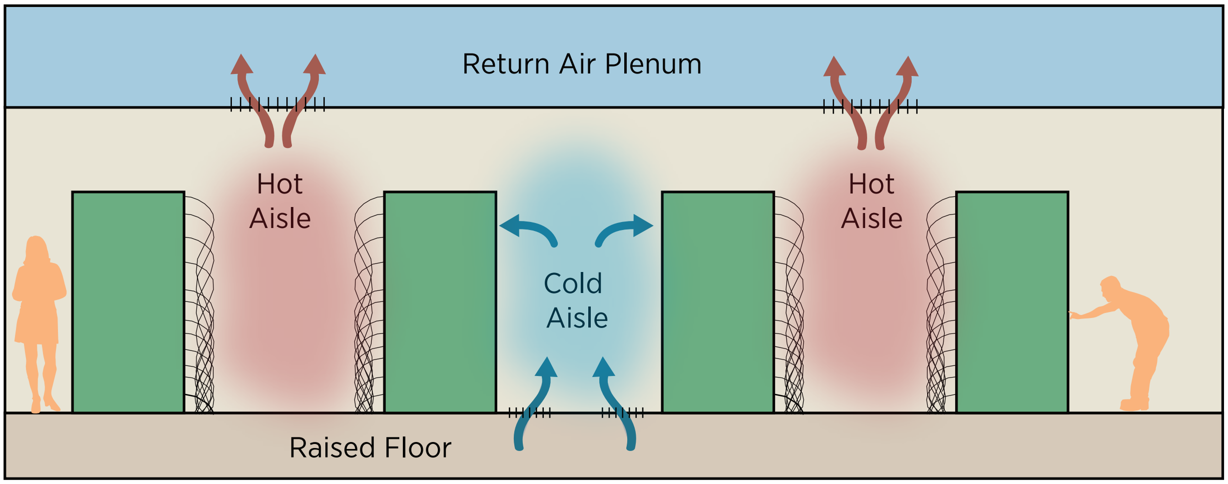

Figure 3. Example of a hot aisle/cold aisle configuration (Source: Rumsey Engineers)

The hot rack exhaust air is not mixed with cooling supply air and, therefore, can be directly returned to the air handler through various collection schemes, returning air at a higher temperature, often 85°F or higher. Depending on the type and loading of a server, the air temperature rise across a server can range from 10°F to more than 40°F. Thus, rack return air temperatures can exceed 100°F when densely populated with highly loaded servers. Higher return temperatures extend economizer hours significantly and allow for a control algorithm that reduces supply air volume, saving fan power. If the hot aisle temperature is high enough, this air can be used as a heat source in many applications. In addition to energy savings, higher equipment power densities are also better supported by this configuration. The significant increase in economizer hours afforded by a hot aisle/cold aisle configuration can improve equipment reliability in mild climates by providing emergency compressor-free data center operation when outdoor air temperatures are below the data center equipment's top operating temperature (typically 90°F to 95°F).

Using flexible plastic barriers, such as plastic supermarket refrigeration covers (i.e. "strip curtains"), or other solid partitions to seal the space between the tops of the rack and air return location can greatly improve hot aisle/cold aisle isolation while allowing flexibility in accessing, operating, and maintaining the computer equipment below. One recommended design configuration, shown in Figure 4, supplies cool air via an under-floor plenum to the racks; the air then passes through the equipment in the rack and enters a separated, semi-sealed area for return to an overhead plenum. This approach uses a baffle panel or barrier above the top of the rack and at the ends of the hot aisles to mitigate "short-circuiting" (the mixing of hot and cold air). These changes should reduce fan energy requirements by 20% to 25%, and could result in a 20% energy savings on the chiller side provided these components are equipped with variable speed drives (VSDs).

Fan energy savings are realized by reducing fan speeds to supply only as much air as a given space requires. There are a number of different design strategies that reduce fan speeds. Among them is a fan speed control loop controlling the cold aisles' temperature at the most critical locations—the top of racks for under-floor supply systems, the bottom of racks for overhead systems, end of aisles, etc. Note that many Direct Expansion (DX) Computer Room Air Conditioners (CRACs) use the return air temperature to indicate the space temperature, an approach that does not work in a hot aisle/cold aisle configuration where the return air is at a very different temperature than the cold aisle air being supplied to the equipment. Control of the fan speed based on the IT equipment needs is critical to achieving savings. Unfortunately variable speed drives on DX CRAC unit supply fans are not generally available despite being a common feature for commercial packaged DX air conditioning units. For more descriptions on CRAC units and common energy-efficiency options, refer to the "DX Systems" subsection of "Cooling Systems" below.

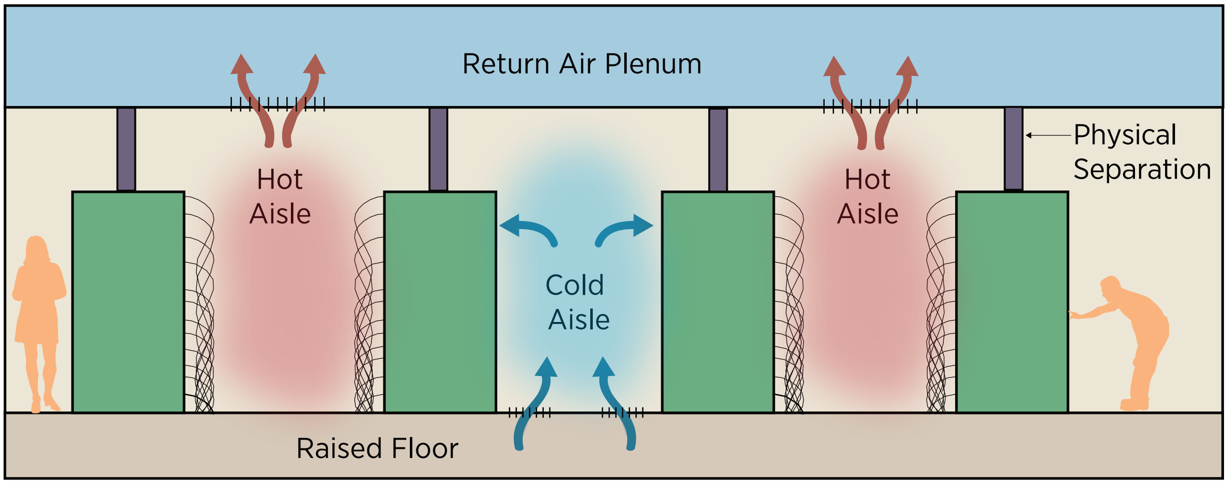

Figure 4. Sealed hot aisle/cold aisle configuration (Source: Rumsey Engineers)

Optimize Supply and Return Air Configuration

Hot aisle/cold aisle configurations can be served by overhead or under-floor air distribution systems. When an overhead system is used, supply outlets that ‘dump’ the air directly down should be used in place of traditional office diffusers that throw air to the sides, which results in undesirable mixing and recirculation with the hot aisles. The diffusers should be located directly in front of racks, above the cold aisle. In some cases return grilles or simply open ducts have been used. The temperature monitoring to control the air handlers should be located in areas in front of the computer equipment, not on a wall behind the equipment. Use of overhead variable air volume (VAV) allows equipment to be sized for excess capacity and yet provide optimized operation at part-load conditions with turn down of variable speed fans. Where a rooftop unit is being used, it should be located centrally over the served area—the required reduction in ductwork will lower cost and slightly improve efficiency. Also keep in mind that overhead delivery tends to reduce temperature stratification in cold aisles as compared to under-floor air delivery.

