Best Practices for Data Centers: Lessons Learned from Benchmarking 22 Data Centers

Steve Greenberg, Evan Mills, and Bill Tschudi, Lawrence Berkeley National Laboratory
Peter Rumsey, Rumsey Engineers
Bruce Myatt, EYP Mission Critical Facilities

ABSTRACT

Over the past few years, the authors benchmarked 22 data center buildings. From this effort, we have determined that data centers can be over 40 times as energy intensive as conventional office buildings. Studying the more efficient of these facilities enabled us to compile a set of "best-practice" technologies for energy efficiency. These best practices include: improved air management, emphasizing control and isolation of hot and cold air streams; right-sizing central plants and ventilation systems to operate efficiently both at inception and as the data center load increases over time; optimized central chiller plants, designed and controlled to maximize overall cooling plant efficiency, central air-handling units, in lieu of distributed units; "free cooling" from either air-side or water-side economizers; alternative humidity control, including elimination of control conflicts and the use of direct evaporative cooling; improved uninterruptible power supplies; high-efficiency computer power supplies; on-site generation combined with special chillers for cooling using the waste heat; direct liquid cooling of racks or computers; and lowering the standby losses of standby generation systems.

Other benchmarking findings include power densities from 5 to nearly 100 Watts per square foot; though lower than originally predicted, these densities are growing. A 5:1 variation in cooling effectiveness index (ratio of cooling power to computer power) was found, as well as large variations in power distribution efficiency and overall center performance (ratio of computer power to total building power). These observed variations indicate the potential of energy savings achievable through the implementation of best practices in the design and operation of data centers.

The Data Center Challenge

The energy used by a typical rack of state-of-the art servers, drawing 20 kilowatts of power at 10 cents per kWh, uses more that $17,000 per year in electricity. Given that data centers can hold hundreds of such racks, they constitute a very energy-intensive building type. Clearly, efforts to improve energy efficiency in data centers can pay big dividends. But, where to start?

"You can't manage what you don't measure" is a mantra often recited by corporate leaders (Tschudi et al 2005). Similarly, the high-availability needs of data center facilities can be efficiently met using this philosophy. This paper presents a selected distillation of the extensive benchmarking in 22 data centers which formed the basis for the development of 10 best-practice guides for design and operation that summarized how better energy performance was obtained in these actual facilities (LBNL 2006).

With typical annual energy costs per square foot 15 times (and in some cases over 40 times) that of typical office buildings, data centers are an important target for energy savings. They operate continuously, which means that their electricity demand is always contributing to peak utility system demands, an important fact given that utility pricing increasingly reflects time-dependent tariffs. Energy-efficiency best practices can realize significant energy and peak-power savings while maintaining or improving reliability, and yield other non-energy benefits.

Data centers ideally are buildings built primarily for computer equipment: to house it and provide power and cooling to it. They can also be embedded in other buildings including high-rise office buildings. In any event, a useful metric to help gauge the energy efficiency of the data center is the computer power consumption index, or the fraction of the total data center power (including computers, power conditioning, HVAC, lighting, and miscellaneous) to that used by the computers (higher is better; a realistic maximum is in the 0.8 to 0.9 range, depending on climate). As shown in Figure 1, this number varies by more than a factor of 2 from the worst centers to the best we've benchmarked to date. Even the best have room for improvement, suggesting that many centers have tremendous opportunities for greater energy efficiency.

