the green grid℠
get connected to efficient IT

SEVEN STRATEGIES TO IMPROVE DATA CENTER COOLING EFFICIENCY

EDITOR:

RICH JONES, CHATSWORTH PRODUCTS

CONTRIBUTORS:

JOHN BEAN, APC

RONALD BEDNAR, EMERSON

DAVID MOSS, DELL

JOHN WALLERICH, DELL

JIM SMITH, DIGITAL REALTY TRUST

DR. MICHAEL PATTERSON, INTEL

Version 1.0 Release Date: 21 October 2008

INTRODUCTION

In 2007, the Technology and Strategy Work Group of The Green Grid consortium established a system architecture Task Force to investigate and provide guidance on existing and emergent cooling technologies that can improve the efficiency of data center cooling architectures. This White Paper provides insight on seven strategies—identified by the Task Force—that, when properly implemented, can have a significant effect on the cooling portion of the Power Usage Effectiveness (PUE) metric[1]. The paper is intended to be used as a guide for individuals responsible for cooling architecture design and operations decisions, including facility engineers, data center architects, designers, managers, and facility owners.

The seven identified strategies are:

1) Developing an air management strategy
2) Moving cooling systems closer to the load
3) Operating at a higher delta-T
4) Installing economizers
5) Using higher-specification and performance equipment
6) Using dynamic controls
7) Maintaining higher operating temperatures

This white paper is a high-level, qualitative overview of the strategies. In order to fully explore the benefits of these strategies, the Task Force will continue to examine them, factoring in quantitative analysis, recommendations, and industry challenges where necessary.

Note: For the purposes of this white paper, the term ACU (Air Cooling Unit) is used to describe any device used to cool air in the data center.

DEVELOPING AN AIR MANAGEMENT STRATEGY

In most cases, a fully developed air management strategy can produce significant and measurable economic benefits and, therefore, should be the starting point when implementing a data center energy savings program. This holds true regardless of the focus—whether simply reducing operating costs, mitigating problem cooling areas in an existing space, upgrading an existing space to higher power densities, converting an existing space into a data center, designing and building a green field project, or isolating the cool supply air from heated return.

Developing an air management strategy can address a multitude of challenges. Without management, air will follow the natural dynamics set up by a facility’s physical layout and the positioning and characteristics of its IT and cooling equipment. This could lead to the hot air and cold air mixing, and it could produce uncertainty in the matching of equipment deployment and relative rack capacity.

Prior to developing an air management strategy, it is important to first understand the airflow characteristics of the data center. Even without an air management strategy, a greater understanding of the airflow dynamics associated with the data center can lead to efficiency improvement opportunities. Throughout the following discussion, it is assumed that a hot/cold aisle layout has already been implemented as one of the first steps toward improving efficiencies in the data center cooling infrastructure.

DATA CENTER AIRFLOW

There are a number of facility dynamics that can affect and impede airflow within the data center. All data centers have variations in raised-floor plenum pressure, which results in variance in the volume of air delivered. These variations are further complicated by raised-floor height and the number of obstacles (cabling, power, piping, etc.) under the raised floor.

Delivered flow rates need to loosely match rack air consumption at the row/aisle level. The dynamics of the data center will determine whether the volumetric delivery has to exceed—or can be a bit less than—the rack consumption. An understanding of areas of low flow rate will lead to better planning and more optimized use of the cooling infrastructure. Perforated floor tiles located close to ACUs generally have less flow. End racks generally have more side recirculation of hot exhaust air. Because of this, the end racks in some designs generally will not support as much equipment as the racks in the middle of a row.

Computational fluid dynamics (CFD) analysis can be used to optimize data center airflow by identifying weak areas of cooling capacity in the data center. CFD enables a designer to take full advantage of available air pressure by plotting existing pressure variables and then employing techniques to efficiently distribute that air pressure throughout the room. A CFD analysis can determine the impact on cooling resources when changes are made to the equipment layout, ACU flow rate, floor tiles, heat load distribution, and other supplementary cooling systems characteristics. It also can offer temperature estimates for given rack loadings. Figure 1 shows an example of a CFD analysis. The cut-plane across the space shows the air temperature in the space. A typical CFD analysis also includes similar displays for airflow velocity.