Under-floor air supply systems have a few unique concerns. The under-floor plenum often serves both as a duct and a wiring chase. Coordination throughout design and into construction and operation throughout the life of the center is necessary since paths for airflow can be blocked by electrical or data trays and conduit. The location of supply tiles needs to be carefully considered to prevent short circuiting of supply air and checked periodically if users are likely to reconfigure them. Removing or adding tiles to fix hot spots can cause problems throughout the system. Another important concern to be aware of is high air velocity in the under-floor plenum. This can create localized negative pressure and induce room air back into the under-floor plenum. Equipment closer to down flow CRAC units or Computer Room Air Handlers (CRAH) can receive too little cooling air due to this effect. Deeper plenums and careful layout of CRAC/CRAH units allow for a more uniform under-floor air static pressure. For more description on CRAH unites as they relate to data center energy efficiency, refer to the "Air Handler" subsection of "Cooling Systems" below.

Raising Temperature Set Points

Higher supply air temperature and a higher difference between the return air and supply air temperatures increases the maximum load density possible in the space and can help reduce the size of the air-side cooling equipment required, particularly when lower-cost mass produced package air handling units are used. The lower required supply airflow due to raising the air-side temperature difference provides the opportunity for fan energy savings. Additionally, the lower supply airflow can ease the implementation of an air-side economizer by reducing the sizes of the penetrations required for outside air intake and heat exhaust.

Air-side economizer energy savings are realized by utilizing a control algorithm that brings in outside air whenever it is appreciably cooler than the return air and when humidity conditions are acceptable (refer to the "Airside Economizer" subsection of "Free Cooling" for further detail on economizer control optimization). In order to save energy, the temperature outside does not need to be below the data center's temperature set point; it only has to be cooler than the return air that is exhausted from the room. As the return air temperature is increased through the use of good air management, as discussed in the preceding sections, the temperature at which an air-side economizer will save energy is correspondingly increased. Designing for a higher return air temperature increases the number of hours that outside air, or a waterside economizer/free cooling, can be used to save energy.

A higher return air temperature also makes better use of the capacity of standard package units, which are designed to condition office loads. This means that a portion of their cooling capacity is configured to serve humidity (latent) loads. Data centers typically have very few occupants and small outside air requirements, and, therefore, have negligible latent loads. While the best course of action is to select a unit designed for sensible-cooling loads only or to increase the airflow, an increased return air temperature can convert some of a standard package unit's latent capacity into usable sensible capacity very economically. This may reduce the size and/or number of units required.

A warmer supply air temperature set point on chilled water air handlers allows for higher chilled water supply temperatures which consequently improves the chilled water plant operating efficiency. Operation at warmer chilled water temperatures also increases the potential hours that a water-side economizer can be used (refer to the subsection on "Water-Side Economizer" of the "Free Cooling" section for further detail).

Cooling Systems

When beginning the design process and equipment selections for cooling systems in data centers, it is important to always consider initial and future loads, in particular part- and low-load conditions, as the need for digital data is ever-expanding.

Direct Expansion (DX) Systems

Packaged DX air conditioners likely compose the most common type of cooling equipment for smaller data centers. These units are generally available as off-the-shelf equipment from manufacturers (commonly described as CRAC units). There are, however, several options available to improve the energy efficiency of cooling systems employing DX units.

Packaged rooftop units are inexpensive and widely available for commercial use. Several manufacturers offer units with multiple and/or variable speed compressors to improve part-load efficiency. These units reject the heat from the refrigerant to the outside air via an air-cooled condenser. An enhancement to the air-cooled condenser is a device which sprays water over the condenser coils. The evaporative cooling provided by the water spray improves the heat rejection efficiency of the DX unit. Additionally, these units are commonly offered with air-side economizers. Depending on the data center's climate zone and air management, a DX unit with air-side economizer can be a very energy-efficient cooling option for a small data center. (For further discussion, refer to section "Raising Temperature Set Points" and subsection "Air-Side Economizer" of the section on "Free Cooling."

Indoor CRAC units are available with a few different heat rejection options. Air-cooled CRAC units include a remote air-cooled condenser. As with the rooftop units, adding an evaporative spray device can improve the air-cooled CRAC unit efficiency. For climate zones with a wide range of ambient dry bulb temperatures, apply parallel VSD control of the condenser fans to lower condenser fan energy compared to the standard staging control of these fans.

CRAC units packaged with water-cooled condensers are often paired with outdoor drycoolers. The heat rejection effectiveness of outdoor drycoolers depends on the ambient dry bulb temperature. A condenser water pump distributes the condenser water from the CRAC units to the drycoolers. Compared to the air-cooled condenser option, this water-cooled system requires an additional pump and an additional heat exchanger between the refrigerant loop and the ambient air. As a result, this type of water-cooled system is generally less efficient than the air-cooled option. A more efficient method for water-cooled CRAC unit heat rejection employs a cooling tower. To maintain a closed condenser water loop to the outside air, a closed loop cooling tower can be selected. A more expensive but more energy-efficient option would be to select an oversized open-loop tower and a separate heat exchanger where the latter can be selected for a very low (less than 3°F) approach. In dry climates, a system composed of water-cooled CRAC units and cooling towers can be designed to be more energy efficient than air-cooled CRAC unit systems. (Refer to the “Efficient Equipment” subsection of the section on “High-Efficiency Chilled Water Systems” for more information on selecting an efficient cooling tower.)

A type of water-side economizer can be integrated with water-cooled CRAC units. A pre-cooling water coil can be added to the CRAC unit upstream of the evaporator coil. When ambient conditions allow the condenser water to be cooled (by either drycooler or cooling tower) to the point that it can provide a direct cooling benefit to the air entering the CRAC unit, condenser water is diverted to the pre-cooling coil. This will reduce, or at times eliminate, the need for compressor-based cooling from the CRAC unit. Some manufacturers offer this pre-cooling coil as a standard option for their water-cooled CRAC units.

Air Handlers

Central vs. Modular Systems

Better performance has been observed in data center air systems that utilize specifically-designed central air handler systems. A centralized system offers many advantages over the traditional multiple distributed unit system that evolved as an easy, drop-in computer room cooling appliance (commonly referred to as a CRAH unit). Centralized systems use larger motors and fans that tend to be more efficient. They are also well suited for variable volume operation through the use of VSDs and maximize efficiency at part-loads.

