
Waterside Economizing in Data Centers: Design and Control Considerations

Jeff Stein, PE
Member ASHRAE

ABSTRACT

Free cooling was not common in data centers in the past for a variety of reasons including the philosophy that data center cooling should be designed for maximum reliability and not for energy efficiency. Recently times have changed. Energy and sustainability are more important to many data center owners now and sophisticated owners and designers know that free cooling can provide a good return on investment while still maintaining adequate reliability. Many data centers are being designed or retrofitted with airside economizers, waterside economizers, and even "wet bulb" economizers (direct evaporative coolers). This paper briefly compares airside and waterside economizers (CRAC and chiller plant) and then focuses on design and control considerations for chiller plant waterside economizers serving data centers.

AIR VERSUS WATER

An airside economizer will generally be more energy efficient than a waterside economizer, if space humidity is not required to be tightly controlled. Airside free cooling is implemented at each air handler and thus the amount of free cooling can be maximized at each air handler. The energy savings of a chiller plant economizer, on the other hand, can be reduced or eliminated for the entire installation by a single "rogue" zone. Furthermore, waterside economizers (WSE) require pump and tower energy and typically three steps of heat transfer (e.g. ambient air to condenser water, to chilled water, to supply air). Airside economizers do not require pump or tower energy or any steps of heat transfer.

If humidity is required to be tightly controlled and humidification is provided by steam, or infrared humidifiers then the

energy savings from airside economizing are significantly reduced but can still be substantial¹. However, if adiabatic humidification (direct evaporation) is used then the savings can be even greater than airside economizing without humidification². Furthermore, more and more data center managers are realizing that humidity control has little if any impact on data center operations and therefore more and more data centers are relaxing or eliminating humidity control.

On the other hand, air economizers have some concerns that waterside economizers do not share including introducing unwanted particulates, and gaseous contaminants. Lawrence Berkeley National Labs recently measured particulate concentrations in several data centers with and without airside economizers. They found that with proper filtration particulate concentrations in data centers with airside economizers are not necessarily higher than in data centers without airside economizers and can easily be maintained below the most conservative standards³. Recent research projects by Microsoft and Intel have also showed that airside economizers do not affect data center reliability^{4,5}. More research is needed but so far there is no research showing that airside economizers compromise data center reliability.

Air side economizers will generally be more expensive than waterside economizers. Furthermore, the deal killer for airside economizers in many installations is the extra space required for outside air intake and exhaust louvers, dampers and ducts, particularly when the data center is buried at the bottom of an office building. In many cases the only option for economizing is a waterside economizer.

TYPES OF WATERSIDE ECONOMIZERS

There are two basic types of waterside economizers: CRAC (computer room air conditioner) unit economizers and

Jeff Stein is a principal at Taylor Engineering in Alameda, CA.

chiller plant economizers. A typical water-side economizer in a CRAC unit is a water-cooled direct expansion CRAC unit with a water coil upstream of the DX coil. The water coil can be served by chilled water or by condenser water (making it a waterside economizer coil). The free cooling provided by the waterside economizer coil is largely offset by the added pressure drop of the extra coil (e.g. selections from one major CRAC manufacturer indicate that the economizer coil increases fan brake horsepower by 40%).

A chiller plant economizer consists of a heat exchanger that allows condenser water to cool the chilled water directly. Such an economizer would be used in a data center where the data floor is served by chilled water computer room air handlers (CRAHs) rather than DX CRAC units. The CRAH units would only have the pressure drop of one coil, not two. Perhaps the biggest difference, however, between a CRAC unit economizer and a chiller plant economizer is that packaged DX CRAC units usually have constant speed fans while chilled water CRAH units can have variable speed drives (VSDs) on the supply fans. Fan energy accounts for about half of data center cooling energy and properly controlled VSD CRAH fans can dramatically reduce fan energy. The rest of this paper is devoted to plant economizers.

INTEGRATED VERSUS NON-INTEGRATED

A chiller plant waterside economizer can be integrated, meaning the economizer can meet all or some of the load while the chiller meets the rest of the load, or non-integrated, meaning the economizer can only operate when it can meet the entire load. Figure 1 shows an integrated waterside economizer in a primary/secondary chiller plant with two chillers. The heat exchanger is in parallel with the chillers on the condenser water side and in series with the chillers on the chilled water return side.