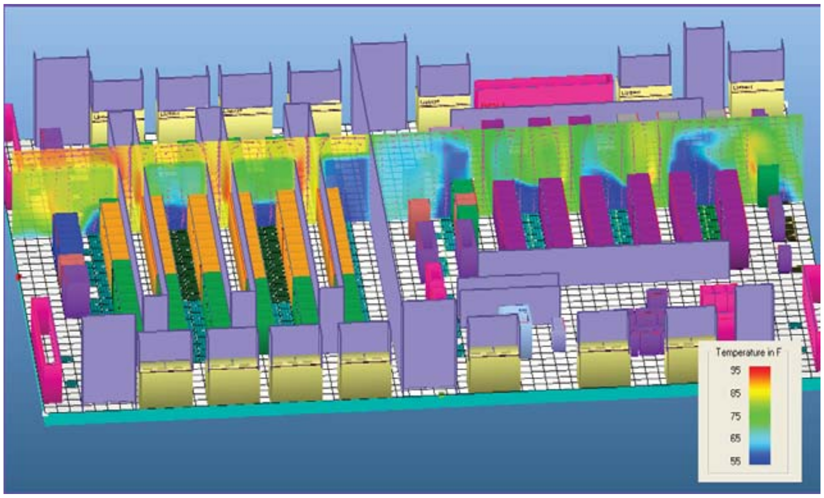

FIGURE 1: EXAMPLE OF AN ENHANCED MULTI-AIRFLOW DESIGN

Graphic courtesy of Intel Corporation

Understanding each data center's unique airflow dynamics at the equipment level and incorporating that understanding into the deployment process can significantly increase the effectiveness and efficiency of the cooling equipment. This will help prevent the facility from having to over-deploy cooling equipment to overcome poor airflow dynamics.

BYPASS AND HOT AIR CONTAMINATION

Bypass air is the volume of cold supply air that enters the room but does not directly enter the IT equipment. Bypass air can be found in cable cut-outs, coming through vents intentionally placed in hot aisles, around racks that are not tightly spaced, and other such areas. Increases in bypass air directly reduce the overall effectiveness of cold air being delivered by the cooling infrastructure and often result in the need to over-provision cold air. Sealing around cable cut-outs or gaps between racks is an efficient method for reducing one of the more common sources of cold air bypass.

Hot air contamination describes the phenomenon in which equipment exhaust leaks back to equipment inlets, either inside or outside of the rack, resulting in raised inlet temperatures. It also results in hot exhaust air being cooled prior to returning to ACUs, which further reduces the efficiency of the cooling infrastructure. Internal contamination paths can be minimized by sealing air leakage around the equipment. This is commonly achieved by using blanking panels to fill spaces between equipment and using air dams on the sides of equipment in racks that allow lateral leakage.

The best way to control external contamination is to implement strategies that maintain separation of hot and cold air. Reduction or elimination of bypass air can have a significant impact on the effectiveness of the cooling infrastructure and related energy costs.

MOVING COOLING SYSTEMS CLOSER TO THE LOAD

ACU COOLING

Raised-floor, room-cooling systems have traditionally proven to be an effective approach to data center environmental management. However, as rack densities exceed 5 kW and load diversity increases across the room, non-traditional cooling architectures should be evaluated for their impact on cooling system performance and efficiency.

Locating cooling closer to IT equipment can reduce data center cooling costs by more than 30% compared with historical approaches to cooling. Air conditioning unit fans are known to consume a significant portion of energy in most data center cooling systems. Mounting the cooling modules as close as possible to the source of heat—such as placing them directly above, alongside, or within high-density racks—reduces the distance that the fans must move air. This can provide up to 70% savings of the energy required to move the air.

This approach may provide existing data centers the flexibility to support new and greater IT loads, especially if there is no upgrade path available. For new data centers, the answer may be a hybrid approach, in which certain racks use high-density supplemental cooling (rack-specific or localized cooling) while others are supported by traditional room cooling (perimeter ACUs with IT equipment on a raised floor).