In Figure 5, the pie charts below show the electricity consumption distribution for two data centers. Both are large facilities, with approximately equivalent data center equipment loads, located in adjacent buildings,

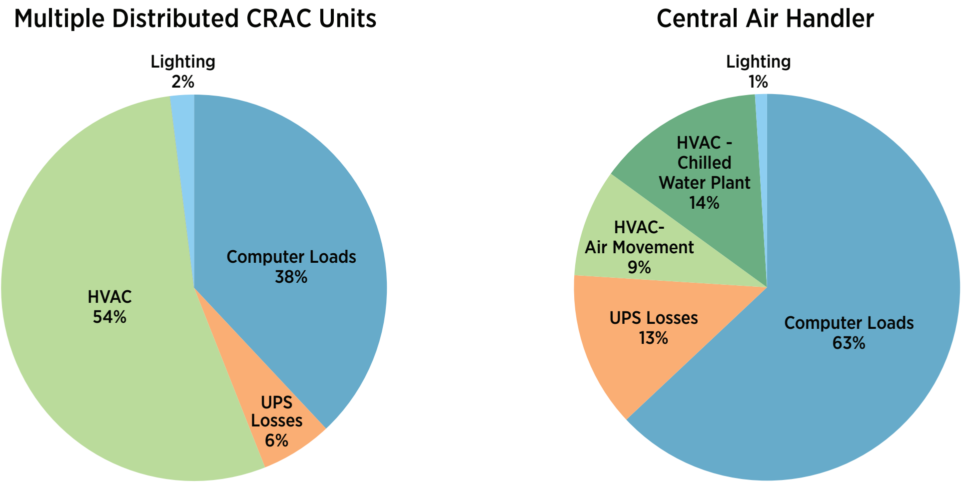

Figure 5. Comparison of distributed air delivery to central air delivery (Source: Rumsey Engineers)

and operated by the same company. The facility on the left uses a multiple distributed unit system based on air-cooled CRAC units, while the facility on the right uses a central air handler system. An ideal data center would use 100% of its electricity to operate data center equipment—energy used to operate the fans, compressors and power systems that support the data center is strictly overhead cost. The data center supported by a centralized air system (on the right) uses almost two thirds of the input power to operate revenue-generating data center equipment, compared to the multiple small unit system that uses just over one third of its power to operate the actual data center equipment. The trend seen here has been consistently supported by benchmarking data. The two most significant energy saving methods are water cooled equipment and efficient centralized air handler systems.

Most data center loads do not vary appreciably over the course of the day, and the cooling system is typically significantly oversized. A centralized air handling system can improve efficiency by taking advantage of surplus and redundant capacity to actually improve efficiency. The maintenance benefits of a central system are well known, and the reduced footprint and maintenance traffic in the data center are additional benefits. Implementation of an airside economizer system is simplified with a central air handler system. Optimized air management, such as that provided by hot aisle/cold aisle configurations, is also easily implemented with a ducted central system. Modular units are notorious for battling each other to maintain data center humidity set points. That is, one unit can be observed to be dehumidifying while an adjacent unit is humidifying. Instead of modular units independently controlled, a centralized control system using shared sensors and set points ensures proper communication among the data center air handlers. Even with modular units humidity control over make-up air should be all that is required.

Low Pressure Drop Air Delivery

A low-pressure drop design (‘oversized’ ductwork or a generous under floor) is essential to optimizing energy efficiency by reducing fan energy and facilitates long-term buildout flexibility. Ducts should be as short as possible in length and sized significantly larger than typical office systems, since 24-hour operation of the data center increases the value of energy use over time relative to first cost. Since loads often only change when new servers or racks are added or removed, periodic manual air flow balancing can be more cost effective than implementing an automated air flow balancing control scheme.

High-Efficiency Chilled Water Systems

Efficient Equipment

Use efficient water-cooled chillers in a central chilled water plant. A high-efficiency VFD-equipped chiller with an appropriate condenser water reset is typically the most efficient cooling option for large facilities. Chiller part-load efficiency should be considered since data centers often operate at less than peak capacity. Chiller part-load efficiencies can be optimized with variable frequency driven compressors, high evaporator temperatures and low entering condenser water temperatures.

Oversized cooling towers with VFD-equipped fans will lower water-cooled chiller plant energy. For a given cooling load, larger towers have a smaller approach to ambient wet bulb temperature, thus allowing for operation at lower cold condenser water temperatures and improving chiller operating efficiency. The larger fans associated with the oversized towers can be operated at lower speeds to lower cooling tower fan energy compared to a smaller tower.

Condenser water and chilled water pumps should be selected for the highest pumping efficiency at typical operating conditions, rather than at full load condition.

Optimize Plant Design and Operation

Data centers offer a number of opportunities in central plant optimization, both in design and operation. A medium-temperature, as opposed to low-temperature, chilled water loop design using a water supply temperature of 55°F or higher improves chiller efficiency and eliminates uncontrolled phantom dehumidification loads (refer to “Humidification” and “Controls” sections). Higher temperature chilled water also allows more water-side economizer hours, in which the cooling towers can serve some or the entire load directly, reducing or eliminating the load on the chillers. The condenser water loop should also be optimized; a 5°F to 7°F approach cooling tower plant with a condenser water temperature reset pairs nicely with a variable speed chiller to offer large energy savings.

Efficient Pumping

A well thought out efficient pumping design is an essential component to a high-efficiency chilled water system. Pumping efficiency can vary widely depending on the configuration of the system, and whether the system is for an existing facility or new construction. Listed below are general guidelines for optimizing pumping efficiency for existing and new facilities of any configuration.

Existing Facilities:

- Reduce the average chilled water flow rate corresponding to the typical load.
- Convert existing primary/secondary chilled water pumping system to primary-only.
- Convert existing system from constant flow to variable flow.
- Reduce the pressure drop of the chilled water distribution system by opening pump balancing valves and allowing pump VFDs to limit the flow rate.
- Reduce the chilled water supply pressure set point.
- Add a chilled water pumping differential pressure set point reset control sequence.
- Eliminate unnecessary bypassed chilled water by replacing 3-way chilled water valves with 2-way valves.

New Construction:

- Reduce the average chilled water flow rate corresponding to the typical load.
- Implement primary-only variable flow chilled water pumping.
- Specify an untrimmed impeller, do not install pump balancing valves and instead use a VFD to limit pump flow rate.
- Design for a low water supply pressure set point.
- Specify a water pumping differential pressure set point reset control sequence.
- Design a low pressure drop pipe layout for pumps.
- Specify 2-way chilled water valves instead of 3-way valves.
- Install VFDs on all pumps and run redundant pumps at lower speeds.

Free Cooling

Air-Side Economizer

The cooling load for a data center is independent of the outdoor air temperature. Most nights and during mild winter conditions, the lowest cost option to cool data centers is an air-side economizer; however, a proper engineering evaluation of the local climate conditions must be completed to evaluate whether this is the case for a specific data center. Due to the humidity and contamination concerns associated with data centers, careful control and design work may be required to ensure that cooling savings are not lost because of excessive humidification and filtration requirements, respectively. Data center professionals are split in the perception of risk when using this strategy. It is standard practice, however, in the telecommunications industry to equip their facilities with air-side economizers. Some IT-based centers routinely use outside air without apparent complications, but others are concerned about contamination and environmental control for the IT equipment in the room. Nevertheless, outside air economizing is implemented in many data center facilities and results in energy-efficient operation. While ASHRAE Standard 90.1 currently does not require economizer use in data centers, a new version of this standard will likely add this requirement. Already some code authorities, such as the Department of Planning and Development for the city of Seattle, have mandated the use of economizers in data centers under certain conditions.

Control strategies to deal with temperature and humidity fluctuations must be considered along with contamination concerns over particulates or gaseous pollutants. For data centers with active humidity control, a dewpoint temperature lockout scheme should be used as part of the air-side economizer control strategy. This scheme prevents high outside air dehumidification and humidification loads by tracking the moisture content of the outside air and locking out the economizer when the air is either too dry or too moist. Mitigation steps may involve filtration or other measures. Other contamination concerns such as salt or corrosive matter should be evaluated. Generally, concern over contamination should be limited to unusually harsh environments such as pulp and paper mills or large chemical spills.

