Typical and Emerging Technologies for Power Flow and Heat Transfer in Data Centers

November 2011

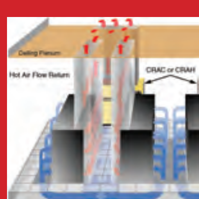

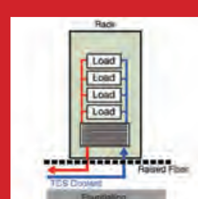

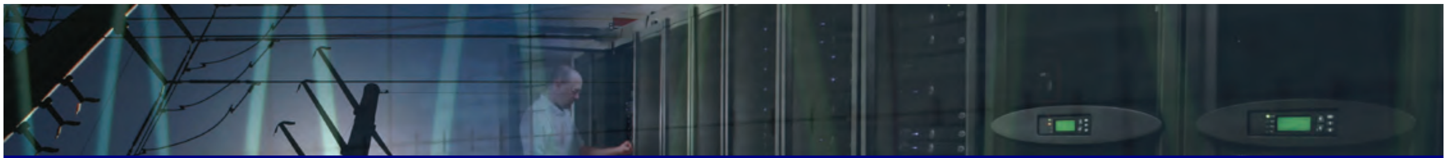

Introduction

A data center is a large concentrated load about 20 times more energy-intensive than the typical commercial building and 80 to 100 times more energy-intensive than the typical home. They operate continuously 24 hours a day, 365 days a year. Paradoxically, data centers are both the result of America's shift to a service economy *and* one of its most powerful drivers. Because information technology (IT) equipment is vital to modern, seamless communication, reliability naturally tends to be the most important criterion for the selection of IT equipment, with first-cost typically next in importance, followed by operating and energy costs. The power density of IT equipment in data centers has been increasing, causing the energy requirements for cooling and other infrastructure to increase.

Between 2000 and 2006, the number of servers being installed in the United States increased from 5.5 million per year to 10.9 million per year. Energy use in data centers was estimated to be 1.5% of the U.S. total electricity use in 2006[1] and was expected to double in five years. Power failures—as well as limitations on space, power, and cooling capacity—were expected to affect almost all data centers in that time period.[2]

Coupled with inefficient air conditioning and other components, a typical data center consumes about two to three times more power than its IT loads.[3] Over half of the power consumed by power systems and cooling systems is fixed (it does not vary with the IT load). Power and cooling systems that can operate at or near capacity and can adjust capacity to meet loads will reduce the effects of fixed losses and transients and increase efficiency.[4] A benchmarking study of 22 data centers by LBNL[5] revealed that energy consumed by a typical data center was distributed across the load classes shown in Table 1.

Table 1 – Data Center Energy Use by Equipment Type (from LBNL)

Application	**Fraction of Facility Power**
Servers	46%
HVAC Cooling	23%
HVAC Fans	8%
UPS	8%
Lighting	4%
Other	11%

Researchers around the world have recently engaged in efforts to identify and develop methods to measure and reduce the power consumed by data centers. This paper highlights existing and emerging opportunities for reducing losses and improving the efficiency and reliability of power flow and heat transfer in data centers.

Data Center IT Equipment and Issues

The electricity supplied to data centers is used to drive the IT equipment and the supporting infrastructure equipment (principally cooling equipment). IT equipment in data centers consist of servers, storage devices, and network equipment. Servers account for a large fraction of the energy consumed by data centers. The modern server, which is energy-intensive and generates abundant heat, continues to drive advances in infrastructure cooling and electric distribution equipment.

Data centers have evolved over the years from one or a few single-corded devices running on 120-volt power sources, to a mix of 120- and 208-volt single- and dual-corded equipment, to rack servers with a high number of 1U and 2U servers running on 208 volt, to high-density data centers with blade servers and techniques that concentrate the power usage (through virtualization, unified networks, or consolidation, as explained later in this section).[6]

It is most efficient to operate servers at high utilization rates because no-load losses represent a large fraction of their power requirements. Virtualization (sometimes called *consolidation*) enables multiple virtual machines to operate on a single physical machine, by sharing the resources of that single computer for multiple applications. Different virtual machines can run different operating systems and multiple applications on the same physical computer. This ensures that each computer is utilized to the fullest extent possible, improving energy

Table of Contents

This white paper was prepared by Mort Blatt, Contractor; Satish Rajagopolan, Dennis Symanski, and Brian Fortenbery of EPRI

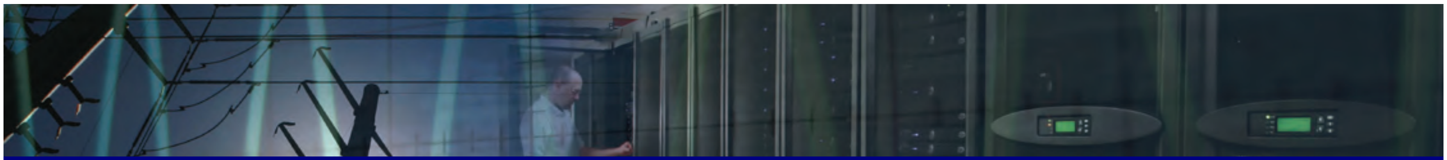

efficiency. (Processors themselves may account for up to 30% of data center energy use. Loading one server at 80% rather than four servers each at 20% clearly provides major energy savings. Because energy is not required for the rotating hard drives, mother board, input/output slots, memory, fans, and power supplies for the three servers, they can be turned off.)[7] On the other hand, consolidating four applications on one server can create reliability issues that need to be addressed (perhaps with some redundancy) in case that server fails.

Each virtual machine contains a complete operating system. In order to dynamically allocate hardware resources, at least one provider of the virtualization software inserts a layer of software containing a supervisor for the virtual machine directly on the computer hardware or on a host operating system. Multiple operating systems and applications run simultaneously on a single physical computer. Because the CPU, memory, operating system, and network devices are encapsulated on a single machine, a virtual machine should be completely compatible with all standard operating systems, applications, and device drivers.[8]

Utilization can also be improved through matrix switching, where equipment is dynamically reallocated to meet requirements, ensuring that equipment is available to meet computing loads and is utilized to the greatest extent possible.[9] Interconnections and duplication of capabilities need to be considered to provide redundancy and associated improvements in availability.

Multiple-core processors can perform far more computations than a single-core processor using the same space and energy. (Dual-core, quad-core, and eight-core processors are available and are more compact, use less power, and produce less heat than single-core processors.)[10] A small reduction in clock frequency causes a small reduction in the amount of work performed but also results in a large reduction in the energy consumed.[11]

Power management should be employed whenever possible to minimize power usage when equipment is idle. One power-management approach simulated energy savings of 26% for the IT equipment over an eight-week period.[12] Power-management software is often not used because of concern for compromising the timeliness of server response that could adversely affect the operations of data centers.[7]

Cloud computing permits users to access pooled resources maintained by a third party and to take advantage of the associated economies of scale. These economies need to be balanced against possible security issues associated with use of these pooled resources. Cloud computing has facilitated functions such as remotely hosted services (applications that run outside corporate data centers), grid computing (deploying computers to simultaneously work on a task), utility computing (permitting users to pay for computing power as they use it), and software as a service (providing software services over the internet).[13] The **MAID** (massive array of idle discs) system saves energy by idling discs and allowing them to be spun up when needed.