When the outdoor air wet-bulb temperature is low, the cooling tower fans are run at high speed to produce cold condenser water (e.g. in the 40's or 50's). This water is pumped through the heat exchanger where it cools chilled water to within a couple degrees of the tower water temperature. If the economizer cannot bring the chilled water temperature down to the supply temperature setpoint then the chiller(s) pick up the remaining load and bring the water leaving the plant down to setpoint.

The HX is located in the secondary loop on the return side just before the common leg. Locating the HX in the secondary loop, rather than the primary loop is important because it allows the heat exchanger to see the warmest possible water which maximizes the hours when the economizer can operate. The secondary loop is also better than the primary loop because it allows the primary pumps to be shut off when the economizer can handle the entire load.

Figure 2 shows an integrated economizer in a primary-only chiller plant. The HX is located on the load side of the common leg so that it sees the warmest return water. If it were on the plant side of the common leg then it could see a blend

of return water from the loads and cold supply water from the common leg [the valve in the common leg is modulated to maintain the required minimum flow rate through the chiller(s)].

Figure 3 shows a non-integrated waterside economizer piped in parallel with the chillers on both the condenser and chilled water sides. If the economizer cannot meet the entire load then it must be shut off. Otherwise, the relatively warm economizer leaving water will be mixed with the cold chiller leaving water and the plant leaving water temperature will be above setpoint. The chillers might be able to compensate by over cooling their leaving water but with chiller(s) operating the primary loop flow will likely exceed the secondary flow which will result in colder primary return water which will reduce or eliminate the economizer capacity. Non-integrated economizers do not save as much energy as integrated economizers and there is no real advantage of non-integrated over integrated. Non-integrated economizers do not meet the economizer requirement in California Title 24 or in ASHRAE 90.1-2007 for some climate zones.

HEAD PRESSURE CONTROL

A critical feature of a chilled water plant with a waterside economizer is head pressure control on all chillers, air conditioners and heat pumps served by the condenser water. To gain benefit from the economizer the cooling towers must produce very cold condenser water but if the condenser water leaving a chiller is too low then the chiller could trip on low head pressure. Most chillers require a minimum head pressure (the pressure differential between the condenser and the evaporator) in order to insure adequate refrigerant flow and adequate oil flow to lubricate the compressor and provide the seal between the rotor or impeller and its housing. Without head pressure control most chillers cannot operate with entering condenser water below about 70°F during normal operation or below about 55°F at start up.

Adequate head pressure can be maintained by modulating the condenser water flow to maintain the head pressure at or above minimum head pressure setpoint. Most chiller controllers have an analog output signal that allow the chiller to directly control a modulating condenser water valve or condenser water pump in order to maintain adequate head pressure for that chiller. If the chiller does not have a head pressure control output, a stand alone controller can easily be added to measure head pressure (or leaving condenser water temperature) and modulate the condenser water valve. As with all control loops, the PID loop modulating the valve/pump must be tuned to avoid hunting.

Some chillers are specifically designed to handle cold entering condenser water (e.g. 55°F during normal operation and 40°F during start up) but even these chillers should have head pressure control to allow the economizer to maximize free cooling. If the chiller has a condenser water flow switch it is a good idea to jumper out the flow switch since modulating

the condenser water flow is likely to cause the flow switch to trip. Condenser flow switches are generally unnecessary because chillers have high head safeties that will protect them from damage if there is insufficient condenser water flow.

Head pressure control on water-cooled air conditions and heat pumps can be achieved in two ways. One option is to use modulating condenser water control valves, like a chiller. A better option is to add a heat exchanger to the condenser water loop serving the heat pumps and control the temperature on the closed loop side of the heat exchanger by modulating the flow on the open loop side. This option has the advantage of protecting the heat pump condenser coils from the particulate matter that accumulates in an open loop condenser water system. Of course another option is to use chilled water fan coils instead of heat pumps. This is likely to be less expensive and more efficient in a building with a data center and a waterside economizer.

LIFECYCLE COST ANALYSIS

Waterside economizers are most cost effective in large data centers in cold or very dry climates where the wet-bulb temperature is often below 50°F but even in moderate climates or relatively small data centers waterside economizers can still “pencil out” in large part due to the 24/7 nature of data center loads. Figure 4 shows annual simulation results for a chilled water plant serving a data center in four diverse climates. Even in a relatively warm climate like El Paso a waterside economizer can reduce total HVAC energy by 30%.