Figure 1. Computer Power Consumption Index

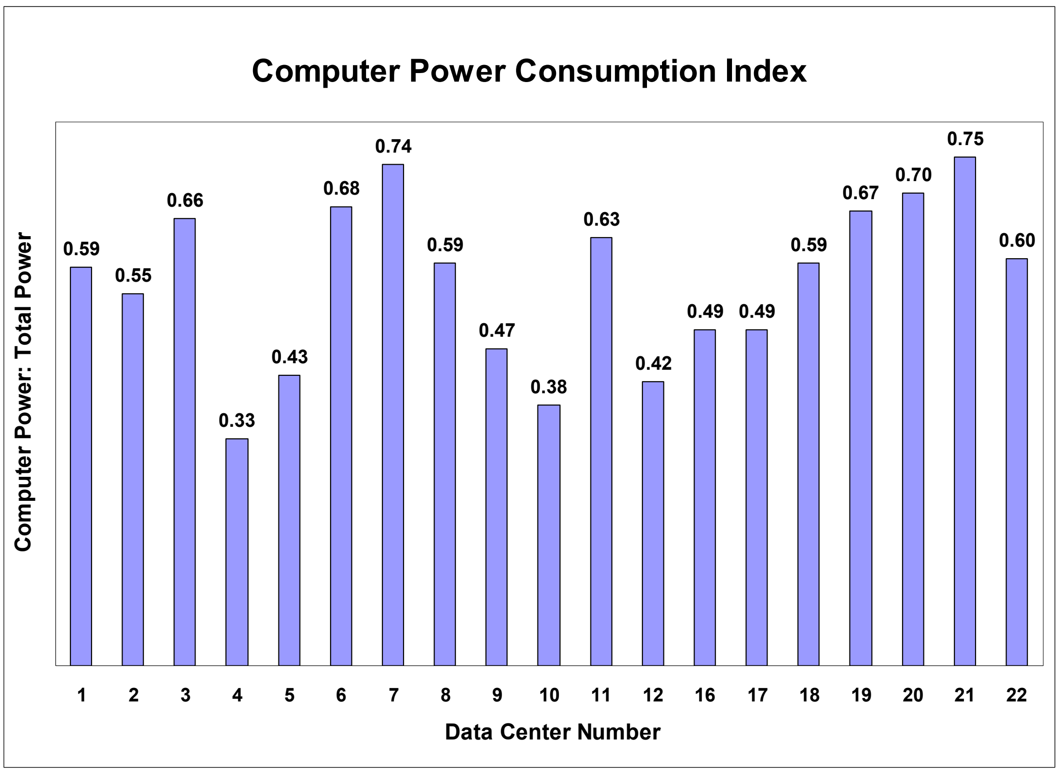

Note: all values are shown as a fraction of the respective data center total power consumption.

The following sections briefly cover data center best practices that have emerged from studying these centers. The references and links offer further details.

Optimize Air Management

As computing power skyrockets, modern data centers are beginning to experience higher concentrated heat loads (ASHRAE 2005a). In facilities of all sizes, from small data centers housed in office buildings to large centers essentially dedicated to computer equipment, effective air distribution has a significant impact on energy efficiency and equipment reliability. (Note that computer equipment is also known as information technology or "IT" equipment.) Energy benchmarking, using a metric that compares energy used for IT equipment to energy expended for HVAC systems (see Figure 2), reveals that some centers perform better than others. For this metric, a higher number indicates that proportionately more of the electrical power is being provided to the computational equipment and less to cooling; in other words, the HVAC system is more effective at removing the heat from the IT equipment. Note that there is a *five-fold* variation from worst to best. This is due a number of factors, including how the cooling is generated and distributed, but air management is a key part of effective and efficient cooling.