Wherever possible, outside air intakes should be located on the north side of buildings in the northern hemisphere where there is significantly less solar heat gain compared to the south side.

Water-Side Economizer

Free cooling can be provided via a waterside economizer, which uses the evaporative cooling capacity of a cooling tower to produce chilled water to cool the data center during mild outdoor conditions. Free cooling is usually best suited for climates that have wet bulb temperatures lower than 55°F for 3,000 or more hours per year. It most effectively serves chilled water loops designed for 50°F and above chilled water or lower temperature chilled water loops with significant surplus air handler capacity in normal operation. A heat exchanger is typically installed to transfer heat from the chilled water loop to the cooling tower water loop while isolating these loops from each other. Locating the heat exchanger upstream from the chillers, rather than in parallel to them, allows for integration of the water-side economizer as a first stage of cooling the chilled water before it reaches the chillers. During those hours when the water-side economizer can remove enough heat to reach the chilled water supply set point, the chilled water can be bypassed around the chillers. When the water-side economizer can remove heat from the hot chilled water but not enough to reach set point, the chillers operate at reduced load to meet the chilled water supply set point.

Thermal Storage

Thermal storage is a method of storing thermal energy in a reservoir for later use, and is particularly useful in facilities with particularly high cooling loads such as data centers. It can result in peak electrical demand savings and improve chilled water system reliability. In climates with cool, dry nighttime conditions, cooling towers can directly charge a chilled water storage tank; using a small fraction of the energy otherwise required by chillers. A thermal storage tank can also be an economical alternative to additional mechanical cooling capacity; for example water storage provides the additional benefit of backup make-up water for cooling towers.

Direct Liquid Cooling

Direct liquid cooling refers to a number of different cooling approaches that all share the same characteristic of transferring waste heat to a fluid at or very near the point the heat is generated, rather than transferring it to room air and then conditioning the room air. One current approach to implementing liquid cooling utilizes cooling coils installed directly onto the rack to capture and remove waste heat. The under-floor area is often used to run the coolant lines that connect to the rack coil via flexible hoses. Many other approaches are available or being pursued, ranging from water cooling of component heat sinks to bathing components with dielectric fluid cooled via a heat exchanger.

Liquid cooling can serve higher heat densities and be much more efficient than traditional air cooling as water flow is a much more efficient method of transporting heat. Energy efficiencies will be realized when such systems allow the use of a medium temperature chilled water supply (55°F to 60°F rather than 44°F) and by reducing the size and power consumption of fans serving the data center. These warmer chilled water supply temperatures facilitate the pairing of liquid cooling with a water-side economizer, further increasing potential energy savings.

Humidification

Low-energy humidification techniques can replace traditional electric resistance humidifiers with an adiabatic approach that uses the heat present in the air or recovered from the computer heat load for humidification. Ultrasonic humidifiers, evaporative wetted media and micro droplet spray are some examples of adiabatic humidifiers. An electric resistance humidifier requires about 430 Watts to boil one pound of 60°F water, while a typical ultrasonic humidifier only requires 30 Watts to atomize the same pound of water. These passive humidification approaches also cool the air, in contrast to an electric resistance humidifier heating the air, which further saves energy by reducing the load on the cooling system.

Controls

More options are now available for dynamically allocating IT resources as computing or storage demands vary. Within the framework of ensuring continuous availability, a control system should be programmed to maximize the energy efficiency of the cooling systems under variable ambient conditions as well as variable IT loads.

Variable speed drives on CRAH and CRAC units (if available for the latter) allow for varying the airflow as the cooling load fluctuates. For raised floor installations, the fan speed should be controlled to maintain an under-floor pressure set point. However, cooling air delivery via conventional raised floor tiles can be ill-suited for responding to the resulting dynamic heat load without either over-cooling the space or starving some areas of sufficient cooling. Variable air volume air delivery systems are a much better solution for consistently providing cooling when and where it is needed. Supply air and supply chilled water temperatures should be set as high as possible while maintaining the necessary cooling capacity.

Data centers often over-control humidity, which results in no real operational benefits and increases energy use. Tight humidity control is a carryover from old mainframe and tape storage eras and generally can be relaxed or eliminated for many locations.

Humidity controls are frequently not centralized. This can result in adjacent units serving the same space fighting to meet the humidity set point, with one humidifying while the other is dehumidifying. Humidity sensor drift can also contribute to control problems if sensors are not regularly recalibrated. One very important consideration to reducing unnecessary humidification is to operate the cooling coils of the air-handling equipment above the dew point (usually by running chilled water temperatures above 50°F), thus eliminating unnecessary dehumidification.

On the chilled water plant side, variable flow pumping and chillers equipped with variable speed driven compressors should be installed to provide energy-efficient operation during low load conditions. Another option to consider for increasing chiller plant efficiency is to actively reset the chilled water supply temperature higher during low load conditions. In data centers located in relatively dry climates and which experience relatively low partial loads, implementing a water-side economizer can provide tremendous savings over the course of the year (see earlier discussion on "Water-Side Economizers").

Electrical Systems

Similar to cooling systems, it is important to always consider initial and future loads, in particular part- and low-load conditions, when designing and selecting equipment for a data center's electrical system.

Power Distribution

Data centers typically have an electrical power distribution path consisting of the utility service, switchboard, switchgear, alternate power sources (i.e. backup generator), paralleling equipment for redundancy (i.e. multiple UPS's and PDU's), and auxiliary conditioning equipment (i.e. line filters, capacitor bank). These components each have a heat output that is tied directly to the load in the data center. Efficiencies can range widely between manufacturers and variations in how the equipment is designed. However, operating efficiencies can be controlled and optimized through thoughtful selection of these components.

Uninterruptible Power Supplies (UPS)

UPS systems provide backup power to data centers, and can be based on battery banks, rotary machines, fuel cells, or other technologies. A portion of all the power supplied to the UPS to operate the data center equipment is lost to inefficiencies in the system. The first step to minimize these losses is to evaluate which equipment, if not the entire data center, requires a UPS system. For instance the percent of IT power required by a scientific computing facility can be significantly lower than the percent required for a financial institution.

Increasing the UPS system efficiency offers direct, 24-hour-a-day energy savings, both within the UPS itself and indirectly through lower heat loads and even reduced building transformer losses. Among double conversion systems (the most commonly used data center system); UPS efficiency ranges from 86% to 95%. When a full data center equipment load is served through a UPS system, even a small improvement in the efficiency of the system can yield a large annual cost savings. For example, a 15,000 square foot data center with IT equipment operating at 100 W/sf requires 13,140 MWh of energy annually for the IT equipment. If the UPS system supplying that power has its efficiency improved from 90% to 95%, the annual energy bill will be reduced by

768,421 kWh, or about $90,000 at $0.12/kWh, plus significant additional cooling system energy savings from the reduced cooling load. For battery-based UPS systems, use a design approach that keeps the UPS load factor as high as possible. This usually requires using multiple smaller units.

Redundancy in particular requires design attention; operating a single large UPS in parallel with a 100% capacity identical redundant UPS unit (n+1 design redundancy) results in very low load factor operation, at best no more than 50% at full design buildout. Consider a UPS system sized for two UPS units with n+1 redundancy, with both units operating at 30% load factor. If the same load is served by three smaller units (also sized for n+1 redundancy), then these units will operate at 40% load factor. This 10% increase in load factor can result in a 1.2% efficiency increase (see Figure 6). For a 100 kW load, this efficiency increase can result in savings of approximately 13,000 kWh annually.