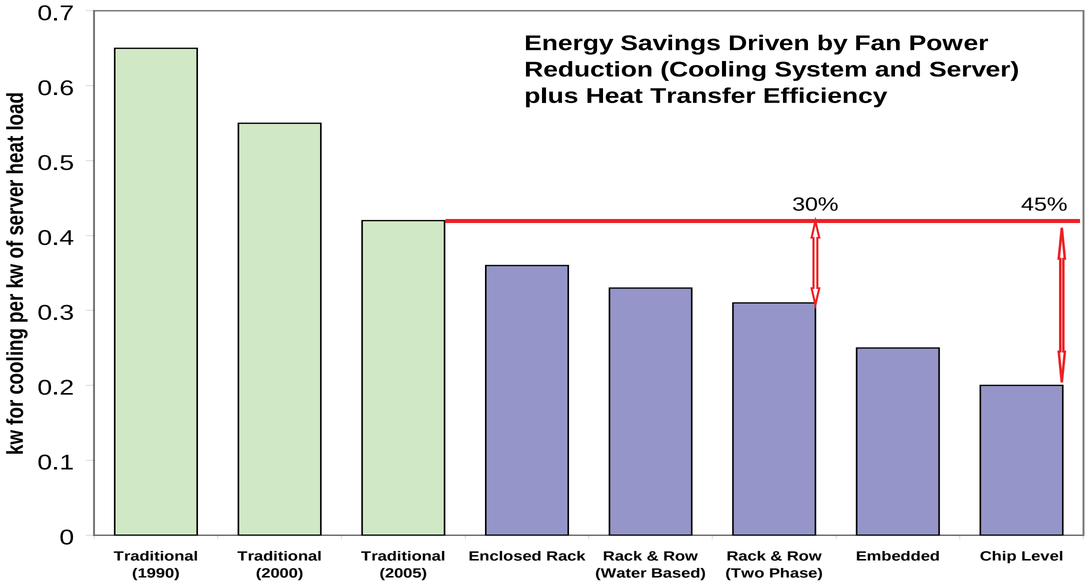

FIGURE 2: RELATIVE EFFICIENCY POTENTIAL OF DIFFERENT COOLING SOLUTIONS

Graphic courtesy of Emerson Network Power

LIQUID-BASED COOLING

Water or refrigerant liquid-based cooling approaches are more effective than air at transferring heat, which is especially beneficial in higher-density applications. In a liquid-based cooling system, hot air passes through an air-to-water or air-to-refrigerant heat exchanger located near the heat load. The heat is transferred to the liquid where it can be more efficiently removed from the building.

Because fluid properties vary, it is important to choose the correct heat-transfer fluid for the application. A single-phase fluid (e.g., water) transfers heat by changing its temperature, while a two-phase fluid (e.g., refrigerant) transfers heat by changing state. More effective heat-transfer fluids (e.g., refrigerants and water) enable cooling systems to be designed to improve overall system energy efficiency (fewer parts, efficient coils and pumps, etc.).

OPEN/CLOSED COOLING ARCHITECTURE

Another important system parameter to consider is open or closed cooling architecture. Open architecture systems use cooling coils that are placed directly above, alongside, behind, or below the server racks. The room's air volume is then used as thermal storage to ride through short power outages. Cool room air enters the front of the server rack cabinet and absorbs heat as it passes over the electronics. It then exits the cabinet, moves back into the room, and passes through the heat exchanger, which removes the heat.

Closed architecture systems involve internal air circulation in server rack cabinets that are completely sealed from room air and environmental conditions. In a closed design, air is circulated through the electronics and passes through a liquid-to-air heat exchanger located in the cabinet. It is then returned back to the electronics. The liquid in the heat exchanger absorbs all the heat.

Table 1 presents some of the advantages and disadvantages of open versus closed cooling systems.

	ADVANTAGES	**DISADVANTAGES**
OPEN	- System redundancy - Initial and operating costs - No limit on rack selection - Self-regulating capacity - Lower cost than closed systems	- Audible noise - Room solution
CLOSED	- Low audible noise - Deployable on a single rack - Minimal bypass air losses	- Short ride-through on cooling loss - Higher cost than open systems - Fixed capacity

TABLE 1: OPEN VS. CLOSED RACK COOLING ARCHITECTURES

Emerging technologies such as embedded cooling and chip-level cooling also have the potential to enhance energy efficiency in higher-density applications. Embedded cooling delivers high-efficiency cooling directly inside the rack through cooling coils that are installed between server modules. Chip-level cooling takes this approach to the next level by helping to move heat directly away from the chip inside the server. This can reduce the energy expended both by the server cooling fans and by the facility cooling system.