Electricity Distribution to a Data Center

Figure 1 shows the elements of power conversion from the utility grid to IT equipment.[14] The uninterruptible power supply (UPS)

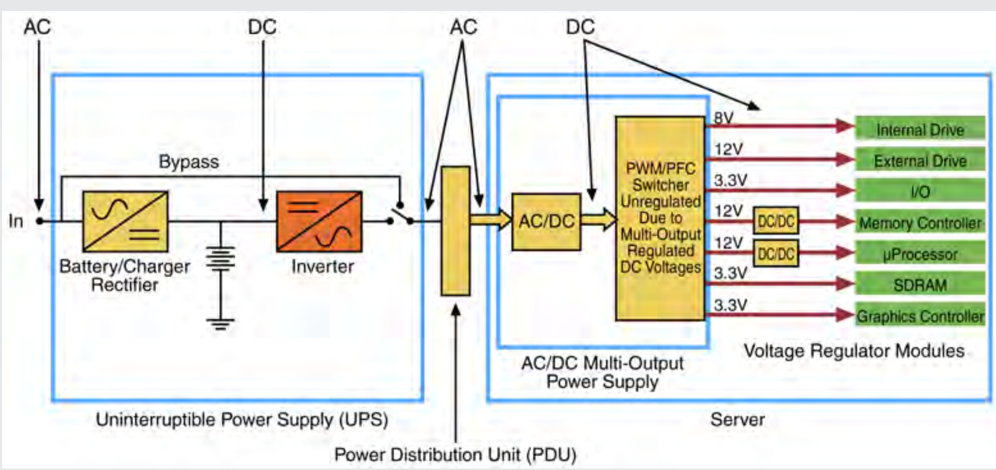

Figure 1 – Steps in Delivering and Conditioning Grid Power to IT Servers

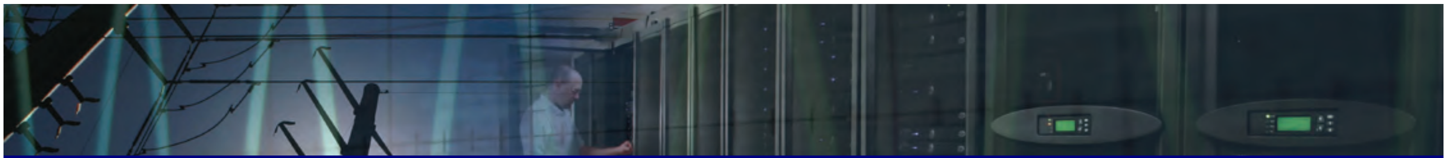

protects the system from disturbances in the supply of electricity. The power distribution unit (PDU) transforms the power to the appropriate voltage. The power supply converts the power to DC, improves the power factor, and regulates the voltage to the level required for the IT equipment.

The power delivered to the data center is converted to DC power so that the batteries in the UPS can be charged and discharged as needed. The power is then converted back to AC and fed to the power supply. Voltage regulators then condition the power for delivery as low-voltage DC to the IT equipment. There are inefficiencies accompanying each of these steps and with the associated equipment that result in losses of more than 20 to 40% in delivering power from the utility to the IT equipment. Using the best available power supplies, processors, and UPS equipment can save almost 30% of the power required by the data center.[15] Each of these steps is discussed in the following paragraphs.

The UPS is designed to provide reliable power to the IT equipment. Batteries are typically used as backup power that can be deployed in the event of a disruption in electric service. UPS efficiency varies with load factor, and fixed losses adversely affect efficiency for load factors below 0.4. Figure 2 (on the following page) shows a collection of UPS efficiency curves as a function of load supplied by manufacturers.[14] Over the past five years, manufacturers have improved their efficiencies, so new test data is needed, but efforts to consolidate loads, adding smaller UPS modules to match the load, or possibly multiplexing the loads should keep efficiencies on the flat part of the curve in between 85 and 95%.

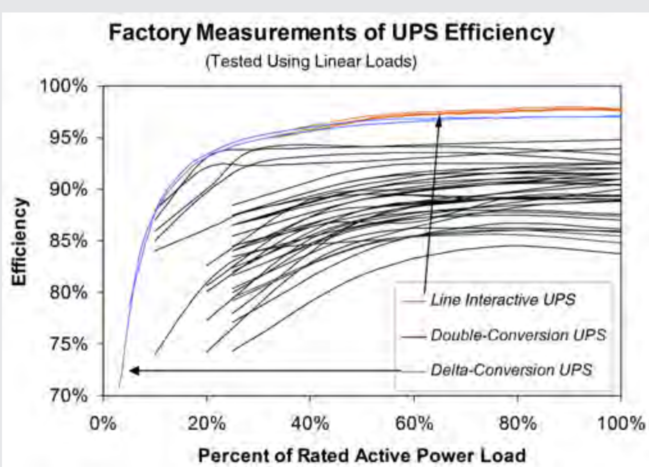

Figure 2 – Typical UPS Performance vs. Load Factor

PDUs transform the power to the desired voltage while balancing the load across the three AC phases and providing protection for each server.[16] Existing PDUs typically have an average efficiency of about 95%,[17] with efficiencies varying from 93 to 99%. The standard power distribution for large data centers in North America is 277/480-volt three-phase power, which supplies PDUs, which in turn convert the voltage to 120/208 volts.

Power supply units (PSUs) convert the 120/208-volt AC to low-voltage DC, which is adjusted to the specific voltage requirement of each DC device (such as the CPU and other chips) by voltage-regulator modules. This DC voltage (12 V or 48 V) is then distributed across a back plane to other voltage regulators, which produce lower voltages for the internal devices, CPU, and memory.[16] Power supplies can also provide power-factor correction to reduce harmonics and thus reduce current requirements and wiring costs. Electrolytic capacitors are provided to permit ride-through of at least one power cycle during an outage. Historically, most power supplies have been much less efficient at 10 to 30% of load (around 60%) compared to their efficiency at 50 to 75% of full load (of around 75%). Today, the efficiencies of power supplies have improved to over 80% for many manufacturers, and over 90% for a few.

Improvements in Power Supply Efficiency

In early 2002, California Energy Commission Public Interest Energy Research (PIER) (http://www.energy.ca.gov/) funded a project to study the active-mode efficiency of desktop computer power supplies. EPRI laboratories performed the efficiency measurement of 50 to 100 randomly chosen computer power supplies bought off the shelf. To its surprise, EPRI discovered that such power supplies were only about 50 to 60% efficient. This was a surprising result for the California Energy Commission (CEC) and other agencies involved in improvement of energy efficiency. This laboratory result led to funding from the CEC to develop a procedure for measuring the efficiency of internal computer power supplies. In 2003, EPRI partnered with Ecos Consulting (www.ecosconsulting.com), a policy-development and research firm, for this effort and developed a general protocol for measuring the efficiency of internal power supplies, which now serves as a *de facto* standard in the power-supply community for efficiency measurements. It is important to note that without this universally accepted test procedure, the industry would be left with no organized method for measurement and verification

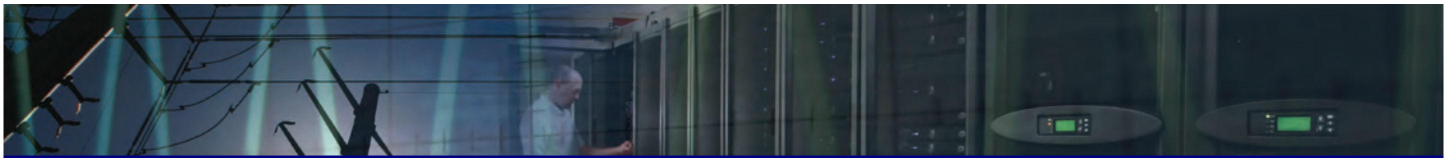

of efficiency. As a result of this work, Ecos launched a program called 80 PLUS, which manages a utility-based incentive pool to help manufacturers meet efficiencies of 80% or more at 20, 50, and 100% of rated load.

In 2005, the U.S. Environmental Protection Agency's ENERGY STAR program funded EPRI and Ecos to measure the efficiencies of the computer power supplies using the latest power-supply efficiency test protocol (available at www.efficientpowersupplies.org/methods.asp) to determine the level with which the power supply industry had responded with higher-efficiency power supplies. The study showed that the market had begun to respond, with a handful of designs that met the 80 PLUS criteria. By the middle of 2006, there were 29 manufacturers offering 123 designs that complied with 80 PLUS. In 2007, ENERGY STAR adopted the 80 PLUS specification, and today, there are over 1200 compliant models on the market. Figures 3 and 4 show the improvement trends for desktop and server power supplies under these programs.