Figure 2. HVAC Effectiveness Index

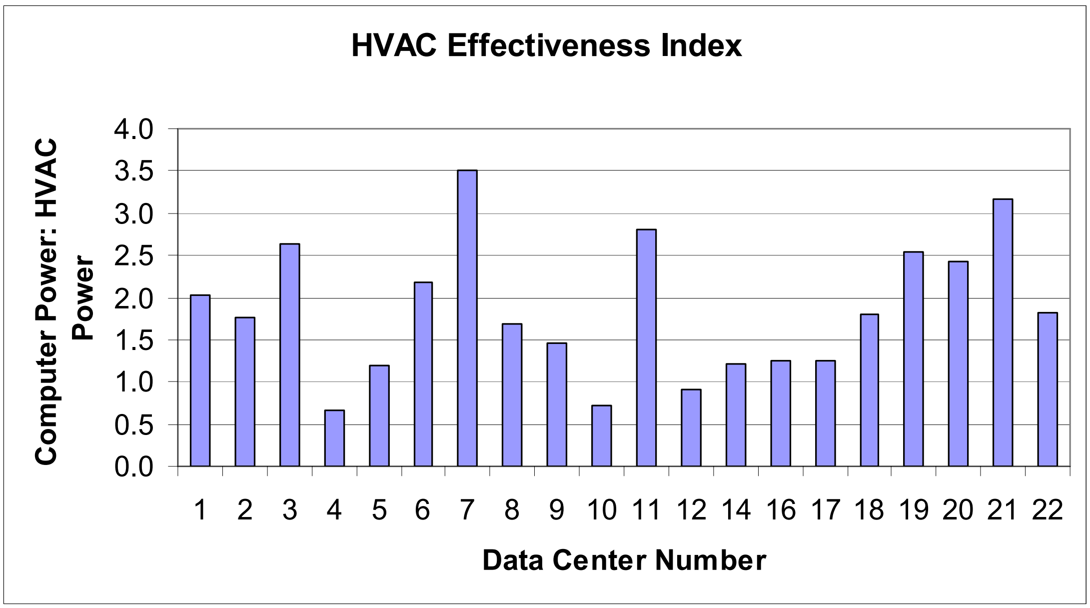

Improving "Air management" - or optimizing the delivery of cool air and the collection of waste heat - can involve many design and operational practices. Air cooling improvements can often be made by addressing:

- short-circuiting of heated air over the top or around server racks
- cooled air short-circuiting back to air conditioning units through openings in raised floors such as cable openings or misplaced floor tiles with openings
- misplacement of raised floor air discharge tiles
- poor location of computer room air conditioning units
- inadequate ceiling height or undersized hot air return plenum
- air blockages such as often happens with piping or large quantities of cabling under raised floors
- openings in racks allowing air bypass (“short-circuit”) from hot areas to cold areas
- poor airflow through racks containing IT equipment due to restrictions in the rack structure
- IT equipment with side or top-air-discharge adjacent to front-to-rear discharge configurations
- Inappropriate under-floor pressurization - either too high or too low

The general goal for achieving better "air management" should be to minimize or eliminate inadvertent mixing between cooling air supplied to the IT equipment and collection of the hot air rejected from the equipment. Air distribution in a well-designed system can reduce operating costs, reduce investment in HVAC equipment, allow for increased utilization, and improve reliability by reducing processing interruptions or equipment degradation due to overheating.

Some of the solutions to common air distribution problems include:

- Use of "hot aisle and cold aisle" arrangements where racks of computers are stacked with the cold inlet sides facing each other and similarly the hot discharge sides facing each other (see Figure 3 for a typical arrangement)

Figure 3. Typical Hot Aisle - Cold Aisle Arrangement

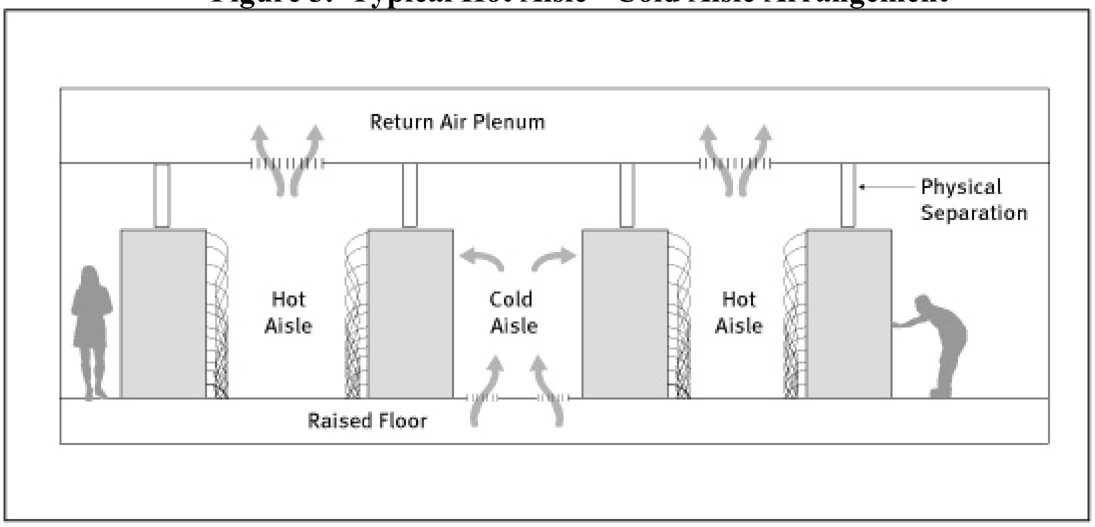

Note: air distribution may be from above or below.