Data center energy efficiencies are realized when heat is removed as close to the source as possible. As embedded and chip-level cooling solutions are deployed in more data centers, a highly efficient, three-tiered approach to data center cooling may emerge. In this approach, heat is effectively moved away from the chip and then cooled in the rack, with stable temperatures and humidity maintained by room air conditioners.

OPERATING AT A HIGHER DELTA-T

The term delta-T refers to the difference in temperatures of two measured points. In this white paper, the term is used to either describe the heating of air as it passes through IT equipment or the cooling of air as it passes through cooling equipment.

Until very recently, all IT equipment operated with constant-speed fans to accommodate worst-case inlet temperatures. These IT systems were designed in an era when peak cooling and energy efficiency were lower priorities. Today, we know that constant-speed fans are very inefficient for IT equipment whose inlet air temperatures are in a lower, more typical temperature range. For applications that require high airflow rates to cool components, a constant-speed fan is inefficient; it basically over-cools in a scenario with low room temperatures or low IT power dissipation[4]. Figure 3 shows system fan power for three servers as it varies with ambient temperature. The exponential power increase from low to high speed is the result of a much smaller linear increase in the airflow rate. The result is significant power savings because of the ability to throttle down the airflow when it is not needed.

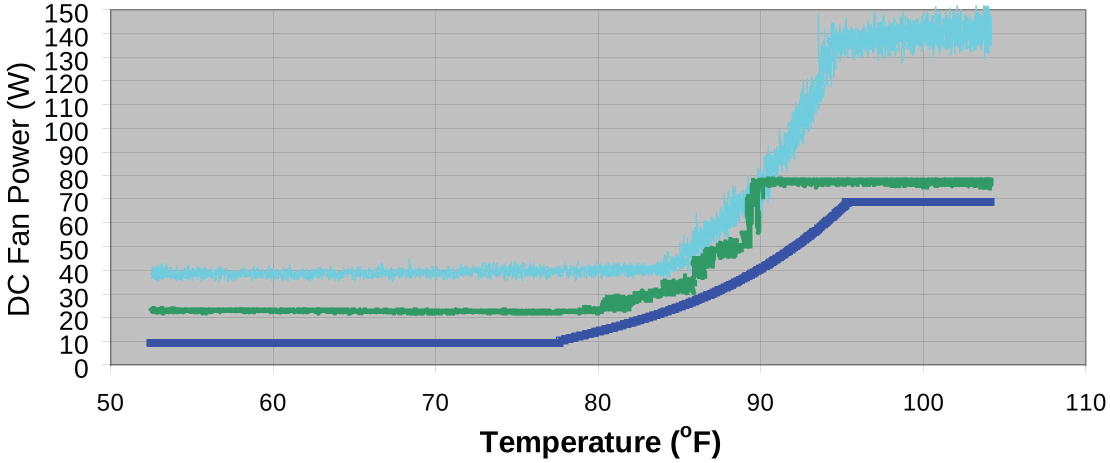

FIGURE 3: DC FAN POWER VERSUS AMBIENT TEMPERATURES FOR THREE DIFFERENT SERVERS

Graphic courtesy of Dell

The delta-T is inversely proportional to the amount of airflow through the IT equipment. For example, if the flow rate is decreased by half, the delta-T is doubled. Thus, one advantage of using IT equipment with higher delta-T is the power savings associated with the fans running at a lower speed.

Another advantage of a higher equipment delta-T is a higher return temperature at the ACU, allowing the ACU to operate closer to the rated cooling capacity. This can then lead to fewer ACUs needed to cool the same load and, in turn, less energy required for air movers in the ACUs. The warmer the air passing through the ACU coil, the less likely the entire coil surface area will experience a temperature below the dew point. If more of the coil surface is positioned above the dew point, condensation is decreased. Energy is wasted in both the condensation process and the subsequent re-humidification process.