Ecos has worked diligently with a new entity called the *Climate Savers Computing Initiative* (CSCI), which is a consortium of end users and vendors. 80 PLUS and CSCI have aligned their specifications for computers and for servers, with three levels of compliance: bronze, silver, and gold. ENERGY STAR will use this work to create its next specification.

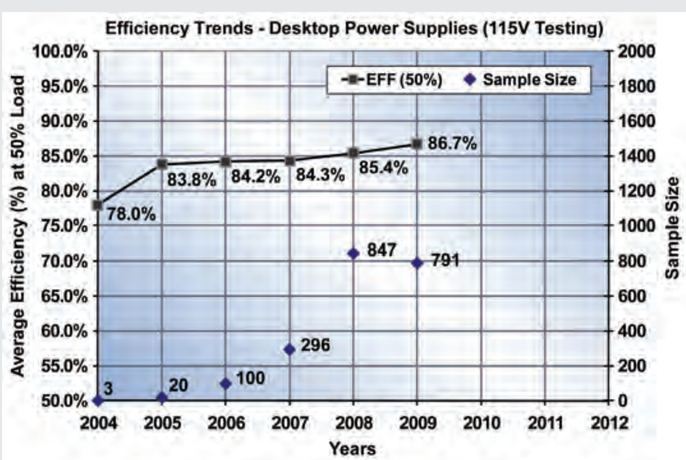

Figure 3 – Desktop Power Supply Efficiency Trend

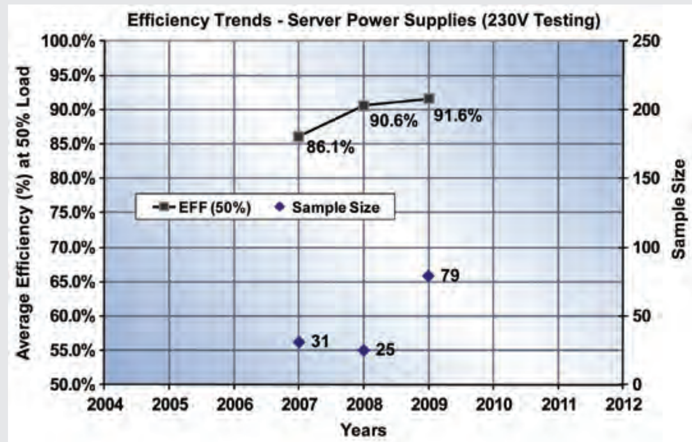

Figure 4 – Server Power Supply Efficiency Trend

DC Versus AC Power

If DC power is delivered by the utility, then steps involving conversion from AC to DC and from DC to AC can be eliminated (see Figure 5 on the following page). In a 2006 joint project, Lawrence Berkeley Labs, Ecos Consulting, and EPRI constructed and tested a proof-of-concept DC system. This system was compared to an AC reference system consisting of the most efficient conversion components available. Test results show that the system achieved energy savings of about 7% compared to the reference AC system and saved 28% compared to a typical current AC system (61% efficient).[18]

Studies by Intel, Emerson, and EYP have shown that benefits that accrue to DC distribution are not limited to efficiency and include:

- No phase balancing, resulting in reduced power strip and wiring complexity.
- No synchronization required to parallel multiple sources.
- No harmonics and therefore no power-factor-correction circuits.
- Fewer breakers are required because of fewer stages.
- Simplified wiring, because only two wires are required.
- No need for complex interlocks, permitting simpler procedures.
- No losses due to distortion.

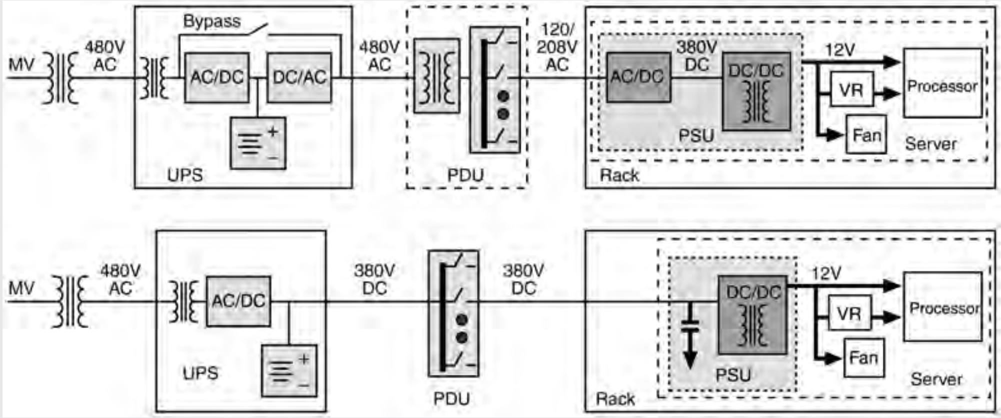

Figure 5 – Typical Data Center Configuration and Efficiencies Achieved in 208-V_{AC} System (Top) vs. 400-V_{DC} System (Source: Delta Products Corp.)

Additional benefits attributable to higher voltage include:

- Simplified wiring, especially at higher power densities.
- Lower currents than at 48 V_{DC} (which is commonly used in telecom central offices). Therefore, smaller wires are used, the use of natural resources is reduced, and less energy is required to process materials.
- Simpler/more efficient connection to renewable energy sources.

In addition, those organizations found about a 7% reduction in energy consumption, twice the reliability, a 33% smaller footprint, and a 15% reduction in capital expense. Other investigators, while acknowledging the efficiency advantages of high-voltage DC power, suggest that these advantages are minor[19] and can be outweighed by factors such as compatibility and availability of IT and cooling equipment.[20] However, there is a burgeoning interest in moving toward DC distribution for all the reasons listed above.

Cooling Data Centers

The classic configuration of a data center with cooling loads of 25 to 40 watts/ft^2 evolved to an approach with the rows facing in the same direction to impart a sense of organization and the computer room air conditioners (CRACs) placed at the ends of the rows. The floor had no perforated tiles. Cold air was directed to the bottom of the IT equipment and discharged out the top. With the introduction of rack-mounted servers, power levels rose, and cabinet orientation became a problem because air that entered in the front of the equipment was discharged out the rear and into the front of the next row.

Because of those shortcomings, the "hot aisle/cold aisle" concept was developed in the 1990s. In this arrangement, CRACs were still located at the perimeter, but in the cold aisles, the cool air was delivered from under the floor via perforated tiles, which required the floor to be raised 12 to 30 inches.[21,22]

Figure 6, on the following page, shows the ideal operation of the hot aisle/cold aisle concept, with the cold air being delivered in cold aisles to the inlets of the racks and the hot air flowing in hot aisles from the racks to the inlet of the CRAC or computer room air handler (CRAH).

Figure 7, on the following page, shows how the hot aisle/cold aisle concept might actually work with some of the hot air short-circuiting back to the rack and creating hot spots near the top of the rack. One way to avoid this short-circuiting is to use fans with variable-speed drives to increase the cold aisle flow to keep the hot air out of the cold aisles. This cooling method is still used in many data centers with deeper raised floors (2 to 4 feet are common), but this approach is cost-effective only up to a certain level of power dissipation (8 to 12 kW per rack or a rack density of 120 to 150 watts/ft^2).[23] Today, a typical server rack can consume about 10 to 20 kW, and typical blade servers require as much as 20 to 30 kW per rack. To simplify cooling, some data centers have spread the loads across half-empty racks to limit power and corresponding cooling requirement to 150 to 200 watts per ft^2 of floor space.[24] This is often referred to as "white space."