- Sealing cable or other openings in under-floor distribution systems.
- Blanking unused spaces in and between equipment racks.
- Careful placement of computer room air conditioners and floor tile openings, often through the use of computational fluid dynamics (CFD) modeling.
- Collection of heated air through high overhead plenums or ductwork to efficiently return it to the air handler(s).
- Minimizing obstructions to proper airflow

Right-Size the Design

Most data centers are designed based upon vague projections of future power needs. Consequently they are often lightly loaded (compared to their design basis) throughout much, if not all, of their life. Given that projections of IT equipment electrical requirements will remain an inexact science, it is nonetheless important to size electrical and mechanical systems such that they will operate efficiently while overall loading is well below design, yet be scalable to accommodate larger loads should they develop.

Providing redundant capacity and sizing for true peak loads can be done in a manner that improves overall system efficiency, but only if a system design approach is used that considers efficiency among the design parameters. Upsizing duct/plenum and piping infrastructure used to supply cooling to the data center offers significant operating cost and future flexibility benefits. The use of variable-speed motor drives on chillers, pumps for chilled and condenser water, and cooling tower fans, can significantly improve part-load performance over standard chiller-plant design, especially when controlled in a coordinated way to minimize the overall chiller plant energy use. Pursuing efficient design techniques such as medium-temperature cooling loops and waterside free cooling can also mitigate the impact of oversizing by minimizing the use of chillers for much of the year, and reducing their energy use when they do operate.

Cooling tower energy use represents a small portion of plant energy consumption and upsizing towers can provide a significant improvement to chiller performance and waterside free cooling system operation. There is little downside to upsizing cooling towers: a slight increase in cost and footprint are required.

Optimize the Central Plant

Data centers offer a number of opportunities in central plant optimization, both in design and operation. A medium-temperature chilled water loop design using 50-60°F chilled water increases chiller efficiency and eliminates uncontrolled phantom dehumidification loads. The condenser loop should also be optimized; a 5-7°F approach cooling tower plant with a condenser water temperature reset pairs nicely with a variable speed (VFD) chiller to offer large energy savings. A primary-only variable volume pumping system is well matched to modern chiller equipment and offers fewer points of failure, lower first cost, and energy savings relatively to a primary-secondary pumping scheme. Thermal energy storage can be a good option, and is particularly suited for critical facilities where a ready store of cooling can have reliability benefits as well as peak demand savings. Finally, monitoring the efficiency of the chilled water plant is a requirement for optimization and basic energy and load monitoring sensors can quickly pay for themselves in energy savings. If chiller plant efficiency is not independently measured, achieving it is almost as much a matter of luck as design.

Cooling plant optimization consists of implementing several important principles:

- Design for medium temperature chilled water (50-55°F) in order to eliminate uncontrolled dehumidification and reduce plant operating costs. Since the IT cooling load is sensible only, operating the main cooling coils above the dewpoint temperature of the air prevents unwanted dehumidification and also increases the chiller plant efficiency, since the thermodynamic "lift" (temperature difference) between the chilled water and the condenser water is reduced.
- Dehumidification, when required, is best centralized and handled by the ventilation air system, while sensible cooling, the large majority of the load, is served by the medium temperature chilled water system.
- Use aggressive chilled and condenser water temperature resets to maximize plant efficiency. Specify cooling towers for a 5-7°F approach in order to improve chiller economic performance.
- Design hydronic loops to operate chillers near design temperature differential, typically achieved by using a variable flow evaporator design and staging controls.

- Primary-only variable flow pumping systems have fewer single points of failure, have a low first cost (half as many pumps are required), are more efficient, and are more suitable for modern chillers than primary-secondary configurations (Schwedler, 2002).
- Thermal storage can offer peak electrical demand savings and improved chilled water system reliability. Thermal storage can be an economical alternative to additional mechanical cooling capacity.
- 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. The VFD optimizes performance as the load on the compressor varies. While data center space load typically does not change over the course of the day, the load on the compressor does change as the condenser water supply temperature varies. Also, as IT equipment is added or swapped for more-intensive equipment, the load will increase over time.
- For peak efficiency and to allow for preventive maintenance, monitor chiller plant efficiency. Monitoring the performance of the chillers and the overall plant is essential as a commissioning tool as well as for everyday plant optimization and diagnosis.