INSTALLING ECONOMIZERS

Economization is a technique where waste heat is rejected outside of the data center environment by reducing or eliminating refrigeration cycle use.

There are primarily two types of economization: airside and waterside.

AIRSIDE ECONOMIZATION

In airside economization, hot return air is vented directly outside and cooler outside air is drawn into the air-handling system for conditioning, humidification, and filtration. It is then delivered to the data center as cool supply air. Airside economization produces energy savings any time the outside air temperature is below the return air temperature.

While traditional airside systems can be deployed with almost any other cooling technology being used for heat rejection, they do require tight control of temperature, contaminants, pollutants, and changing humidity conditions when directly supplying outside air to data centers. To achieve maximum benefit, airside systems should ideally be managed by an external control or building automation system.

Most large-scale, airside-economized systems are easier to implement in new data centers as opposed to existing ones because moving large volumes of air in and out of the data center requires a large cross section of airflow. The volume of open space needed will have an impact on the architectural and structural design considerations of the data center.

Airside economization is best suited in climates with a stable and moderate temperature and humidity profile. However, airside economized projects in colder, drier climates also can be efficient. These applications work best at sites that can tolerate wider humidity ranges for their technology equipment than is typical for many enterprise users.

WATERSIDE ECONOMIZATION

In waterside economization, heat in the return air is transferred to a chilled water or glycol loop with traditional air-handling equipment. It is then transported outside and rejected to the atmosphere via multiple different heat exchange mechanisms. Waterside economization is limited to use in chilled water systems (with a few exceptions) and is often deployed at relatively large scales. Many consider it less risky than airside economization because the heat exchange mechanism provides a physical barrier between outside and inside systems.

The two main advantages of waterside economization are its minimal impact on the data center space and its overall effectiveness. In waterside systems, the “inside-the-data center” architecture is identical to non-economized systems; heat in the return air stream is transferred across a coil connected to the chilled water loop.

In general, waterside economization is better suited for colder climates. Many data centers operate chilled water loops between 40^0 and 50^0 F. When the outside temperature drops even just a few degrees below this point, direct heat transfer to the surrounding air provides all the cooling necessary without use of the refrigeration cycle.

USING HIGHER-SPECIFICATION AND PERFORMANCE EQUIPMENT

The design choices made for the selection of equipment for a new, expanded, or retrofit data center will affect not only capital and operational costs, but also the efficiency of the cooling system and data center as a whole. Therefore, it is vital that those decisions be made by the entire data center team, including facility owners, IT owners, and those individuals responsible for capital and operational budgets.

As density and the cost of energy continue to increase, it is also important that total cost of ownership (TCO) and return on investment (ROI) be factored into the decision-making process. Fundamental engineering economics can even be used to analyze capital alternatives, their initial costs, and operational cost implications. For example, something as simple as a high-efficiency pump or fan motor may cost 25% more than a non high-efficiency alternative. However, very often the power savings of that pump or fan motor will soon offset the increased initial expenditure.

There also exist potential external sources for funding these types of incremental choices. Many local and state governments and utility providers offer rebates, credits, and incentives for making a choice that results in a more energy-efficient data center. For instance, these types of funding could fully or partially offset the initial cost of the aforementioned fan motor.

The following are examples of some of the opportunities data center teams may encounter.

Design Agent: This often overlooked choice potentially can have the greatest impact. Hiring a highly capable, experienced architecture and engineering (A&E) firm ensures that available efficiency options will be considered. The A&E scope of work usually includes an analysis of energy options with regard to capital equipment selection, including TCO calculations. Qualified A&E firms typically include a CFD analysis of the room in their proposals as well. Most A&E firms also are likely to be aware of current energy codes and local energy programs (rebates and incentives) and are knowledgeable about the most efficient, thorough and detailed design packages.

High-Efficiency Motors: As mentioned previously, high-efficiency motors can potentially pay for themselves quickly. They will reduce energy use and, in most cases, provide an attractive ROI. In cases where a motor comes as part of a complete, packaged ACU, it is still worth pursuing a high-efficiency motor choice and specifying it if offered by the supplier.