As heat loads have increased, techniques for removing heat have improved. These techniques include placing cooling devices closer to the heat load (reducing fan power), concentrating the cooling

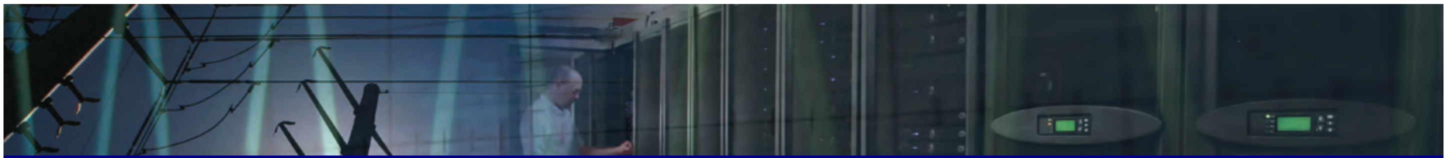

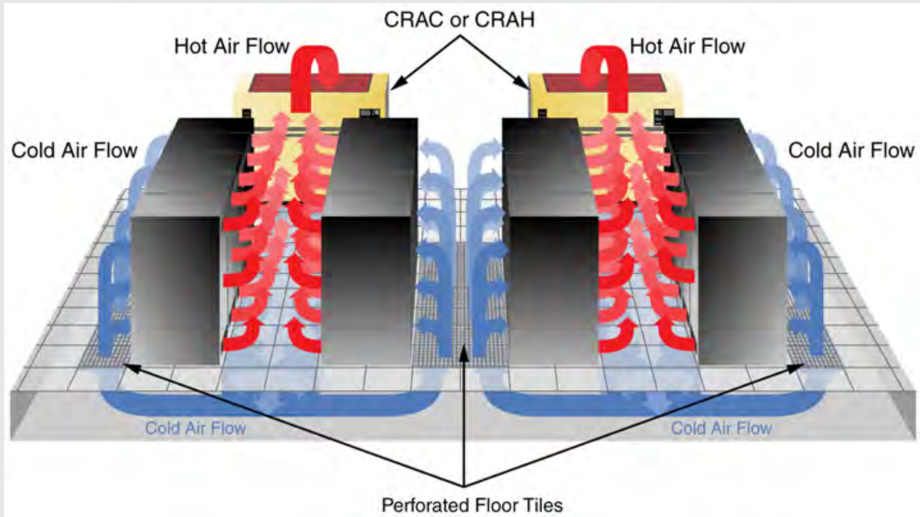

Figure 6 – Example of Ideal Hot Aisle/Cold Aisle Orientation

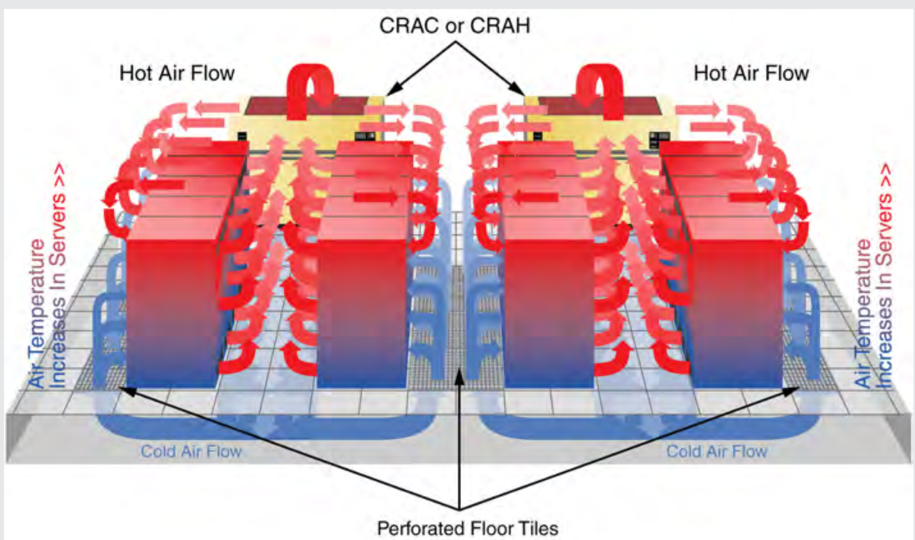

Figure 7 – Example of Typical Hot Aisle/Cold Aisle Orientation and Possible Issues

on the most temperature-sensitive portions of the IT equipment (minimizing overcooling), and reducing mixing of the cold inlet air and hot exhaust air. The latter increases the efficiency of the cooling system by maximizing the temperature difference between the inlet air and exhaust air. Other techniques that can be employed are using close-coupled, in-row, overhead, enclosed systems (hot-aisle or cold-aisle containment), and direct cooling of server components.

Another approach is to use hot aisle or cold aisle containment. Aisle containment constrains the air to the cold aisle or the hot aisle, minimizing air short-circuiting, recirculation, and mixing. This can significantly increase the temperature difference between the air leaving the cooling equipment and the air entering the cooling equipment, thereby improving the efficiency of the cooling system. (Mixing and short-cycling were found to be a major cause of thermal problems in data centers, resulting in an average of two to three times more cooling than was required.)[25,26] Cold aisle containment as shown in Figure 8 on the following page, has the advantage of being able to use a free return for the exhaust air to be drawn into the cool-

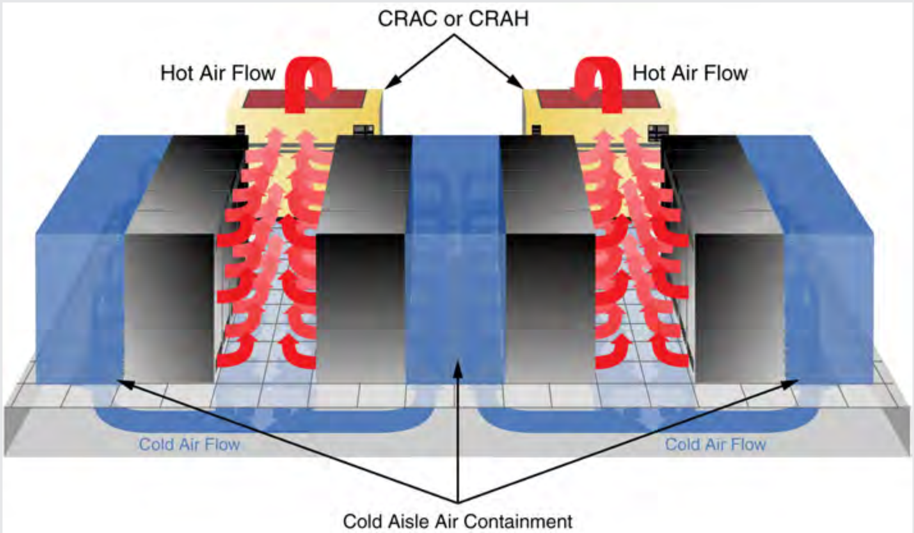

Figure 8 – Schematic of Cold Aisle Containment Configuration

ing equipment, but it keeps most of the data center uncomfortably hot for personnel working in this environment (80°F to as much as 100°F).[27] If equipment other than servers is contained in the data center, then a separate means must be provided for their cooling.