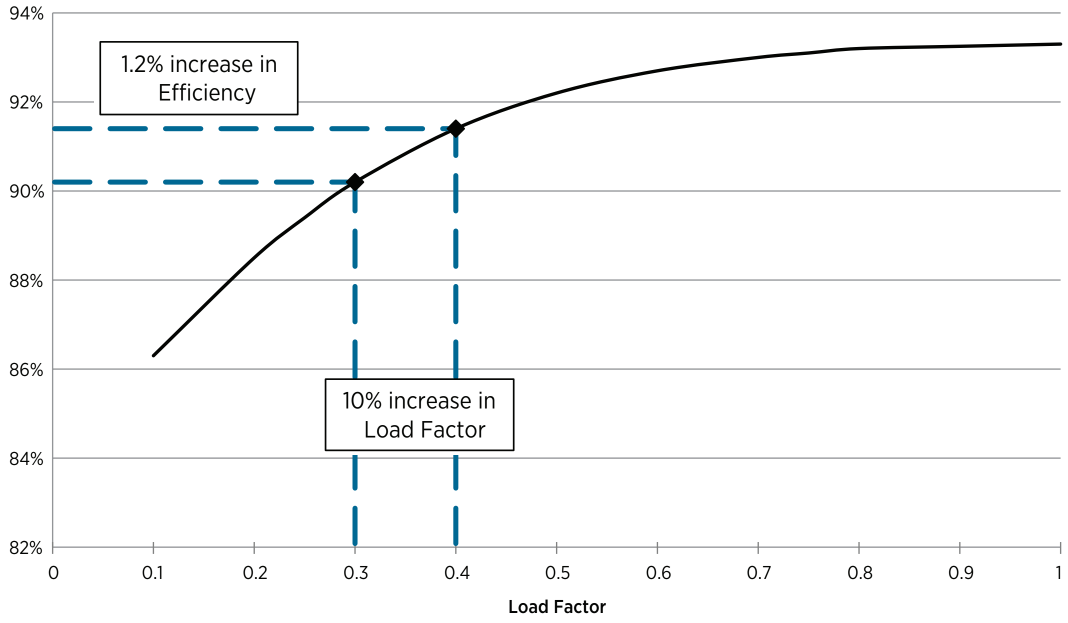

Figure 6. Typical UPS efficiency curve for 100 kVA capacity and greater (Source: Rumsey Engineers)

Evaluate the need for power conditioning. Line interactive systems often provide enough power conditioning for servers at a higher efficiency than typical double conversion UPS systems. Some traditional double conversion UPS systems (which offer the highest degree of power conditioning) have the ability to operate in the more efficient line conditioning mode, usually advertised as ‘economy’ or ‘eco’ mode.

New technologies currently being proven on the market, such as flywheel systems, should also be considered. Such systems eliminate replacement and disposal concerns associated with conventional lead-acid battery based UPS systems, the added costs of special ventilation systems and, often, conditioning systems required to maintain temperature requirements to ensure battery life.

Power Distribution Units (PDU)

A PDU passes conditioned power that is sourced from a UPS or generator to provide reliable power distribution to multiple pieces of equipment. It provides many outlets to power servers, networking equipment and other electronic devices that require conditioned and/or continuous power. Maintaining a higher voltage in the source power lines fed from a UPS or generator allows for a PDU to be located more centrally within a data center. As a result, the conductor lengths from the PDU to the equipment are reduced and less power is lost in the form of heat.

Specialty PDU's that convert higher voltage (208V AC or 480V AC) into lower voltage (120V AC) via a built-in step-down transformer for low voltage equipment are commonly used as well. Transformers lose power in the form of heat when voltage is being converted. The parameters of the transformer in this type of PDU can be specified such that the energy efficiency is optimized. A dry-type transformer with a 176°F temperature rise uses 13% to 21% less energy than a 302°F rise unit. The higher efficiency 176°F temperature rise unit has a first-cost premium; however, the cost is usually recovered in the energy cost savings. In addition, many transformers tend to operate most efficiently when they are loaded in the 20% to 50% range. Selecting a PDU with a transformer at an optimized load factor will reduce the loss of power through the transformer. Energy can also be saved by reducing the number of installed PDU's with built-in transformers.

Distribution Voltage Options

Another source of electrical power loss for both AC and DC distribution is that of the conversions required from going from the original voltage supplied by the utility (usually a medium voltage of around 12 kV AC or more) to that of the voltage at each individual device within the data center (usually a low voltage around 120V AC to 240V AC). Designing a power distribution network that delivers all of the required voltages while minimizing power losses is often a challenging task. Following are general guidelines for delivering electrical power in the most energy-efficient manner possible:

- Minimize the resistance by increasing the cross-sectional area of the distribution path and making it as short as possible.
- Maintain a higher voltage for as long as possible to minimize the current.
- Use switch-mode transistors for power conditioning .
- Locate all voltage regulators close to the load to minimize distribution losses at lower voltages.

Demand Response

Demand response refers to the process by which facility operators voluntarily curb energy use during times of peak demand. Many utility programs offer incentives to business owners that implement this practice on hot summer days or other times when energy demand is high and supply is short. Demand Response programs can be executed by reducing loads through a building management system or switching to backup power generation.

For reducing loads when a demand response event is announced, data center operators can take certain reduction measures such as dimming a third of their lighting or powering off idle office equipment. Through automated network building system solutions this can be a simple, efficient, inexpensive process, and has been known to reduce peak loads during events by over 14% (refer to Cisco's site, Enterprise Automates Utility Demand Response http://www.cisco.com/en/US/prod/collateral/ps6712/ps10447/ps10454/case_study_c36-543499.html).

DC Power

In a conventional data center power is supplied from the grid as AC power and distributed throughout the data center infrastructure as AC power. However, most of the electrical components within the data center, as well as the batteries storing the backup power in the UPS system, require DC power. As a result, the power must go through multiple conversions resulting in power loss and wasted energy.

One way to reduce the number of times power needs to be converted is by utilizing a DC power distribution. This has not yet become a common practice and, therefore, could carry significantly higher first costs, but it has been tested at several facilities. A study done by Lawrence Berkeley National Labs in 2007 compared the benefits of adopting a 380V DC power distribution for a datacom facility to a traditional 480V AC power distribution system. The results showed that the facility using the DC power had a 7% reduction in energy consumption compared to the typical facility with AC power distribution. Other DC distribution systems are available including 575V DC and 48V DC. These systems offer energy savings as well.

Lighting

Data center spaces are not uniformly occupied and, therefore, do not require full illumination during all hours of the year. UPS, battery and switch gear rooms are examples of spaces that are infrequently occupied. Therefore, zone based occupancy sensors throughout a data center can have a significant impact on reducing the lighting electrical use. Careful selection of an efficient lighting layout (e.g. above aisles and not above the server racks), lamps and ballasts will also reduce not only the lighting electrical usage but also the load on the cooling system. The latter leads to secondary energy savings.

Other Opportunities for Energy-Efficient Design

On-Site Generation

The combination of a nearly constant electrical load and the need for a high degree of reliability make large data centers well suited for on-site electric generation. To reduce first costs, on-site generation equipment should replace the backup generator system. It provides both an alternative to grid power and waste heat that can be used to meet nearby heating needs or harvested to cool the data center through absorption or adsorption chiller technologies. In some situations, the surplus and redundant capacity of the on-site generation plant can be operated to sell power back to the grid, offsetting the generation plant capital cost.