Variable Speed Drives (VSDs): Also referred to as variable frequency drives (VFDs) or adjustable frequency drives (AFDs), this choice in equipment design is being considered more often. Significant power savings are potentially obtainable through off-peak and turn-down performance of much of a data center's utility infrastructure, particularly because many data centers operate at less than full load for much of the time. VFDs can be used on fans, chilled water pumps, air-conditioning chillers, cooling towers, and more. They generally offer excellent ROI and energy savings.

Instrumentation and Controls (I&C): Proper choice of I&C components and design can lead to better efficiency. Sufficient monitoring for tracking and analyzing cooling system operation and energy use is vital to measuring and optimizing efficiency. To determine the value of a site's PUE, the energy expended in the cooling system must be known. If an I&C system does not have the ability to monitor cooling system energy use, the efficiency benefits of a real-time PUE trend will not be realized. The next section discusses further economic benefits of a capable I&C system.

Humidifiers: Choosing among these often-overlooked energy consumers typically comes down to a TCO analysis. Generally, the least expensive option, initially, is a heat-based humidifier that converts water to steam and adds it to the airflow. Yet this model may carry a significant long-term energy cost. An option that costs more initially but is far more efficient is an ultrasonic humidifier, which uses a high frequency diaphragm to create a fine mist of water that is easily evaporated into the air stream. These units demand far less energy per unit mass of liquid absorbed by the air than do thermally driven units. The number of hours the humidifier will be used each year also should be considered.

Chiller Type: Chiller type is often one of the most complex decisions to make because of the many factors involved in addition to TCO analysis. For instance, a scroll or centrifugal compressor is nearly always more efficient than a reciprocating compressor. However, the compressor decision is often driven by the chiller capacity (e.g., scroll compressors may not always be available in smaller chillers).

Another consideration is the choice of liquid-cooled or air-cooled condenser for the cooling plant. The air-cooled condenser generally costs less than the liquid-cooled alternative, but it is also less efficient. The hours of usage are also a factor. If the data center design makes extensive use of economizers (see preceding section) and the chiller system is only used a limited number of hours per year, then the air-cooled condenser may have a compelling TCO advantage in spite of a lower efficiency when in operation.

USING DYNAMIC CONTROLS

Another strategy for improving data center efficiency is to take advantage of dynamic controls that allow power demands of cooling equipment to scale with the thermal load of the IT equipment.

Modern servers frequently employ multiple power states for the CPU cores along with variable airflow through the chassis. Virtualization is a driver for even greater power diversity of servers within the same space[5]. With virtualization, many machines might be running multiple concurrent applications, while others might be in either idle or sleep states or powered off entirely. Techniques to increase or decrease compute loads based upon hourly compute needs will cause swings in the heat that needs to be removed from IT equipment in the data center.

This section considers two separate capacity modulation schemes: airflow management and compressor mass flow management. Both of the schemes offer relatively short payback periods on the incremental price increase of purchasing the better technology.

AIRFLOW MANAGEMENT

Adopting a dynamic fan speed for data center cooling equipment can lead to significant energy savings. Even a modest change to cooling equipment airflow will provide a profound power reduction.

Table 2 offers a quick representation of fan drive power reductions based upon various airflow reductions.

AIRFLOW FROM NOMINAL	POWER FROM NOMINAL
100%	100%
80%	51%
60%	22%
40%	6.4%

TABLE 2: RELATION OF FAN SPEED TO POWER CONSUMPTION

An important consideration for the selection of dynamic cooling equipment is the allowable maximum turn-down ratio. Turn-down ratio is the full airflow rate in cubic feet per minute (CFM) divided by the lowest allowed CFM. Cooling equipment can experience operational and/or reliability issues if the turn-down ratio is exceeded. Equipment manufacturers supporting dynamic fan management have built in control limits to avoid less-than-desirable operating points.

The control system that manages the fan speed is an equally important consideration. Ideally, the control system would seamlessly scale the ACU's airflow in real time as a function of the IT equipment load. While there may be some successful cases of ad-hoc systems, it is best to consider products with such capabilities provided by the OEM. This will assure component suitability, balance, and stable operation, as well as characterized system performance.