Hot aisle containment provides a comfortable working environment (~70°F) outside the hot aisles, but this approach requires a chimney or other ducting to ensure that the hot air from the contained aisles is not directed back to the inlet of the cooling equipment. More cold air is available for emergency cooling during a disruption of the cooling system, extending operating time in this mode to approximately seven times more than with cold aisle containment (to perhaps a few minutes rather than just seconds).[28] This approach is shown in Figure 9.[4]

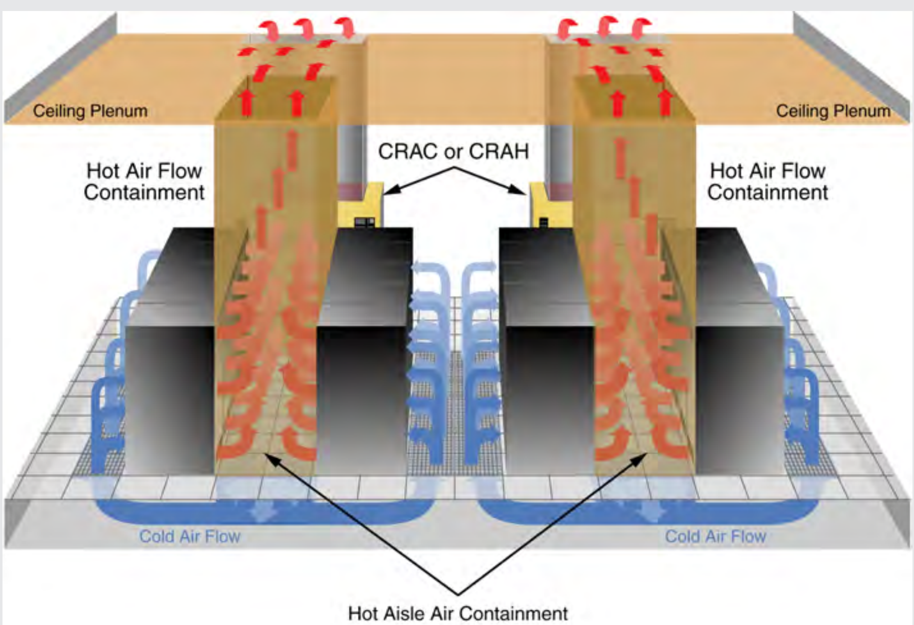

Figure 9 – Schematic of Hot Aisle Containment System

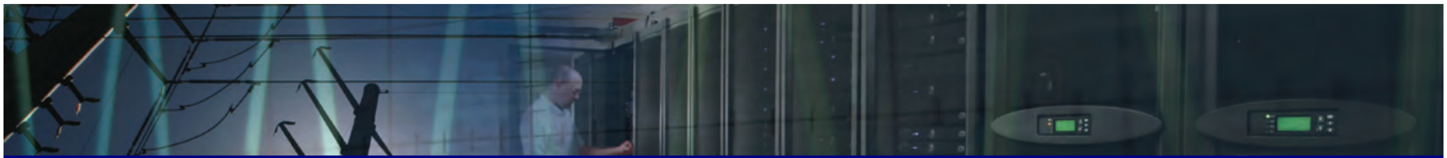

Rack containment as shown in Figure 10 is a refinement of hot aisle containment, with ducted exhaust-rack cabinets to draw air through the cabinet to keep operating temperatures within acceptable limits. Managing the exhaust air-flow completely segregates the hot exhaust air by directing it up an exhaust duct at the top rear of the cabinet. This approach allows perforated tiles to be used anywhere in the room. Heat loads of 17 to 25 kW have been handled with this system. Systems of this type are being designed to handle 30 kW per cabinet.[29]

With row-based cooling, the cooling equipment is placed in spaces between racks or above the racks, minimizing the distance between the load and the cooling source and reducing fan energy compared to the room-based cooling approach. Row-based cooling does not require raised floors. With rack-based cooling, the distance between the cooling equipment and the load is reduced further by putting the cooling system inside the rack, closer to the heat load, enabling up to 50 kW to be removed from the rack.[30] The principal disadvantage of this approach is the cost and large number of air-conditioning devices and associated piping compared to the room and row approaches. CRAC or CRAH equipment can be located on the floor of the data center or above the ceiling. Floor-mounted units seem to be typical for larger data centers, with ceiling-mounted units sometimes used in computer rooms and small data centers.

Cooling Equipment

Currently available cooling systems designed for data centers range from air-cooled systems using CRACs and CRAHs, to systems where liquid (refrigerant or water) is pumped to the racks and used to cool the air in the racks, to direct-liquid-cooled systems with refrigerant or dielectrics injected onto cold plates attached directly to the server electronics. Systems that use a vapor-compression cycle can be included as part of the liquid-delivery configurations. There is a large array of possible permutations and combinations of cooling configurations. These include aisle containment, as mentioned above, in-row positioning of the CRAHs, overhead or top-mounted CRAHs, liquid-cooled rear doors that pull hot air in from the servers and decrease power consumption of the fan, and spraying liquid into the server enclosure.

More and more data centers are implementing liquid cooling to handle the increasing loads, resulting from the higher power density of IT equipment such as blade servers. Liquid cooling takes advantage of the better transport and heat-transfer properties of liquids (water, refrigerant, or dielectric fluids) compared to gases (air) and takes advantage of the consequent ability to more easily remove high heat loads from racks. The most common approach is to integrate a chilled-water loop into the rack itself. Capturing heat directly from the racks allows for greater water-side economizer free-cooling, which can reduce the energy

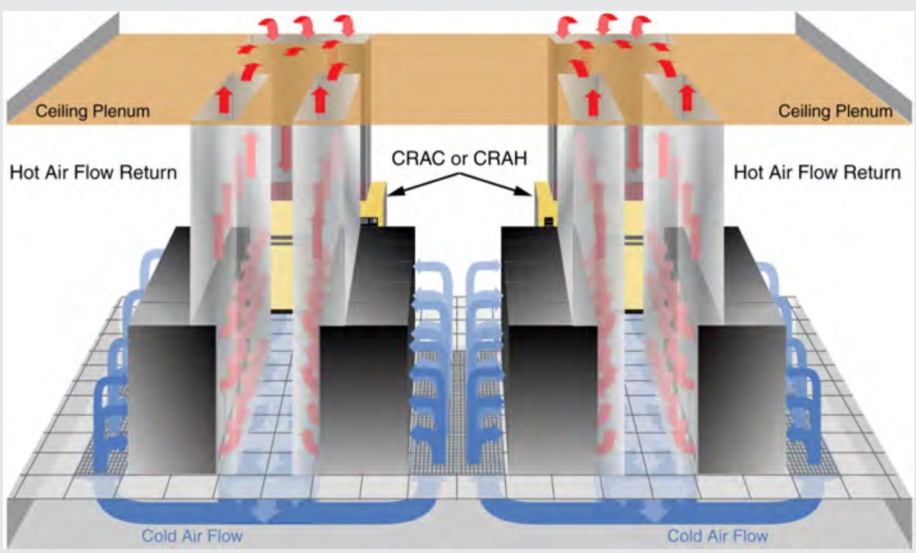

Figure 10 – Schematic of Rack Containment System

use of the cooling system by 70% or more. It may be possible to eliminate fans entirely with direct liquid cooling. Designs currently offered vary from simple chilled-water systems with integrated fans to systems that incorporate the chilled-water coils into the coolers of the rack cabinet. Many of these systems are designed with sophisticated leak-detection systems to minimize the negative consequences of leakage. These systems can accommodate a much higher power density than air cooling systems. These liquid cooling systems could use a lot less floor space than systems with CRACs or CRAHs.[31] Figure 11 shows a schematic of a typical liquid-cooled facility with four separate cooling loops: a condenser-water system (CWS), a chilled-water system (CHWS), a technology cooling system (TCS), and a Datacom equipment cooling system (DECS).[32]

These loops are interrelated and need to be designed and controlled to effectively move heat from the electronics through the loops to the external environment. Multiple coolant loops provide advantages in reducing the quantity of fluid that can leak into the data center, providing some control opportunities that can reduce energy and increase flexibility. Systems should be designed to provide ease of installation, commissioning, operation, maintenance, and troubleshooting; redundancy for critical components; and flexibility to scale the system up to accommodate growth in the IT heating load.[33]

