MORRISON HERSHFIELD

Optimizing Existing Data Center Chilled Water Plants

January 25, 2013

This **White Paper** describes factors to be considered when optimizing an existing data center chilled water plant.

Authored by:
Michael McKenna, PE
Principal
Senior Mechanical Engineer

For more information contact:

Morrison Hershfield
1500 S Edgewood Street | Baltimore, MD 21227
Phone: 410-525-0010

www.morrisonhershfield.com

Optimizing Existing Data Center Chilled Water Plants

Data centers have always emphasized system reliability as a primary goal. However, recent emphasis on the energy consumption of data centers has renewed a focus on cooling efficiency. As system efficiency and ultra-low Power Utilization Effectiveness (PUE) numbers assume a higher profile, particularly in new purpose-built data centers, the PUE of legacy data centers can sometimes be dismissed or overlooked. But even among data centers a decade old or more, substantial improvements in operating efficiency can be made, often with little or no new equipment.

The primary energy users in chilled water systems are roughly, in decreasing order: the chiller; the cooling tower fan; the chilled water pumps (whether primary only, primary/secondary or variable primary); and the condenser water pumps. The fan energy of the air distribution systems is another component which needs to be considered when trying to optimize total net cooling system efficiency instead of just concentrating on wire-to-water plant efficiency, or least effective of all, just chiller operating efficiency. All the components combine and work in concert.

Establish The Right Temperature

The temperature trend in data centers is warmer. In 2008, ASHRAE Committee TC9.9 raised the maximum recommended temperature from 77°F (25°C) to 80.6°F (27°C). The expanded allowable range for most equipment increased to 89.6°F (32°C) with new equipment being designed by manufacturers for even higher tolerances.

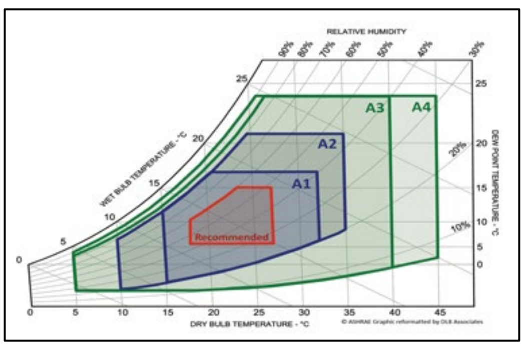

Many legacy data centers were designed for space (actually return air) temperatures of 72° F or less. With the newer, expanded temperature ranges and more widespread implementation of hot aisle/cold aisle and other best practices, the space setpoints in many existing computer rooms can be raised with little to no effect on IT equipment reliability. The days of the 'meat locker' computer room are over. Implementation of even partial containment, either in hot or cold aisles can also allow for tighter space tolerances.

When modifying established space conditions, Computational Fluid Dynamic (CFD) modeling can be used to ensure that ASHRAE recommended temperatures are maintained in all areas. Often, a CFD analysis can justify more aggressive setpoints by confirming that proper airflow and IT cabinet inlet temperatures can be met in all areas. CFD models can also determine which containment strategies can be successfully retrofitted to an existing equipment configuration.

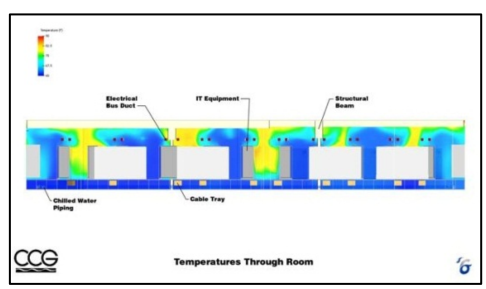

Reset Chilled Water Temperatures

A common practice in commercial applications such as office buildings and schools is to reset the chilled water supply temperature upwards during periods of lower loads such as during the evenings or during the winter. Data centers differ from commercial applications in that the cooling load tends to remain fairly constant throughout both the day and the year. Weather and climate conditions have only minor impact on data center cooling loads. For this reason, the chilled water supply temperature is usually kept constant. If the space temperature of the data center is allowed to increase, the chilled water temperature could also be raised. If the data center space temperature is increased, the chilled water temperature can be raised by the same amount with no loss in Computer Room Air Handler (CRAH) capacity. For example, if the space temperature is raised from 72°F to 77°F, the chilled water temperature could be increased from 45°F to 50°F with no loss of critical load capability.

Raising the supply chilled water temperature of a chiller plant has two effects. It tends to increase the efficiency of the chillers and to increase their capacity as well. Chiller efficiency increases about 2% for each 1°F rise in setpoint.