Design for Efficient Air Handling

Better performance was observed in data center air systems that utilize custom-designed central air handler systems. A centralized system offers several advantages over the traditional multiple-distributed-unit system that evolved as an easy, drop-in computer room cooling appliance. Centralized systems use fewer, larger motors and fans, which are generally more efficient. They are also well-suited for variable-volume operation through the use of variable frequency drives, to take advantage of the fact that the server racks are rarely fully loaded. A centralized air handling system can improve efficiency by taking advantage of surplus and redundant capacity. The maintenance-saving benefits of a central system are well known, and the reduced footprint and maintenance personnel traffic in the data center are additional benefits. Since central units are controlled centrally, they are less likely to “fight” one another (e.g. simultaneous humidification and dehumidification) than distributed units with independent and uncoordinated controls. Another reason that systems using central air handlers tend to be more efficient is that the cooling source is usually a water-cooled chiller plant, which is typically significantly more efficient than the cooling source for water-cooled or air-cooled computer room units. Central air-handlers can also facilitate the use of outside-air economizers, as discussed in the following section.

Capitalize on Free Cooling

Data center IT equipment cooling loads are almost constant during the year. "Free cooling" can be provided through water-side economizers using evaporative cooling (usually provided by cooling towers) to indirectly produce chilled water to cool the datacenter during mild outdoor conditions found in many climates or at night. Free cooling is best suited for climates that have wet bulb temperatures lower than 55°F for 3,000 or more hours per year. Free cooling can improve the efficiency of the chilled water plant by lowering chilled water approach temperatures (i.e. precooling the chilled water before it enters the chiller(s)), or completely eliminate the need for compressor cooling depending upon the outdoor conditions and overall system design. While operating with free cooling, chilled water plant energy consumption can be reduced by up to 75% and there are related improvements in reliability and maintenance through reductions in chiller operation. And since this solution doesn't raise concerns over contamination or humidity control for air entering the IT equipment, it can be an economic alternative to air-side economizers (see below) for retrofitting existing chilled water cooled centers with free cooling.

Air-side economizers can also effectively provide "free" cooling to data centers. Many hours of cooling can be obtained at night and during mild conditions at a very low cost. However this measure is somewhat more controversial. 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. We have observed the use of outside air in several data center facilities, and recognize that outside air results in energy-efficient operation. Some code authorities have mandated the use of economizers in data centers. Currently LBNL is examining the validity of contamination concerns that are slowing its implementation. ASHRAE's data center technical committee, TC 9.9, is also concerned with this issue and plans to develop guidance in the future (ASHRAE 2005b). For now, simply using a standard commercial building economizer should include an engineering evaluation of the local climate and contamination conditions. In many areas, use of outside air may be beneficial, however local risk factors should be known.

Control strategies to deal with temperature and humidity fluctuations must be considered along with contamination concerns over particulates or gaseous pollutants. Mitigation steps may involve filtration or other measures.

Data center architectural design requires that the configuration provide adequate access to the outside if outside air is to be used for cooling. Central air handling units with roof intakes or sidewall louvers are most commonly used, although some internally-located Computer Room Air Conditioning (CRAC) units offer economizer capabilities when provided with appropriate intake ducting.

Improve Humidification Systems and Controls

Data center design often introduces potential inefficiencies when it comes to humidity. Tight humidity control is a carryover from old mainframe and tape storage eras and generally can be relaxed or eliminated for many locations. Data centers attempting tight humidity control were found to be simultaneously humidifying and dehumidifying in many of the benchmarking case studies. This can take the form of either: inadvertent dehumidification at cooling coils that run below the dew point of the air, requiring re-humidification to maintain desired humidity levels; or some computer-room air-handling or air-conditioning units (independently controlled) intentionally operating in dehumidification mode while others humidify. Fortunately, the ASHRAE data center committee addressed this issue by developing guidance for allowable and recommended temperatures and humidity supplied to the inlet of IT equipment (ASHRAE 2004). This document presents much wider recommended and allowable ranges than previous designs required. Since there is very little humidity load from within the center, there is very little need for such control. Some centers handle the issue very effectively by performing humidity control on the make-up air unit only. Humidifiers that use evaporative cooling methods (reducing the temperature of the recirculated air while adding moisture) can be another very effective alternative to conventional steam-generating systems. 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 degrees F), thus eliminating unnecessary dehumidification.