The Coolant Distribution Units (CDU) and TCS isolate the electronics in the data center from the chilled-water system and provide enhanced control to prevent condensation, minimize the consequences of a liquid leak, and provide flexibility with regard to coolant fluid flow and heat transfer. The CDU can control the loop temperature to above the dew point, isolate the cooling fluid for the data center to minimize the negative consequences of a leak, and minimize the filtration requirements of the building chiller water. The CDU also affords isolation of the CHWS and CWS from the electronics and permits use of a different cooling fluid. Some commonly used approaches to liquid cooling of cabinets are shown in Figure 12 on the following page.[32,33]

Other approaches use a base plate or cold plate that is attached directly to the server internal components. The cold plate is cooled directly by liquid flowing over the cold plate surface[34] or sprayed onto the surface of the cold plate.[35,36] The electronics are cooled via the conduction path to the cold plate. Liquid does not enter the server enclosure with either approach. These approaches, while requiring planning in configuring the server racks, clearly have a lot of potential for removing large quantities of heat, but there is resistance to this technology (due perhaps to the need for such advanced planning and coordination and also due to the fear of placing coolant near the electronics).[37]

If the outside temperatures are low enough, outside air can be used for cooling, eliminating or reducing the need for vapor-compression loops. Avoiding the use of compressor energy provides very high cooling efficiency and associated energy savings. The equipment used to take advantage of these low outside air temperatures is called an *economizer*. Economizers can be configured as air-side economizers, water-side economizers, and adiabatic economizers.

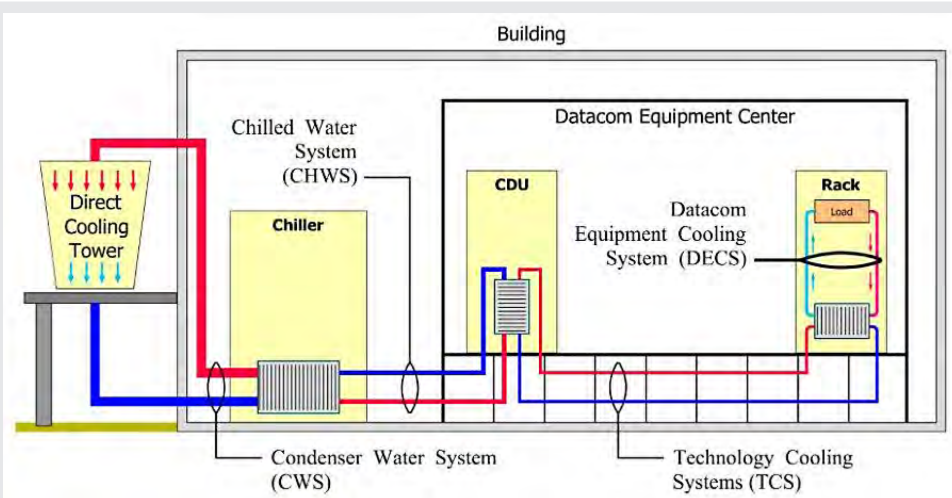

Figure 11 – Schematic of Liquid Cooling Loops in a Typical Liquid-Cooled Facility[38]

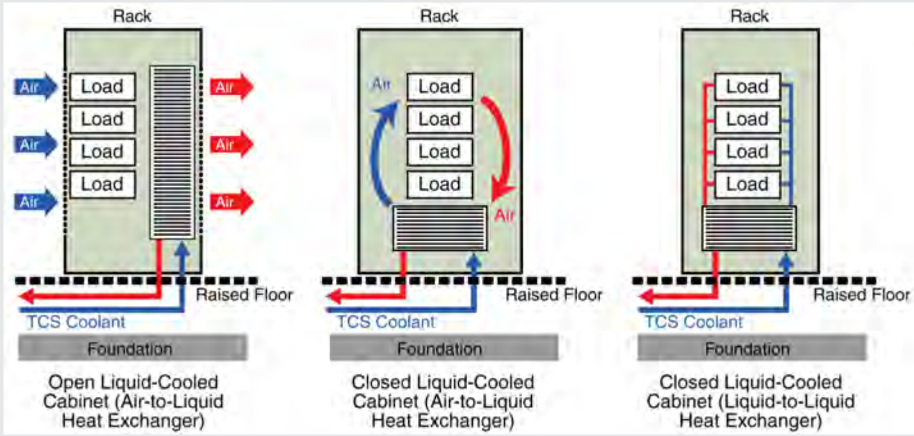

Figure 12 – Typical Configurations for Liquid Cooling of Racks

Air-side economizers bring outside air directly into the conditioned space of the data center and as such can adversely affect contamination and moisture levels in the data center. Therefore, humidity and contamination need to be monitored and controlled. Enthalpy control should be used to avoid the need for humidification.[38]

When the outdoor enthalpy is below the required supply air enthalpy, cooling via a vapor-compression cycle may not be required. (Full "free cooling" can be attained when the outside air is at least 2°F below the entering temperature of the cold aisle.) When the outdoor enthalpy is between the required enthalpy of the supply air and the required enthalpy of the return air, some cooling by way of a vapor-compression cycle will be required to deliver the air at the required supply-air enthalpy. (Partial benefit can be attained when the outside air temperature is at least 2°F below the air temperature returning to the air handler of the computer room.)[39] Some humidification may be required as well. Some studies have shown that higher contaminant levels are present in the air environment of a data center when using 100% outside air. To keep contamination levels under control, the MERV rating when using outside air needs to be 10 or 11. (This compares to a MERV of 8 or 9 with 100% recirculation air.)[40] Commissioning and recommissioning are advised to ensure that contaminants are kept under control.[41]

Water-side economizers use water rather than air to deliver free cooling when outdoor conditions are favorable. Water-side economizers have the advantage over air-side economizers in not introducing additional outside air and associated contaminants and moisture to the data center space during economizer operation. However, using water as an intermediate fluid introduces additional temperature differences that reduce some of the availability of "free cooling."

Integrated economizers are piped in series with the chiller and can be used to unload the chiller when full economizer operation is not possible. Parallel water-side economizers can be used to replace the chiller when the wet bulb temperature of the outside air permits complete economizer operation. When the economizer is operated in series with the chiller, additional free cooling benefits can be delivered, with the chiller making up for the cooling when outdoor temperatures are not low enough to permit the economizer to deliver water of a sufficiently low temperature to do all of the cooling.

Both direct and indirect systems are used. With the direct system, the water is cooled in a cooling tower or in a dry cooler by the ambient air. This water passes directly into the facility cooling coil in parallel with the chilled-water piping that would be used when the chiller is operating. The most common type of direct water-side economizer uses dry (air) coolers to cool the condenser water.[42]

With the indirect system, the water-side economizer piping has a flat-plate heat exchanger between the water in the cooling coils and the water in the condenser or dry cooler loop. The heat exchanger isolates the chilled-water loop from the condenser water to prevent fouling of the coils. Indirect systems require the outdoor wet bulb to be 7 to 10°F below the designed chilled-water temperature.

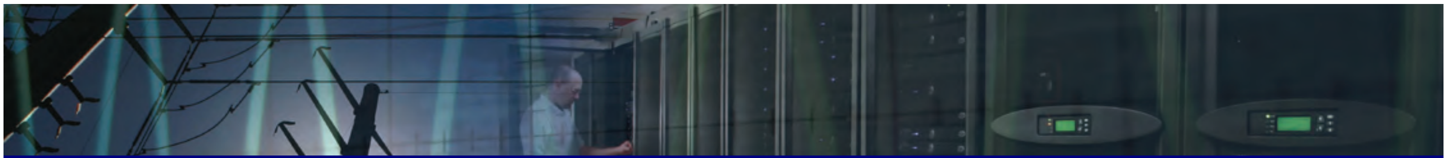

Water-side economizers are best suited to climates where the wet-bulb temperatures are 45°F or less (and therefore can be retrofitted to use most existing chillers designed for 42°F water) for more than 3000 hours per year. Water-side economizers can also be integrated into CRAC DX units with a free coil (which is cooled with a cooling tower, dry cooler, or spray cooler). The coil is inserted as an alternate path for the refrigerant.[16] At 55°F, they can be used in many CRACs with appropriate chilled-water reset.