Balance Capacities

Three primary factors affect the capacity of a chilled water CRAH, entering air temperature, entering chilled water temperature, and chilled water flow rate. As computer room densities vary, the selection of setpoints can be adjusted to avoid stranded capacity while still meeting space requirements. If the critical load has grown higher than the original design basis, raising the space temperature setpoint can squeeze additional capacity out of existing CRAH units. If the space is under-populated or if the cooling load has decreased either through migration or virtualization, the chilled water supply temperature can be increased to take advantage of more optimum operating conditions.

The key is that space temperature changes and changes in chilled water temperature must be balanced against each other.

Reduce Chilled Water Flows

In many large distributed chilled water plants, particularly primary/secondary systems serving multiple types of occupancies, low chilled water return temperatures literally rob the chiller plant of capacity. If in a system designed for a 12°F difference between the return and supply water temperatures, the actual system temperature increase is only 9°F, the chillers are limited to 75% of their design capacity.

If the space setpoint can be increased, the design load can be met with the same equipment at a higher system temperature difference. Since pump energy follows the cube law pump relationship, even minor decreases in required chilled water flow can result in substantial pumping energy savings. Adjusting chilled water flow rates affects the cooling capacity of both the chiller and the terminal cooling devices (such as CRAHs), so care must be taken to keep equipment capacities in line with load.

Operate Chillers At Their Sweet Spot

Chillers are often selected, designed, and tested at their peak loading conditions, but even in data center operations, they operate much of the year at reduced loads or off-design conditions. The peak efficiency point for most chillers is at approximately 75% to 80% of their design capacity. For chillers equipped with variable speed drives, this optimum point can be even lower.

In large chiller plants, it can be beneficial to operate more chillers than required to meet load. All that is necessary is for the efficiency savings of the combined chiller operation to be greater than the incremental pumping energy of the additional chilled water and condenser water pumps. In variable primary systems, this almost always pays off as running an additional chiller line also results in lowered pump heads through the chillers. Even in primary/secondary systems, this can be a net savings under the proper conditions.

Reset Condenser Water Temperature

Besides load, the other factor that most affects chiller efficiency is the entering condenser water temperature. Within limits, the cooler the condenser water temperature, the more efficiently the chiller operates. The lowest condenser water temperature possible from a cooling tower is based on the outside wet bulb temperature. A condenser water temperature reset schedule should account for the fact that the cooling tower approach increases as the wet bulb temperature decreases. Condenser water temperature must also maintain sufficient refrigerant head pressure at the chiller. This is particularly important if chilled water supply setpoints have been reset upwards.

Operate Additional Towers

One strategy that works well in conjunction with an aggressive condenser water reset schedule is to operate more towers than chillers. If the condenser water system is properly manifolded, the condenser water flow for the chillers can be split among multiple towers provided the minimum flow to each tower is met. If necessary, weir dams can be added to the tower hot water basins to allow for lowered flow conditions. Often this can result in lower entering condenser water temperatures which increases chiller efficiency at a net decrease in tower fan energy.

Economizers

If the goal is to lower PUE, the most effective way once all other systems have been optimized is to increase the amount of free cooling hours available to the system. For chiller plants with economizers, raising the chilled water supply temperature increases the number of hours dramatically. Many older systems were designed as full economizers where the system is either on economizer or not. Systems like this are poorly suited for data centers as they have several operational issues which adversely affect mission critical operations. In these systems the condenser water temperature must be driven down to below the chilled water supply setpoint before change-over can be initiated. This limits the number of free-cooling hours available and requires a service interruption during the changeover.

With an integrated economizer, partial free-cooling is available as soon as the condenser water temperature is below the chilled water return temperature, dramatically increasing the number of hours the system is available. Provided there are measures to maintain the head pressure of the chillers, the changeover can be done seamlessly and while he system is operating. More importantly, backing out of economizer operation can be done with the chiller restarted at low load and then gradually re-establishing the load under mechanical cooling.

Older full-economizer systems can be retrofitted to integrated economizers with modifications to the piping and controls. If practical, this is the single most effective conservation measure available, particularly in cold climates.

Summary

Optimization of an existing data center chilled water plant requires assessing a myriad of factors including reviewing system setpoints, reviewing current and projected system loads, maintaining accurate system trend logs, analyzing climate data, and validating operating sequences. It requires staff and management dedicated to the consistent monitoring and review of the system. Properly optimized, a legacy data center can be kept competitive with newer facilities at little or new additional capital cost.
Because adjusting one set of parameters often has a ripple effect, a holistic approach must be maintained to avoid negating the savings of a strategy with increased energy consumption in another component.