Specify Efficient Uninterruptible Power Supply (UPS) Systems and Information Technology (IT) Equipment Power Supplies

There are two overarching strategies for reducing energy use in any facility: reducing the load, and serving the remaining load efficiently. In the first category, one of the best strategies for improving data center energy efficiency is to reduce the heat loads due to power conversion both within IT equipment and the data center infrastructure. Power conversion from AC to DC and back to AC occurs with battery UPS systems, resulting in large energy losses, which are then compounded by the energy for cooling needed to remove the heat generated by the loss. To make things worse, the efficiency of the power conversion drops off steeply when they are lightly loaded - which is almost always the case due to desires to maintain redundancy and keep each UPS loaded at around 40%. LBNL, along with subcontractors Ecos Consulting and EPRI Solutions, measured UPS efficiencies throughout their operating range and found a wide variation in performance; the delta and flywheel systems showed the highest efficiencies (Figure 4).

Figure 4. Tested UPS Efficiencies

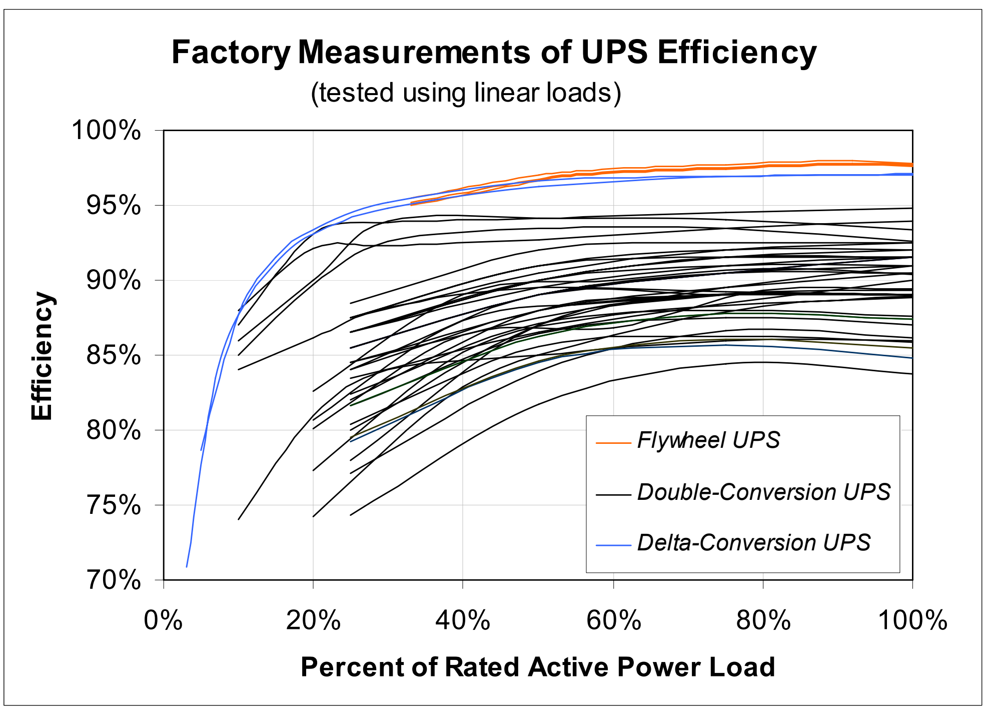

Simply by specifying a more efficient UPS system a 10-15% reduction in the powering of IT equipment is possible along with a corresponding reduction in electrical power for cooling. In other words, an immediate overall reduction of 10-15% in electrical demand for the entire data center is possible. In addition to ongoing energy savings, downsizing HVAC and upstream electrical systems (in new construction), results in reduced capital cost savings, or reserved capacity for future load growth or additional reliability.

A similar phenomenon occurs *within* the IT equipment (e.g. servers), where multiple power conversions typically occur. Conversion from AC to DC, and then multiple DC conversions each contribute to the overall energy loss in IT equipment. LBNL found that a similar wide range of efficiencies exist in power supplies used in servers. By specifying more efficient power supplies, additional savings in both energy and capital costs can be obtained. And any savings at the IT level are magnified by savings in the UPS and the balance of the power distribution system.