Adiabatic-cooled economizers are derivatives of the air-side economizer and use a wetted media as a humidifier component. An adiabatic humidifier exposes water to an air stream and uses the heat energy, for example, of a waste-heat air stream to evaporate the water.[31]

Air-to-air indirect evaporative coolers can address the cooling load with a cooling energy-efficiency ratio of greater than 50.[43] Hot aisle return air is sensibly cooled by the ambient air stream drawn across the wetted exterior of the tubes. The units are sized to yield a 70% approach of the 100°F temperature of the hot aisle return to the ambient wet-bulb condition of the outdoor air. When coupled with a refrigeration system during high ambient humidity conditions, the EERs will be in the range of 12 to 15. These units, called *recirculation air cooling by evaporation* (RACE), offer the opportunity for efficient heat rejection while limiting contamination from outside air. During winter operation, it should be possible to operate without the spray pump.

Several studies have estimated the availability of free cooling in different locations under different conditions. Work by LBNL[44] citing Green Grid tools[45] shows that full economizer operation can be used for most of the year (over 98% of the time in San Francisco, 93% of the time in London, 75% of the time in New York, and 51% of the time in Dallas), with a 68°F supply air temperature and 50°F dew point. Partial use of an economizer (some use of the chiller to augment the free cooling) increases economizer availability to 100% in San Francisco, 96% in London, 81% in New York, and 57% in Dallas. Research shows that the adiabatically cooled economizer provides much greater availability than the air-side economizer, which, in turn, provides greater availability than the water-cooled economizer.[16] Supply temperature increases provide substantial increases in economizer availability.

What's in the Future?

Improvements are expected to continue in the efficiency of servers, electric-distribution equipment, and cooling equipment. Virtualization will increase in acceptance, along with other techniques and equipment used to consolidate server utilization. This will increase the power use and power density of servers, consequently requiring cooling equipment to remove greater heat loads, logically creating a shift to more liquid-cooling systems, including direct liquid cooling such as spray cooling. Modular server systems with packaged cooling systems and other ancillary components are likely to increase in acceptance, affording increased flexibility in meeting increases in facility computing loads.

In the power-conversion or electric-distribution areas, higher-efficiency power supplies, UPS systems, high-voltage power distribution, and DC systems will increase in market penetration. Encouraging this growth will be programs like ENERGY STAR and 80 PLUS programs, as well as efforts to improve DC components and CHP technologies

Studies to understand the effects of humidity and contamination on IT systems could result in the relaxing of limitations on humidity and contaminants in data centers or could encourage making the IT equipment more tolerant to these environmental conditions. Either of these outcomes would result in saving energy by reducing filtration, humidification, and dehumidification requirements as well as permitting greater use of air-side economizers.

Use of intelligent controls, including greater use of power management that does not impair the functionality of IT equipment, is expected to increase.

All of these improvements, as well as the best practices described in this paper, will be more widely accepted as well-documented case studies appear in the literature. Additionally, there is a distinct need for more carefully constructed and executed tests that isolate the effects of individual equipment attributes and operational alternatives, particularly on data center performance and energy use.

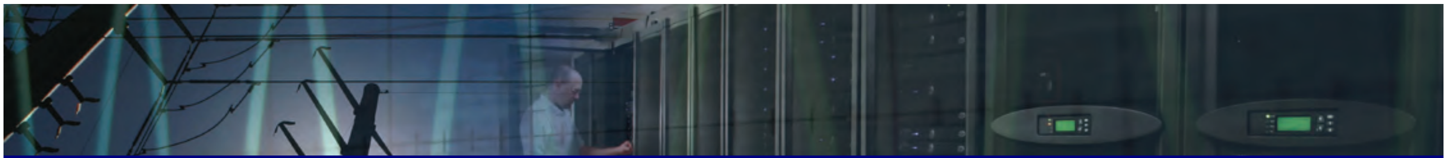

References

1 "*Saving Energy at Data Centers*," Fact Sheet, U.S. Department of Energy, Energy Efficiency and Renewable Energy, October 15, 2007 http://www1.eere.energy.gov/industry/saveenergynow/pdfs/data_center.pdf

2 Tschudi, William and Sartor, Dale, "*Data Center Efficiency*," Lawrence Berkeley National Laboratory, November 30, 2008 http://hightech.lbl.gov/presentations/Austin-Nov08-Part1.pdf

3 Uhlman, Kenneth, and Underpas, James, "*Out of Power, Out of Cooling*," Mission Critical, Data center and emergency backup solutions, pp 26-29, March/April, 2009 http://www.missioncriticalmagazine.com/digitalmagazine/SWF/MC/0409/index.html

4 Shouldice, Dennis, "*Extending the useful life of your data center*," APC by Schneider Electric, Datacenter Dynamics, San Francisco, July 17, 2009, http://www.datacenterdynamics.com/Media/WebPages/SF09%20H1%201230%20Dennis%20Shouldice%20-%20APC%20by%20Schneider%20Electric_0.pdf

5 Schmidt, Roger, et al., "*Increasing Energy Efficiency in Data Centers, Designing Better Data Centers*," ASHRAE Journal, Volume 49, Number 12, pp 18-25, December 2007, http://findarticles.com/p/articles/mi_m5PRB/is_12_49/ai_n25125554?tag=content;col1

6 "*The Vector Approach to Data Center Power Planning*," Eaton, July 7, 2009 http://viewer.media.bitpipe.com/1239644075_149/1247253200_553/The-Vector-Approach4.2.pdf

7 Moore, Bob, "*The Green Information Technology Revolution*," Mission Critical, Data center and emergency backup solutions, pp 42-48, March 1, 2008, http://www.missioncriticalmagazine.com/Articles/Features/BNP_GUID_9-5-2006_A_10000000000000269526

8 "*Virtualization Basics*," VMware, Downloaded July 15, 2009 http://www.vmware.com/technology/virtualization.html

9 Perkowitz, Jason, "*Cutting Energy Costs in Enterprise Data Centers*," APCON, Mission Critical, Data center and emergency backup solutions, Volume 2, Number 3, p41, May/June 2009 http://www.missioncriticalmagazine.com/CDA/Articles/Features/BNP_GUID_9-5-2006_A_10000000000000595610

10 Williams, Ben, "*Maximizing Data Center Efficiencies through Microprocessor Innovations*," EPA Conference on Enterprise Servers and Data Centers: Opportunities for Energy Savings, San Jose, CA, Advanced Micro Devices, January 31, 2006 http://hightech.lbl.gov/DCTraining/docs/server-conference/bwilliams-keynote.pdf

11 "*Reducing Data Center Energy Consumption, A summary of strategies used by CERN, the world's largest physics laboratory*," Intel, September 9, 2008 http://download.intel.com/products/processor/xeon5000/CERN_Whitepaper_r04.pdf

12 "*Case Study: Cassatt resource Optimization through Active Power Management*," July 30, 2008, https://microsite.accenture.com/svlgreport/Documents/pdf/case%20study_Cassatv2.pdf

13 Bulkeley, William, "*How Well Do You Know… The Cloud?*,"

14 Symanski, Dennis, "*Electrical Systems, Powering the data center*," EPRI for ASHRAE and Save Energy Now, April 28, 2009

15 Spears, Ed, "*Is your data center running out of power or cooling?*" Eaton, May 30, 2009, http://viewer.media.bitpipe.com/1239644075_149/1244606383_105/Is-your-data-center-running-out-of-power-or-coolingfinal.pdf

16 "*Best Practices for Datacom Facility Energy Efficiency*," ASHRAE Datacom Series, Second Edition, 2009, Available for purchase from ASHRAE at: http://www.ashrae.org/publications/