Co-generation Plants

Co-generation systems, also known as combined heat and power, involve the use of a heat engine or power station to simultaneously produce electricity and useful heat. In data centers it is very common to see a diesel generator as the source of backup power, which can easily be utilized as a co-generation system. The waste heat produced by the generator can be used to run an absorption chiller which provides cooling to the data center. (Refer to the “Different Uses of Waste Heat” section.) Due to the significant air pollution impact of diesel generators, the number of hours of generator operation can be limited by air quality regulations.

Reduce Standby Losses

Standby generators are typically specified with jacket and oil warmers that use electricity to maintain the system in standby at all times; these heaters use more electricity than the standby generator will ever produce. Careful consideration of redundancy configurations should be followed to minimize the number of standby generators. Using waste heat from the data center can minimize losses by block heaters. Solar panels could be considered as an alternative source for generator block heat. Another potential strategy is to work with generator manufacturers to reduce block heater output when conditions allow.

Use of Waste Heat

Waste heat can be used directly or to supply cooling required by the data center through the use of absorption or adsorption chillers, reducing chilled water plant energy costs by well over 50%. The higher the cooling air or water temperature leaving the server, the greater the opportunity for using waste heat. The direct use of waste heat for low temperature heating applications such as preheating ventilation air for buildings or heating water will provide the greatest energy savings. Heat recovery chillers may also provide an efficient means to recover and reuse heat from data center equipment environments for comfort heating of typical office environments.

Absorbers use low-grade waste heat to thermally compress the chiller vapor in lieu of the mechanical compression used by conventional chillers. Single stage, lithium bromide desiccant based chillers are capable of using the low grade waste heat that can be recovered from common onsite power generation options including micro-turbines, fuel cells, and natural gas reciprocating engines. Although absorption chillers have low coefficient of performance (COP) ratings compared to mechanical chillers, utilizing “free” waste heat from a generating plant to drive them increases the overall system efficiency. Earlier absorption chiller model operations have experienced reliability issues due to lithium bromide crystallization on the absorber walls when entering cooling tower water temperatures were not tightly controlled. Modern controls on newer absorption chiller models, though more complicated, remedy this problem. However, start-up and maintenance of absorption chillers have often been viewed as significantly more involved than that for electric chillers.

A potentially more efficient and more reliable thermally driven technology that has entered the domestic market is the adsorption chiller. An adsorbent is a silica gel desiccant based cooling system that uses waste heat to regenerate the desiccant and cooling towers to dissipate the removed heat. The process is similar to that of an absorption process but simpler and, therefore, more reliable. The silica gel based system uses water as the refrigerant and is able to use lower temperature waste heat than a lithium bromide based absorption chiller. Adsorption chillers include better automatic load matching capabilities for better part-load efficiency compared to absorption chillers. The silica gel adsorbent is non-corrosive and requires significantly less maintenance and monitoring compared to the corrosive lithium bromide absorbent counterpart. Adsorption chillers generally restart more quickly and easily compared to absorption chillers. While adsorption chillers have been in production for more than 20 years, they have only recently been introduced in the U.S. market.

Data Center Metrics and Benchmarking

Energy efficiency metrics and benchmarks can be used to track the performance of and identify potential opportunities to reduce energy use in data centers. For each of the metrics listed in this section, benchmarking values are provided for reference. These values are based on a data center benchmarking study carried out by Lawrence Berkeley National Laboratories. The data from this survey can be found under LBNL's *Self-Benchmarking Guide for Data Centers*: http://hightech.lbl.gov/benchmarking-guides/data.html

Power Usage Effectiveness (PUE) and Data Center Infrastructure Efficiency (DCiE)

PUE is defined as the ratio of the total power to run the data center facility to the total power drawn by all IT equipment:

$$PUE = \frac{Total\ Facility\ Power}{IT\ Equipment\ Power}$$

Standard	Good	Better
2.0	1.4	1.1

An average data center has a PUE of 2.0; however, several recent super-efficient data centers have been known to achieve a PUE as low as 1.1.

DCiE is defined as the ratio of the total power drawn by all IT equipment to the total power to run the data center facility, or the inverse of the PUE:

$$DCIE = \frac{1}{PUE} = \frac{IT\ Equipment\ Power}{Total\ Facility\ Power}$$

Standard	Good	Better
0.5	0.7	0.9

The Green Grid developed benchmarking protocol for these two metrics for which references and URLs are provided at the end of this guide.

It is important to note that these two terms—PUE and DCiE—do not define the overall efficiency of an entire data center, but only the efficiency of the supporting equipment within a data center. These metrics could be alternatively defined using units of average annual power or annual energy (kWh) rather than an instantaneous power draw (kW). Using the annual measurements provides the advantage of accounting for variable free-cooling energy savings as well as the trend for dynamic IT loads due to practices such as IT power management.

PUE and DCiE are defined with respect to site power draw. Another alternative definition could use a source power measurement to account for different fuel source uses.

Energy Star defines a similar metric, defined with respect to source energy, Source PUE as:

$$\text{Source PUE} = \frac{Total\ Facility\ Energy\ (kWh)}{UPS\ Energy\ (kWh)}$$

As mentioned, the above metrics provide a measure of data center infrastructure efficiency in contrast to an overall data center efficiency. Several organizations are working on developing overall data center efficiency metrics with a protocol to account for the useful work produced by a data center per unit of energy or power. Examples of such metrics are the Data Center Productivity and Data Center Energy Productivity metrics proposed by The Green Grid, and the Corporate Average Data Center Efficiency metric proposed by The Uptime Institute.

Rack Cooling Index (RCI) and Return Temperature Index (RTI)

RCI measures how effectively equipment racks are cooled according to equipment intake temperature guidelines established by ASHRAE/NEBS. By using the difference between the allowable and recommended intake temperatures from the ASHRAE Class 1 (2008) guidelines, the maximum (RCI_{HI}) and minimum (RCI_{LO}) limits for the RCI are defined as follows:

$$RCI_{HI} = \left[1 - \frac{\sum_{T_x > 80} (T_x - 80)}{(90 - 80)n} \right] \times 100\ [\%] \qquad RCI_{LO} = \left[1 - \frac{\sum_{T_x < 65} (65 - T_x)}{(65 - 59)n} \right] \times 100\ [\%]$$

where,

T_x = Mean temperature at equipment intake x

n = Total number of intakes

An RCI of 100% represents ideal conditions for the equipment, with no over or under temperatures.
An RCI < 90% is often considered to portray poor conditions.

RTI evaluates the energy performance of the air management system. RTI is defined as:

$$RTI = \frac{\Delta T_{AHU}}{\Delta T_{EQUIP}} \times 100\%$$

where,

ΔT_{AHU} is the typical (airflow weighted) air handler temperature drop

ΔT_{EQUIP} is the typical (airflow weighted) IT equipment temperature rise

Deviations from an RTI of 100% indicate declining performance in the air management system; over 100% suggests recirculation of air which results in sporadic "hot spots" being significantly hotter than the average space temperature thus elevating return air temperatures; less than 100% suggests by-pass of air where the cold air does not contribute to cooling the electronic equipment and returns directly to the air handler thus decreasing the return air temperature. Therefore, an RTI of 100% should be the target goal for an efficient air management system. Since the air temperature rise across IT equipment can range from 10°F to more than 40°F, the equipment delta-T (ΔT) used in the RTI calculation is an airflow weighted average. Measuring a precise temperature rise across all IT equipment in a data center can be a challenging and often impractical task. Suggested methods for measuring and estimating the airflow weighted equipment ΔT are provided in the Air-Management Data Collection Guide and Engineering Reference found at DOE's DC Pro Software Tool Suite: http://www1.eere.energy.gov/industry/datacenters/software.html.