Consider On-Site Generation

The combination of a nearly constant electrical load and the need for a high degree of reliability make large datacenters well suited for on-site 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 other near-by heating needs or harvested to cool the datacenter through absorption or adsorption chiller technologies. In some situations, the surplus and redundant capacity of the on-site generation station can be operated to sell power back to the grid, offsetting the generation plant capital cost.

There are several important principles to consider for on-site generation:

- On-site generation, including gas-fired reciprocating engines, micro-turbines, and fuel cells, improves overall efficiency by allowing the capture and use of waste heat.
- Waste heat can be used to supply cooling required by the datacenter through the use of absorption or adsorption chillers, reducing chilled water plant energy costs by well over 50%. In most situations, the use of waste heat is required to make site generation financially attractive. This strategy reduces the overall electrical energy requirements of the mechanical system by eliminating electricity use from the thermal component, leaving only the electricity requirements of the auxiliary pump and motor loads.
- High-reliability generation systems can be sized and designed to be the primary power source while utilizing the grid as a backup, thereby eliminating the need for standby generators and, in some cases, UPS systems. Natural gas used as a fuel can be backed up by propane.
- Where local utilities allow, surplus power can be sold back into the grid to offset the cost of the generating equipment. Currently, the controls and utility coordination required to configure a datacenter-suitable generating system for sellback can be complex but efforts are underway in many localities to simplify the process.

Employ Liquid Cooling of Racks and Computers

Liquid cooling is a far more efficient method of transferring concentrated heat loads than air, due to much higher volumetric specific heats and higher heat transfer coefficients. A high-efficiency approach to cooling datacenter equipment is to transfer the waste heat from racks to a liquid loop at the highest possible temperature. The most common current approach is to use a chilled water coil integrated in some manner into the rack itself. Liquid cooling is adopted for reasons beyond efficiency; it can also serve higher power densities (W/sf). Energy efficiencies are typically realized because such systems allow the use of a medium temperature chilled water loop (50-60°F rather than ~45°F) and by reducing the size and power consumption of fans serving the data center. With some systems and climates, there is the possibility of cooling without the use of mechanical refrigeration (cooling water circulated from a cooling tower would be enough).

In the future, products may become available that allow for direct liquid cooling of chips and equipment heat sinks more directly, via methods ranging from fluid passages in heat sinks to spray cooling with refrigerant (e.g. SprayCool 2006) to submersion in a dielectric fluid. While not currently widely available, such approaches hold promise and should be evaluated as they are commercialized.

There are several key principles in direct liquid cooling:

- Water flow is a very efficient method of transporting heat. On a volume basis, it carries approximately 3,500 times as much heat as air, and moving the water requires an order of magnitude less energy. Water-cooled systems can thus save not just energy but space as well.
- Cooling racks of IT equipment reliably and economically is the main purpose of the data center cooling system; conditioning the remaining space in the data center room without the rack load is a minor task in both difficulty and importance.
- Capturing heat at a high temperature directly from the racks allows for much greater use of waterside economizer free cooling, which can reduce cooling energy use by 60% or more when operating.
- Transferring heat from a small volume of hot air directly off the equipment to a chilled water loop is more efficient than mixing hot air with a large volume of ambient air and removing heat from the entire mixed volume. The water-cooled rack provides the ultimate hot/cold air separation and can run at very high hot-side temperatures without creating uncomfortably hot working conditions for occupants.
- Direct liquid cooling of components offers the greatest cooling system efficiency by eliminating airflow needs entirely. When direct liquid component systems become available, they should be evaluated on a case-by-case basis.

Reduce Standby Loss of Standby Generation

Most data centers have engine-driven generators to back up the UPS system in the event of a utility power failure. Such generators have relatively small but significant and easily addressed standby losses. The biggest standby loss is the engine heater.

The engine heater, used to ensure rapid starting of the engine, typically uses more energy than the generator will ever produce over the life of the facility. These heaters are usually set to excessively high temperatures, based on specifications for emergency generators that have much faster required start-up and loading time (seconds) rather than the minutes allowed for standby generators in data center applications. Reducing setpoints from the typical 120°F to 70°F has shown savings of 79% with no reduction in reliability; delaying the transfer switch transition to the generator greatly reduces engine wear during the warm-up period (Schmidt, 2005).