A waterside economizer was bid as an add alternate on two recent office/data center projects for which the author’s firm designed the mechanical system. Project A consisted of a chiller plant serving a 500,000 ft² office building and a 2,000 ton data center. Project B consisted of a plant serving a 150,000 ft² office building and a 110 ton data center. In both cases the office air handlers had airside economizers so the

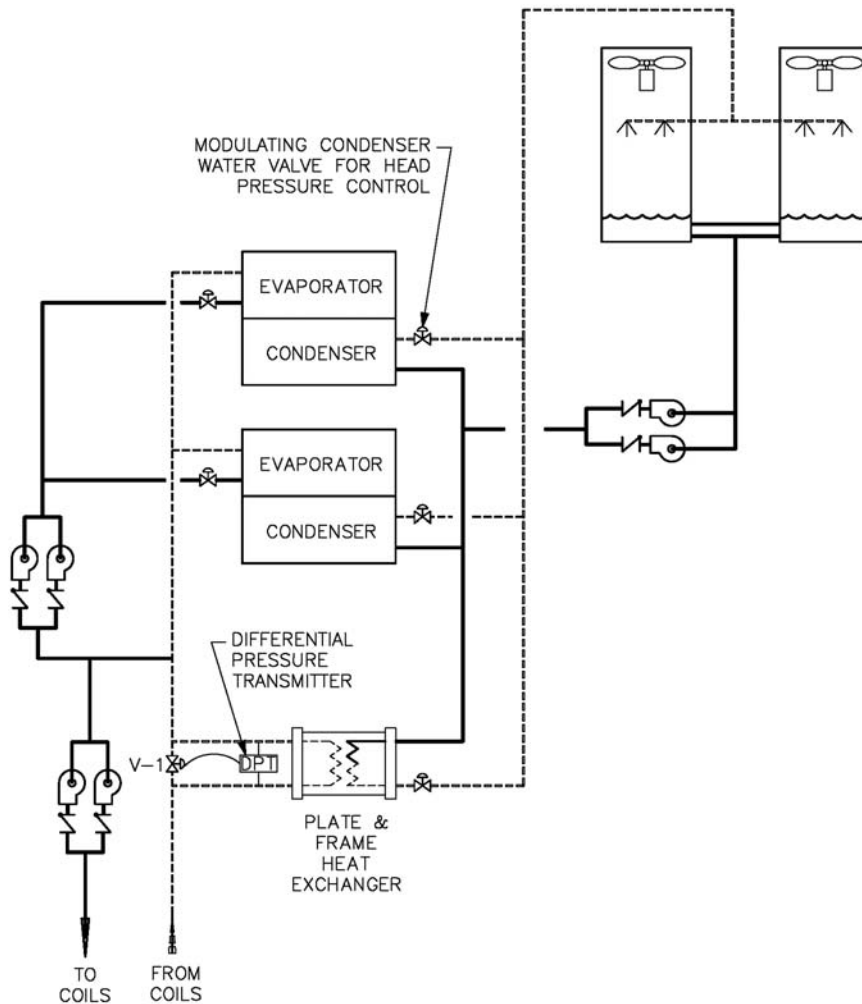


Figure 1 Integrated economizer—primary/secondary.