The RCI and RTI parameters allow an objective method of measuring the overall performance of a data center air management system. They should be used in tandem to ensure the best possible design. The supply and return air temperature difference, commonly referred to as the "air-side ΔT" is commonly used as a metric for air management effectiveness. RTI is a better indicator of air management effectiveness because it accounts for the temperature differences at the servers (which can range from 10°F to over 40°F, depending on server loading) and at the air handlers. However, the air-side ΔT can provide additional guidance in terms of how heavily to load a rack. That is, the more densely populated a rack is, the higher the equipment ΔT, and, therefore, one can design for a higher air-side ΔT to realize fan energy savings.

Heating, Ventilation and Air-Conditioning (HVAC) System Effectiveness

This metric is defined as the ratio of the annual IT equipment energy to the annual HVAC system energy:

$$\text{Effectiveness} = \frac{kWh/yr_{\ IT}}{kWh/yr_{\ HVAC}}$$

Standard	Good	Better
0.7	1.4	2.5

For a fixed value of IT equipment energy, a lower HVAC system effectiveness corresponds to a relatively high HVAC system energy use and, therefore, a high potential for improving HVAC system efficiency. Note that a low HVAC system effectiveness may indicate that server systems are far more optimized and efficient compared to the HVAC system. Thus, this metric is a coarse screen for HVAC efficiency potential. According to a database of data centers surveyed by Lawrence Berkeley National Laboratory, HVAC system effectiveness can range from 0.6 up to 3.5.

Airflow Efficiency

This metric characterizes overall airflow efficiency in terms of the total fan power required per unit of airflow. This metric provides an overall measure of how efficiently air is moved through the data center, from the supply to the return, and takes into account low pressure drop design as well as fan system efficiency.

$$\frac{\text{Total Fan Power (W)}}{\text{Total Fan Airflow (cfm)}}$$

Standard	Good	Better
1.25W/cfm	0.75 W/cfm	0.5 kW/cfm

Cooling System Efficiency

There are several metrics that measure the efficiency of an HVAC system. The most common metric used to measure the efficiency of an HVAC system is the ratio of average cooling system power usage (kW) to the average data center cooling load (tons). A cooling system efficiency of 0.8 kW/ton is considered good practice while an efficiency of 0.6 kW/ton is considered a better benchmark value.

$$\frac{\text{Average Cooling System Power (kW)}}{\text{Average Cooling Load (ton)}}$$

Standard	Good	Better
1.1 kW/ton	0.8 kW/ton	0.6 kW/ton

On-Site Monitoring and Continuous Performance Measurement

Ongoing energy-usage management can only be effective if sufficient metering is in place. There are many aspects to monitoring the energy performance of a data center that are necessary to ensure that the facility maintains the high efficiency that was carefully sought out in the design process. Below is a brief treatment of best practices for data center energy monitoring. For more detail, refer to the Self Benchmarking Guide for Data Center Energy Performance in the “Bibliography and Resources” section.

Energy-efficiency benchmarking goals, based on appropriate metrics, first need to be established to determine which measured values need to be obtained for measuring the data center’s efficiency. The metrics listed above provide a good starting point for high-level energy-efficiency assessment. A more detailed assessment could include monitoring to measure losses along the electrical power chain equipment such as transformers, UPS and PDUs with transformers. (For a list of possible measured values, refer to the *Self-Benchmarking Guide for High-Tech Buildings: Data Centers* at Lawrence Berkeley National Laboratory’s Web site: http://hightech.lbl.gov/benchmarking-guides/data.html.)

The accuracy of the monitoring equipment should be specified, including calibration status, to support the level of desired accuracy expected from the monitoring. The measurement range should be carefully considered when determining the minimum sensor accuracy. For example, a pair of +/- 1.5°F temperature sensors provides no value for determining the chilled water ΔT if the operating ΔT can be as low as 5°F. Electromagnetic flow meters and "strap-on" ultrasonic flow meters are among the most accurate water flow meters available. Three phase power meters should be selected to measure true root mean square (RMS) power.

Ideally, the Energy Monitoring and Control System (EMCS) and Supervisory Control and Data Acquisition (SCADA) systems provide all of the sensors and calculations required to determine real-time efficiency measurements. All measured values should be continuously trended and data archived for a minimum of one year to obtain annual energy totals. An open protocol control system allows for adding more sensors after initial installation. IT equipment often includes on-board temperature sensors. A developing technology includes a communications interface which allows the integration of the on-board IT sensors with an EMCS.

Monitoring for performance measurement should include temperature and humidity sensors at the air inlet of IT equipment and at heights prescribed by ASHRAE's Thermal Guidelines for Data Processing Environments, 2009. New technologies are becoming more prevalent to allow a wireless network of sensors to be deployed throughout the IT equipment rack inlets.

Supply air temperature and humidity should be monitored for each CRAC or CRAH unit as well as the dehumidification/humidification status to ensure that integrated control of these units is successful.

Bibliography and Resources

General Bibliography

- *Design Recommendations for High Performance Data Centers.* Rocky Mountain Institute, 2003.
- *High Performance Data Centers – A Design Guidelines Sourcebook.* Pacific Gas and Electric, 2006. http://hightech.lbl.gov/documents/data_centers/06_datacenters-pge.pdf. Accessed December 3, 2009.
- *Best Practices for Datacom Facility Energy Efficiency.* ASHRAE Datacom Series, 2008.
- *Design Considerations for Datacom Equipment Centers.* ASHRAE Datacom Series, 2005.
- U.S. Department of Energy, Energy Efficiency and Renewable Energy Industrial Technologies Program, *Saving Energy in Data Centers.* http://www1.eere.energy.gov/industry/datacenters/index.html. Accessed December 3, 2009.

Resources

IT Systems

- Server System Infrastructure (SSI) Forum. http://www.ssiforum.org. Accessed December 3, 2009.
- Efficiency of Power Supplies in the Active Mode. EPRI. http://www.efficientpowersupplies.org. Accessed December 3, 2009.
- 80 PLUS® Energy-Efficient Technology Solutions. http://www.80plus.org http://www.hightech.lbl.gov. Accessed December 3, 2009.
- *Program Requirements for Computer Servers, Version 1.0.* Energy Star, 2009.
- *Data Processing and Electronic Areas,* Chapter 17. ASHRAE HVAC Applications, 2007.
- The Green Data Center 2.0, Chapter 2, Energy-Efficient Server Technologies, 2009. http://www.processor.com/editorial/article.asp?article=articles%2Fp3008%2F32p08%2F32p08.asp. Accessed December 3, 2009.
- The Green Grid, *Quantitative Efficiency Analysis of Power Distribution Configurations for Data Center.* http://www.thegreengrid.org/en/Global/Content/white-papers/Quantitative-Efficiency-Analysis. Accessed December 3, 2009.
- Juniper Networks. E*nergy Efficiency for Network Equipment: Two Steps Beyond Greenwashing.* http://www.juniper.net/us/en/local/pdf/whitepapers/2000284-en.pdf. Accessed December 3, 2009.
- Energy Star. E*nterprise Server and Data Center Energy Efficiency Initiatives.* http://www.energystar.gov/datacenters. Accessed December 3, 2009.
- Energy Star. E*nterprise Servers for Consumers.* http://www.energystar.gov/index.cfm?c=ent_servers.enterprise_servers. December 3, 2009.