Improve Design, Operations, and Maintenance Processes

Achieving best practices is not just a matter of substituting better technologies and operational procedures. The broader institutional context of the design and decision-making processes must also be addressed, as follows:

- Create a process for the IT, facilities, and design personnel to make decisions together. It is important for everyone involved to appreciate the issues of the other parties.
- Institute an energy management program, integrated with other functions (risk management, cost control, quality assurance, employee recognition, and training).
- Use life-cycle cost analysis as a decision-making tool, including energy price volatility and non-energy benefits (e.g. reliability, environmental impacts).
- Create design intent documents to help involve all key stakeholders, and keep the team on the same page, while clarifying and preserving the rationale for key design decisions.
- Adopt quantifiable goals based on Best Practices.
- Minimize construction and operating costs by introducing energy optimization at the earliest phases of design.
- Include integrated monitoring, measuring and controls in the facility design.
- Benchmark existing facilities, track performance, and assess opportunities.
- Incorporate a comprehensive commissioning (quality assurance) process for construction and retrofit projects.
- Include periodic “re-commissioning” in the overall facility maintenance program.
- Ensure that all facility operations staff receives site-specific training on identification and proper operation of energy-efficiency features.

Summary

Through the study of 22 data centers, the following best practices have emerged:

- Improved air management, emphasizing control and isolation of hot and cold air streams.
- Right-sized central plants and ventilation systems to operate efficiently both at inception and as the data center load increases over time.
- Optimized central chiller plants, designed and controlled to maximize overall cooling plant efficiency, including the chillers, pumps, and towers.
- Central air-handling units with high fan efficiency, in lieu of distributed units.
- Air-side or water-side economizers, operating in series with, or in lieu of, compressor-based cooling, to provide “free cooling” when ambient conditions allow.
- Alternative humidity control, including elimination of simultaneous humidification and dehumidification, and the use of direct evaporative cooling.
- Improved configuration and operation of uninterruptible power supplies.
- High-efficiency computer power supplies to reduce load at the racks.

- On-site generation combined with adsorption or absorption chillers for cooling using the waste heat, ideally with grid interconnection to allow power sales to the utility.
- Direct liquid cooling of racks or computers, for energy and space savings.
- Reduced standby losses of standby generation systems.
- Processes for designing, operating, and maintaining data centers that result in more functional, reliable, and energy-efficient data centers throughout their life cycle.

Implementing these best practices for new centers and as retrofits can be very cost-effective due to the high energy intensity of data centers. Making information available to design, facilities, and IT personnel is central to this process. To this end, LBNL has recently developed a self-paced on-line training resource, which includes further detail on how to implement best practices and tools that can assist data center operators and service providers in capturing the energy savings potential. See http://hightech.lbl.gov/DCTraining/top.html

References

ASHRAE 2004. *Thermal Guidelines for Data Processing Environments*. TC 9.9 Mission Critical Facilities. Atlanta, Geor.: American Society of Heating, Refrigerating, and Air-Conditioning Engineers.

ASHRAE 2005a. *Datacom Equipment Power Trends and Cooling Applications*. TC 9.9.

ASHRAE 2005b. *Design Considerations for Datacom Equipment Centers*. TC 9.9.

LBNL 2006. http://hightech.lbl.gov/datacenters.html. "High-Performance Buildings for High-Tech Industries, Data Centers". Berkeley, Calif.: Lawrence Berkeley National Laboratory.

RMI 2003. *Design Recommendations for High-Performance Data Centers* – Report of the Integrated Design Charrette. Snowmass, Colo.: Rocky Mountain Institute.

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

Schmitt, Curtis. (Emcor Group). 2004. Personal communication. June 15.

SprayCool 2006. "Spray cooling technology for computer chips". http://www.spraycool.com/ Liberty Lake, Wash.: Isothermal Systems Research.

Tschudi, W., P. Rumsey, E. Mills, T. Xu. 2005. "Measuring and Managing Energy Use in Cleanrooms." *HPAC Engineering*, December.