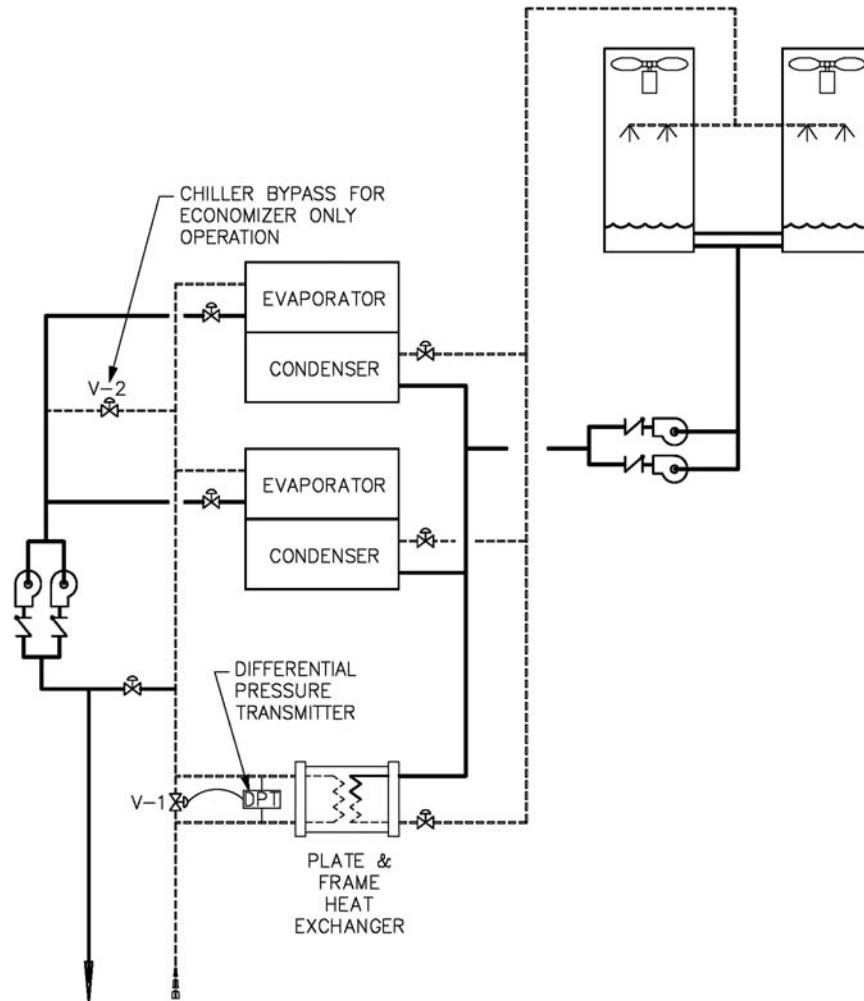


Figure 2 Integrated economizer—primary only.

waterside economizers were really only serving the data center CRAH units. Detailed lifecycle cost analyses were performed in both cases to determine the cost effectiveness of the waterside economizer. The simple payback for Project A was less than one year and the simple payback for Project B was about 5 to 10 years. In both cases, the owner elected to proceed with the waterside economizer.

In order to accurately calculate the energy savings from the economizer the plant must be simulated with an annual simulation program which allows accurate modeling of waterside economizers and water coil response to chilled water temperature setpoint reset. The DOE-2.2 DesignDay simulation engine has these features. It is available for free from www.DOE2.com. One of the most sensitive assumptions in the model is the expected data center load as a percentage of design load. A data center might be designed for 100 Watts/ft² but the actual load is not likely to match the design load, at least not initially. If the data center is too lightly loaded then the savings potential is too small to justify the economizer but

if the data center is fully loaded then the chilled water temperature reset is limited and the economizer will not be able to run as many hours. It is a good idea therefore to run several parametrics at different data center load levels (e.g. 25%, 50%, 75% and 100%) to see how load affects the payback.

HX SELECTION

There isn't necessarily a right way to select a waterside economizer heat exchanger (HX). The larger the heat exchanger the greater the savings but there is obviously a diminishing return on investment. The main defining variables with a heat exchanger are the approach and the capacity. The approach is the difference between the entering condenser water and the leaving chilled water. The smaller the approach the bigger the HX. Ideally one would iterate on the HX size and re-run the HX lifecycle cost analysis for each selection to determine the lifecycle cost optimum. Unfortunately most heat exchanger manufacturers only allow their