Environmental Conditions

- *Thermal Guidelines for Data Processing Environments,* 2nd Edition, ASHRAE Datacom Series 1, 2009.
- *Thermal Guidelines for Data Processing Environments*, TC9.9 Mission Critical Facilities, ASHRAE, 2004.

Air Management

- *Thermal Guidelines for Data Processing Environments,* TC9.9 Mission Critical Facilities, ASHRAE, 2004.
- *Data Processing and Electronic Areas,* Chapter 17, ASHRAE HVAC Applications, 2003.

Cooling Systems

- *Best Practices Guide for Variable Speed Pumping in Data Centers,* Ernest Orlando, Lawrence Berkeley National Laboratory, 2009.

- *Data Processing and Electronic Areas,* Chapter 17, ASHRAE HVAC Applications, 2003.

- *Variable-Primary-Flow Systems Revisited,* Schwedler P.E., Mick, Trane Engineers Newsletter, Volume 31, No.4, 2002.

- *Thermal Guidelines for Data Processing Environments,* TC9.9 Mission Critical Facilities, ASHRAE, 2004.

- The Green Grid. *Free Cooling Estimated Savings.* http://cooling.thegreengrid.org/namerica/WEB_APP/calc_index.html. Accessed December 3, 2009.

- *Supervisory Controls Strategies and Optimization, Chapter 41,* ASHRAE Applications Handbook, 2003.

- *ARI Standard 550/590- 2003, Water Chilling Packages Using the Vapor Compression Cycle,* Air-Conditioning and Refrigeration Institute, 2003.

- *High-Density Server Cooling,* Engineered Systems, Christopher Kurkjian and Doug McLellan. http://www.esmagazine.com/Articles/Cover_Story/48aac855d6ea8010VgnVCM100000f932a8c0. Accessed December 3, 2009

- Processor. *Throwing Water at the Heat Problem*, Bruce Gain. http://www.processor.com/editorial/article.asp?article=articles%2Fp2736%2F31p36%2F31p36.asp. Accessed December 3, 2009.

- *Liquid Cooling Guidelines for Datacom Equipment Centers,* ASHRAE Datacom Series, 2006.

- IBM Heat eXchanger water cooled rack, https://www-03.ibm.com/systems/x/hardware/options/cooling.html.

- Knurr CoolTherm. http://www.knuerr.com/web/zip-pdf/en/cooltherm/Knuerr-CoolTherm-4-35kW.pdf. Accessed December 3, 2009.

- Fujitsu. *Water-Cooled Primecenter LC Rack*. https://sp.ts.fujitsu.com/dmsp/docs/ds_primecenter_lc.pdf. Accessed December 3, 2009

- *Psychometrics, Chapter 6, ASHRAE HVAC Fundamentals Handbook*, 2005.

- *Best Practices for Data Centers: Lessons Learned from Benchmarking 22 Data Centers, ACEEE Summer Study on Energy Efficiency in Buildings, 2006.*

- Seattle.Gov. *Seattle Nonresidential Energy Code, Chapter 14 Mechanical Systems,* Section 1433. http://www.seattle.gov/DPD/Codes/Energy_Code/Nonresidential/Chapter_14/default.asp#Section1413. Accessed December 3, 2009.

- High-Performance for High-Tech Buildings. *Data Center Energy Benchmarking Case Study*, LBNL and Rumsey Engineers, February 2003, Facility Report 8. http://hightech.lbl.gov/dc-benchmarking-results.html. Accessed December 3, 2009.

- Waste Heat Systems. *Absorption vs. Adsorption Chillers.* http://www.wasteheat.com/Library_files/WHS%20Adsorption%20vs%20Absorption.pdf. Accessed December 3, 2009.

Electrical Systems

- Echelon Corporation. *Automatic Demand Response: Driving Energy Efficiency for Global Climate Change Using LonWorks® Control Networks,* Steve Nguyen. http://www.echelon.com/Solutions/demandresponse/documents/Echelon_DemandResponse.pdf. Accessed December 3, 2009.

- *Data Processing and Electronic Areas, Chapter 17,* ASHRAE HVAC Applications, 2003.

- *High Performance Buildings: Data Centers Uninterruptible Power Supplies,* Ecos Consulting, EPRI Solutions, LBNL. 2005.

- High-Performance Buildings for High-Tech Industries. *Data Centers*. http://datacenters.lbl.gov. Accessed December 3, 2009.

- Dixon, Gregg (2008). *Demand Response for Today's Data Centers,* Focus Magazine: The International Data Center Design and Management Magazine.

Other Opportunities for Energy-efficient Design

- *Data Processing and Electronic Areas, Chapter 17, ASHRAE HVAC Applications,* 2003.

- Data Center Energy Management. Electrical Infrastructure. http://hightech.lbl.gov/DCtraining/strategies/ei.html. Accessed December 3, 2009.

Data Center Metrics and Benchmarking

- *Self Benchmarking Guide for Data Center Energy Performance Version 1.0,* Ernest Orlando, Lawrence Berkeley National Laboratory, 2006. http://hightech.lbl.gov/documents/DATA_CENTERS/self_benchmarking_guide-2.pdf. Accessed December 3, 2009.

- The Green Grid. *The Green Grid Data Center Power Efficiency Metrics: PUE and DCiE, 2007.* http://www.thegreengrid.org/en/Global/Content/white-papers/The-Green-Grid-Data-Center-Power-Efficiency-Metrics-PUE-and-DCiE. Accessed December 3, 2009.

- The Green Grid. Library and Tools.http://www.thegreengrid.org/library-and-tools.aspx?category=MetricsAndMeasurements&range=Entire%20Archive&type=All&lang=en. Accessed December 3, 2009.

- *Airflow and Cooling Performance of Data Centers: Two Performance Metrics, ASHRAE Transactions, Vol. 114, Part 2, 2008.*

- Industrial Technologies Program. Saving Energy in Data Centers. *DC Pro Software Tool Suite.* http://www1.eere.energy.gov/industry/datacenters/software.html. Accessed December 3, 2009.

- *Self-Benchmarking Guide for High-Tech Buildings: Data Centers*, Ernest Orlando Lawrence Berkeley National Laboratory. http://hightech.lbl.gov/benchmarking-guides/data.html. Accessed February 4, 2010.

Federal Energy Management Program

The Department of Energy's Federal Energy Management Program's (FEMP) mission is to facilitate the Federal Government's implementation of sound, cost-effective energy management and investment practices to enhance the nation's energy security and environmental stewardship.
femp.energy.gov

FEMP Resources

FEMP provides assistance through project transaction services, applied technology services, and decisions support services.

U.S. DEPARTMENT OF ENERGY | Energy Efficiency & Renewable Energy

EERE Information Center
1-877-EERE-INF (1-877-337-3463)
eere.energy.gov/informationcenter

Prepared by the National Renewable Energy Laboratory (NREL)
NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency and Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC

DOE/GO-102010-2956 • February 2010

Printed with a renewable-source ink on paper containing at least 50% wastepaper, including 10% post-consumer waste.