COMPRESSOR MASS FLOW MANAGEMENT

Whether the cooling system is based upon fixed or dynamic airflow, incremental efficiency also may be gained by modulating the compressor capacity as needed to achieve a balance between cooling capacity and IT equipment thermal load. There are a variety of compressor modulation schemes available, including VFD, pulse width modulation (digital scroll), variable inlet geometry slides/vanes, and cylinder unloading. All of the options, except cylinder unloading, offer a virtually continuous variable response from the minimum turn down to maximum capacity.

MAINTAINING HIGHER OPERATING TEMPERATURES

Most data centers are run at temperatures much lower than is necessary for IT equipment. There are a number of reasons that temperatures are often kept too cold:

- There is misguidance on which temperature is important to the IT equipment. Common practice is for ACU operation to be governed by return air temperature. If the (return) set point meets ASHRAE recommendations between 68^{0} F and 77^{0} F, the supply temperature to the IT equipment must be much colder than needed—with an unwarranted energy penalty[3].

- The room is being kept cold to achieve a longer ride-through time during a cooling outage. A few degrees' increase in set point can have a large energy effect; however, it is unlikely that lowering a room's temperature by a few degrees will offer any meaningful increase in ride-through time. For example, if the maximum server temperature is 95 degrees F, running the room at 68^{0} F versus 70^{0} F will offer very little extra time in the event of a cooling failure.

- There is a fear that higher IT equipment temperatures will affect reliability. Consider the graphs in Figures 4 and 5, which compare the temperature of a typical server component and the wattage required to run the system fans with an increasing system ambient temperature. In Figure 4, where the system fans are held at high speed, component temperature tracks fairly closely to inlet ambient temperature. However, the system fans are using a large amount of power; they represent about 20% of the system's power requirement.

In Figure 5, the fans are allowed to vary according to normal control algorithms. System fan power is drastically lower due to variable speed control. Component temperatures are higher with variable speed fans, but they remain fairly constant with respect to increasing ambient temperature.