17 "*Opportunities for Transformer/PDU Savings in Data Centers*," PowerSmiths, July 2, 2008, http://svlg.net/campaigns/datacenter/ppt/%5bRT4%5dHigh_Efficiency_Power_Transformation.pdf

18 Ton, My et al, "*DC Power for Improved Data Center Efficiency*," Ecos Consulting, EPRI and Lawrence Berkeley National Laboratory, January 2007, http://hightech.lbl.gov/documents/DATA_CENERS/DCDemoFinalReport.pdf

19 Rasmussen, Neil, "*AC vs. DC Power Distribution for Data Centers*," White Paper #63, Revision 4, American Power Conversion, May 6, 2008 http://www.scribd.com/doc/4854614/AC-vs-DC-Power

20 Rasmussen, Neil, "*AC vs. DC Power Distribution for Data Centers*," White Paper #63, Revision 5, American Power Conversion, January 26, 2009 http://www.scribd.com/doc/4854614/AC-vs-DC-Power

21 "*Installing on raised floors or non-raised floors*," IBM Corporation, Downloaded September 3, 2009, http://publib.boulder.ibm.com/infocenter/dsichelp/ds8000ic/index.jsp?topic=/com.ibm.storage.ssic.help.doc/f2c_raisednonraised_1xobba.html

22 Neudorfer, Julius, "*Data center cooling optimization in the virtualized server world*," North American Access Technologies, SearchDataCenter, April 21, 2008, http://searchdatacenter.techtarget.com/tip/0,289483,sid80_gci1310600,00.html

23 Cronin, Dennis, "*Imagining the Perfect Data Center Power and Cooling Solution*," Mission Critical, Data center and emergency backup solutions, pp 8-10, March 1, 2008, http://www.missioncriticalmagazine.com/CDA/Articles/Column_Cronin/BNP_GUID_9-5-2006_A_10000000000000268163

24 Neudorfer, Julius, "*Virtual Servers Need REAL Cooling*," North American Access Technologies, Mission Critical Magazine, pp 28 – 30, July/August 2009 http://www.missioncriticalmagazine.com/CDA/Articles/Features/BNP_GUID_9-5-2006_A_10000000000000636703

25 Phelps, Wally, "*Cooler Servers with Less Cooling*," AdaptivCool, Mission Critical, Data center and emergency backup solutions, Volume 2, Number 3, p26, May/June 2009, http://www.missioncriticalmagazine.com/Articles/Features/BNP_GUID_9-5-2006_A_10000000000000595564

26 Sullivan, Robert, "*Data Center Energy Efficiency*," Keynote Presentation at EPA Conference on Enterprise Servers and Data Centers: Opportunities for Energy Savings,, San Jose, CA, The Uptime Institute, Inc., January 31, 2006 http://hightech.lbl.gov/DCTraining/docs/server-conference/bsullivan-keynote.pdf

27 "*Focused Cooling Using Cold Aisle Containment*," Emerson, Liebert, July 7, 2009, http://www.liebert.com/common/ViewDocument.aspx?id=1295

28 Niemann, John, "*Hot Aisle vs. Cold Aisle Containment*," White Paper # 135, APC, October 7, 2008, http://www.apcmedia.com/salestools/DBOY-7EDLE8_R0_EN.pdf

29 Donowho, Brian, "*High-Performance Data Centers Call for Innovative Cooling Techniques*," Mission Critical, Data center and emergency backup solutions, pp 60-63, March 1, 2008, http://www.missioncriticalmagazine.com/Articles/Features/BNP_GUID_9-5-2006_A_10000000000000269486

30 Dunlap, K., and Rasmussen, N., "*The Advantages of Row and Rack - Oriented Architectures for Data Centers*," APC by Schneider Electric, June 23, 2006 http://www.apcmedia.com/salestools/VAVR-6J5VYJ_R0_EN.pdf

31 "*High Performance Data Centers, A Design Guidelines Sourcebook*," Pacific Gas and Electric Company, January 2006 http://www.scribd.com/doc/11806979/High-Performance-Data-Centers-A-Design-Guidelines-Sourcebook

32 Whitenack, Katherine, "*Liquid Cooling for Datacom Equipment Centers*," Electronics Cooling, 14 (3), Lytron, August 2008 http://www.electronics-cooling.com/html/2008_aug_a3.php

33 "*Liquid Cooling Guidelines for Datacom Equipment Centers*," ASHRAE Datacom Series, 2006, Available for purchase from ASHRAE at http://www.ashrae.org/publications/

34 "*Benefits of Direct Cooling in the Data Center*," Clustered Systems, Downloaded October 29, 2009 http://www.clusteredsystems.com/494/6801.html

35 "*SprayCool® Sealed Transit Case Solutions*," Downloaded May 14, 2009 http://www.spraycool.com/products/racks.asp

36 Finch, Andy, "*Mechanics of Spray Cool direct spray*," SprayCool, January 2009. http://www.spraycool.com/products/whitepapers/Mechanics_of_SprayCool_direct_spray.pdf

37 Dogra, Rakesh, "*Choices in liquid cooling for your data center*," The Data Center Journal, March 31, 2009 http://datacenterjournal.com/index.php?option=com_content&task=view&id=2671&Itemid=41

38 Gembinski, William, et al., "*Outside Air Inside Data Centers*," Engineered Systems Magazine, pp 76-80, September 2008 http://www.nxtbook.com/nxtbooks/bnp/es_200809/

39 "*Energy Efficient Solutions for Data Centers, Introducing the Emerson Network Power tool for evaluating cooling system total cost of ownership*," Liebert/Emerson, 2007, http://www.liebert.com/common/ViewDocument.aspx?id=879

40 Tschudi, William, "*Data Center Economizer Contamination and Humidity Study*," LBNL/PG&E, 2007 http://hightech.lbl.gov/documents/DATA_CENTERS/EconomizerDemoReportMarch13.pdf

41 Sorell, Vali, "*OA Economizers for Data Centers, Designing Better Data Centers*," ASHRAE Journal, Volume 49, Number 12, pp 26-31, December 2007, http://findarticles.com/p/articles/mi_m5PRB/is_12_49/ai_n25125555?tag=content;col1

42 "*Design Considerations for Datacom Centers*," ASHRAE Datacom Series, Second Edition, 2009, Available for purchase from ASHRAE at http://www.ashrae.org/publications/

43 Dunnavant, Keith, et al, "*Reduce Data Center Cooling Costs by 75%*," Engineered Systems, pp 34 - 41, April 2009, http://www.esmagazine.com/Articles/Feature_Article/BNP_GUID_9-5-2006_A_10000000000000561587

44 Mathew, Paul, et al. "*How Does Your Data Center Measure Up?*" Lawrence Berkeley National Laboratory, HPAC Engineering, Volume 81, Number 5, pp 16-21, May 2009, http://hpac.com/fastrack/how-data-center-measure-0509/

45 Frier, Eric, "*Calculated Free Cooling Savings, Free cooling estimated savings calculation tool*," Weatherbank Inc. for Green Grid, February 2009 http://cooling.thegreengrid.org/calc_index.html

EPRI Resources

Mort Blatt, Contractor
650.855.2457, mblatt@earthlink.com

Brian Fortenbery, Program Manager, EPRI
865.218.8012, bfortenbery@epri.com

Satish Rajagopolan, Sr. Project Engineer/Scientist, EPRI
865.218.8043, srajagopolan@epri.com

Dennis Symanski, Sr. Project Manager, EPRI
650.855.1000, dsymanski@epri.com

Energy Efficiency, Program 170

1024624 November 2011

Electric Power Research Institute
3420 Hillview Avenue, Palo Alto, California 94304-1338 • PO Box 10412, Palo Alto, California 94303-0813 USA
800.313.3774 • 650.855.2121 • askepri@epri.com • www.epri.com