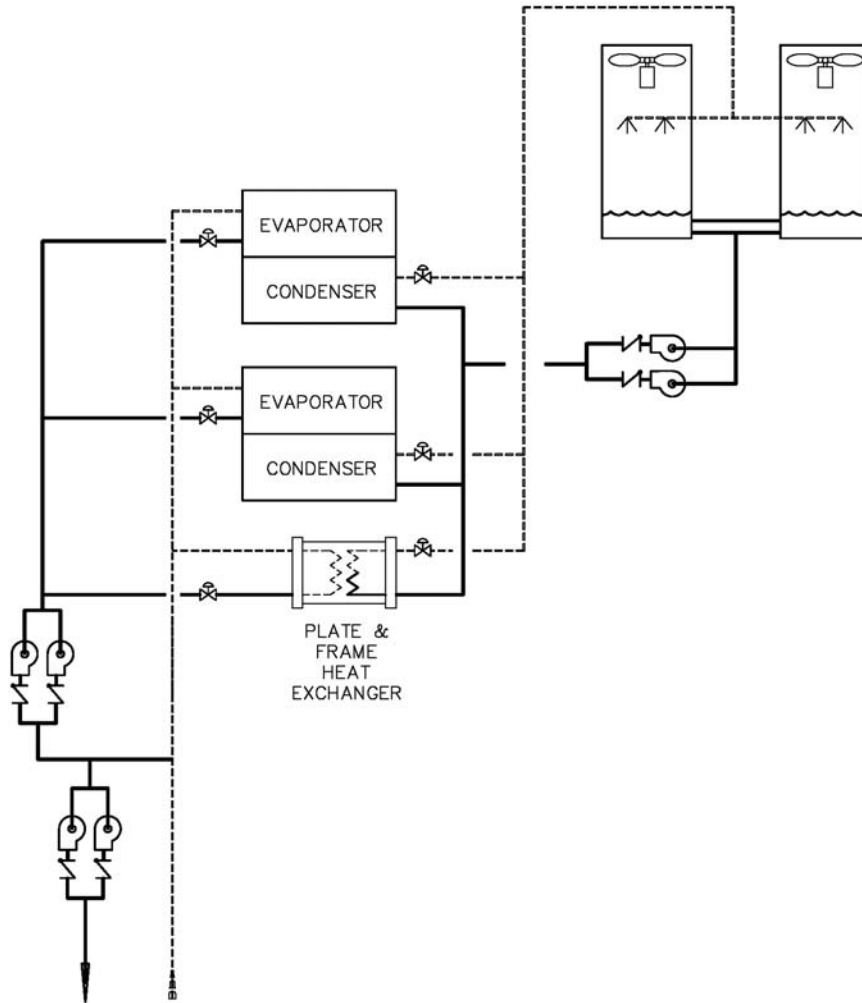


Figure 3 Non-integrated economizer.

sales representatives (and not consulting engineers) to run their selection software, so iterating on heat exchanger selection parameters can be cumbersome.

The following procedure is one way to select a heat exchanger:

1. CW Flow—Pick a condenser water flow and dP to match the number of chiller(s) needed for the expected data center load. The HX will typically be served by the same condenser water pumps as the chillers so it makes sense to select a HX with a condenser water flow and pressure drop similar to one chiller, or to the sum of multiple chillers. At one extreme you could size the heat exchanger with the same condenser water flow rate as all chillers combined but this is not likely to be cost effective if the plant also serves air handlers with airside economizers (e.g. offices) or if the data center is not likely to be fully loaded or if there are redundant chillers.

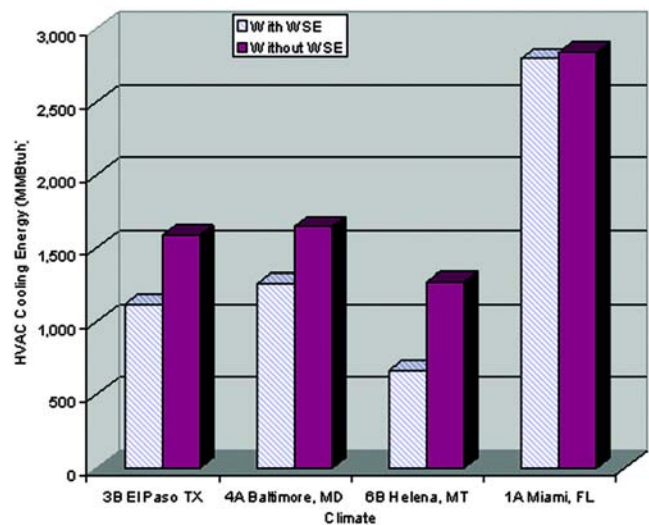


Figure 4 Waterside economizer savings in 4 climates.

2. CW ΔT —Assume the range (ΔT) on the condenser water side is equal to the cooling tower design range since this is about the range the cooling tower can produce for the design flow. At this point the design capacity (Btu/hr) is defined.
3. CW/CHW ΔP —Limit the maximum HX ΔP on the CW side to that across the chiller condensers. This will effectively limit the ΔP on the chilled water side as well for plate & frame heat exchangers.
4. CHWST/CHWRT—Because CHW supply temperature will be reset, assume the leaving CHW temperature will be 5°F warmer than the design temperature. For example, with a 42°F design CHWST, assume 47°F CHWST for the HX selection. Using the CRAH selection software, determine what CHW return temperature would result assuming the coil has design air conditions, design CHW flow rate, but 5°F warmer CHWST. (As a rule of thumb the CHW return temperature goes up about 0.5°F for every 1.0°F increase in CHW supply temperature, so a coil with a 20°F ΔT at 42 CHWST might have a 17.5°F ΔT at 47°F.).
5. Approach—Select the entering CW temperature to be about 3°F below the CHWST. A smaller approach will improve performance but cost will rapidly increase. The leaving CW temperature is then calculated from the
6. CHW Flow—Chilled water flow is calculated to match the heat transfer on the condenser water side.