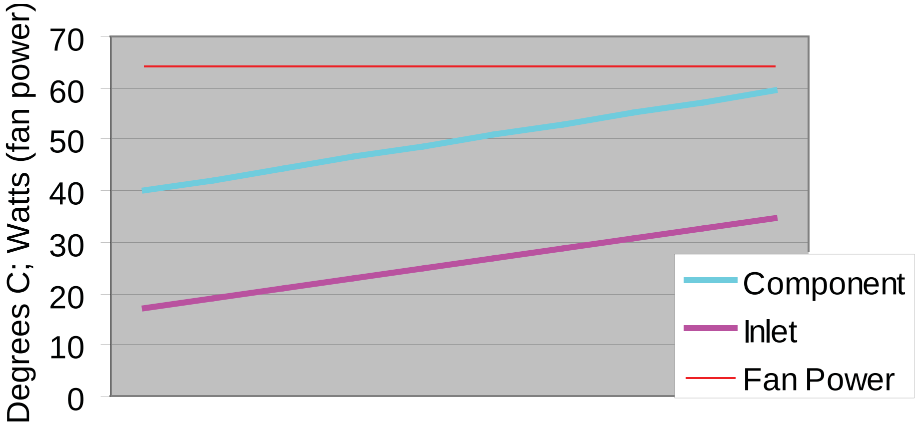

FIGURE 4: TYPICAL COMPONENT TEMPERATURE RESPONSE TO CONSTANT- SPEED FAN: INCREASING INLET AMBIENT TEMPERATURE

Graphic courtesy of Dell

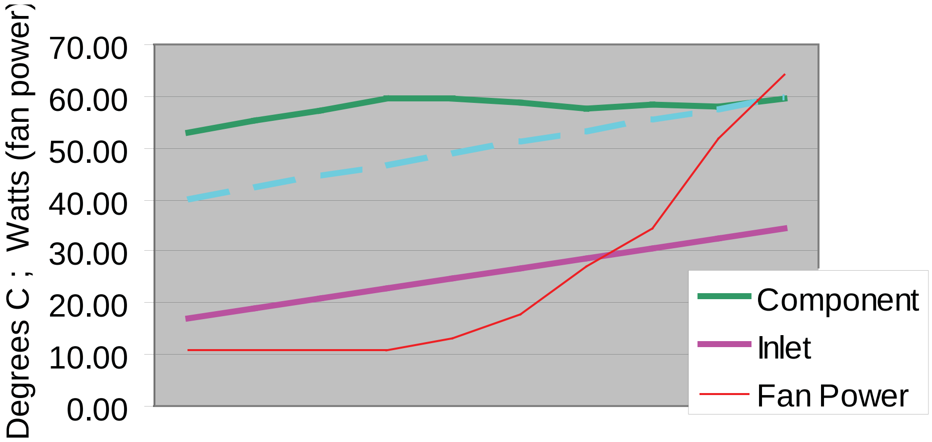

FIGURE 5: TYPICAL COMPONENT TEMPERATURE RESPONSE TO VARIABLE-SPEED FAN: INCREASING INLET AMBIENT TEMPERATURE

Graphic courtesy of Dell

- The room is being kept colder to achieve marginal inlet temperatures in worst-case locations. Nearly every data center has areas with poor airflow dynamics and experiences elevated temperatures for a subset of its equipment. Any attempt to raise operating temperatures should take these locations into account. A variety of actions can be taken to lessen the risk to these areas, such as moving tiles to rebalance airflows, moving actual IT systems to cooler areas, implementing supplemental cooling for specific hot spots, and establishing containment strategies for cold or hot aisles to ensure appropriate airflow segregation.

- The room temperature is based on personnel comfort level. Obviously personnel comfort will have to be weighed against the opportunities for energy savings. Some of the more aggressive efforts to operate at high temperatures have been accompanied by separate air-conditioned spaces adjacent to the data center, as well as service strategies that limit the amount of time spent in the data center. Perforated tiles may even be used temporarily in a hot-aisle work area during maintenance activities.

A higher operational temperature should be a consideration in the search for increased data center efficiency. Coupled with ACU energy-saving options, an increase in operating temperatures offers the opportunity of saving large amounts of energy—5% or more at the facility level is not uncommon[2].

For reasons stated in this paper, IT equipment fan power should be taken into account when considering any increase in set points. The Green Grid recognizes and endorses the recommended operational limits set forth by ASHRAE TC 9.9 3.

SUMMARY

The cooling systems of today's typical legacy data center are often highly inefficient, primarily due to the fact that cooling high-density equipment was not originally a requirement. However, by properly implementing the seven strategies identified in this white paper—along with developing a deep understanding of data center layout, operation, and new technologies—significant data center efficiency gains can be realized.

REFERENCES

1. The Green Grid (2007). *Existing Metrics, Guidelines and Programs Affecting Data Center and IT Energy Efficiency*

2. Miller, R. (2007). *Data Center Cooling Set Points Debated,* Retrieved 1 October, 2008 from http://www.datacenterknowledge.com/archives/2007/Sep/24/data_center_cooling_set_points_debated.html

3. ASHRAE Technical Committee 9.9 et al. (2004). *Thermal Guidelines for Data Processing Environments.*

4. Patterson, M. K, Steinbrecher, R. and Montgomery, S. Comparing Data Center & Computer Thermal Design (ASHRAE Journal), 4/06, pg 38-43.

5. Blackburn, M. (2008). *Five Ways to Reduce Server Power Consumption* The Green Grid.

D. Fenton and W. Richards. (2001). *ASHRAE Standard 15-2001 "Safety Standard for Refrigeration Systems"*

ASHRAE (2008) *Best Practices in Datacom Facility Energy Efficiency*

Sorrel, Vali. OA *Economizers for Datacenters (ASHRAE Journal)*, 12/07, pg 32-37.

Telecommunications Industry Association's. (2005). *TIA-942 Telecommunications Industry Association Standard for Data Centers.*

Tschudi, W. (2007). *Data Center Economizer Contamination and Humidity Study. LBNL/Pacific Gas & Electric.* http://tinyurl.com/3y4es7 (or http://hightech.lbl.gov/documents/DATA_CENTERS/EconomizerDemoReportMarch13.pdf).

Version 1.0 Release Date: 21 October 2008