Example Selection

1. Suppose the plant has (3) 775 ton chillers and an expected data center load of 1000 tons. The chillers have CW flow rates of 1760 GPM each. The HX would be selected to have the same CW flow rate as two chillers, i.e. 3520 GPM.
2. The cooling towers are designed for 12°F ΔT so we assume the HX CW ΔT is also 12°F. The design HX capacity is then $500 \times 3520 \times 12/12,000 = 1760$ tons.
3. The chillers have a design CW ΔP of 5 psi. Thus the HX CW and CHW ΔP would be specified to be ≤ 5 psi.
4. The CRAH units are selected at 43°F entering CHW and 20°F ΔT . At 48°F entering CHW the ΔT is 17.5°F and the leaving CHW is 65.5°F.
5. At a 3°F approach the entering CW temperature is $48 - 3 = 45$ °F. The leaving CW temperature is then $45 + 12 = 57$ °F.
6. The CHW flow is calculated from 1760 tons and 17.5°F ΔT to be 2414 GPM.

	Chilled Water (Hot) Side	Condenser Water (Cold) Side
Flow (GPM)	2414	3520
Entering Water (°F)	65.5	45
Leaving Water (°F)	48	57
Max. Pressure drop (psi)	≤ 5	≤ 5

DESIGNING FOR HIGH CHILLED WATER RETURN TEMPERATURE

To maximize the benefit of the waterside economizer the entire system should be designed and operated to achieve a high chilled water return temperature. The higher the return water temperature the more hours the economizer will operate and the greater its cooling output. Here are some design considerations:

- Use 2-way chilled water coil control valves (no 3-way valves).
- Select all chilled water coils for high design chilled water ΔT (e.g. 20°F).
- Generously size all coils and air handlers—IDF closet fan coils, for example, should be oversized to prevent them from limiting the chilled water temperature set-point reset. (Note that oversizing air handlers with variable speed fans improves energy efficiency due to lower coil velocities and cube law fan savings at actual operating conditions.)
- Run all redundant CRAH units—Data centers will typically have one or more spare CRAH units to meet the redundancy requirements. Rather than allowing redundant units to sit idle, all units should be run at lower fan speed. This reduces the load on each coil and allows greater chilled water temperature reset. If a CRAH unit fails the others speed up to compensate.
- Design for high airside ΔT —Many data centers operate at very low airside ΔT at the CRAH units. The higher the return air temperature the warmer the return water temperature back to the plant. Low airside ΔT has two main causes: short circuiting of CRAH supply air directly back to the CRAH unit and computer servers with high air flow rates (low ΔT). The best way to limit short circuiting is to effectively duct the return air from the hot aisles back to the CRAHs using ceiling return air plenums and/or hot or cold aisle containment systems. The IT department should also be brought into the design process and educated on the benefits of selecting servers with high ΔT (low air flows) and 2-speed or variable speed server fans.

DESIGNING FOR LOW CONDENSER WATER SUPPLY TEMPERATURE

Oversizing the towers by selecting a low design approach will increase economizer hours but excessive oversizing is not practical or cost-effective. A reasonable approach for a tower selection is 5 to 8°F. Another way to achieve low condenser water temperatures is to select or configure the cooling towers for low flow. Most towers can operate at 50% of design flow without causing excessive scaling of the fill and many can operate at 30% flow. Operating as many tower cells as the minimum flow will permit will allow lower condenser water than operating the minimum number of cells. In locations with very good water quality (e.g. San Francisco) the tower manufacturers'

minimum flow can be ignored entirely, with little or no risk of tower scaling.

CONTROL SEQUENCES

Enabling/Disabling the Economizer

When the economizer is disabled the towers are not necessarily producing their coldest water so the economizer is enabled based on the predicted condenser water supply temperature when the towers are sped up. For example: “WSE is enabled if it has been disabled for a minimum off time and $CHWRT > \text{predicted CWST} + \text{deadband}$.” (CHWST is chilled water supply temperature, CHWRT is chilled water return temperature and CWST is condenser water supply temperature). The predicted CWST can be as simple as Wet-bulb + Approach. The achievable approach, however, is a function of the wet-bulb and the load on the tower. Approach goes down with wet-bulb and up as load decreases. A more sophisticated strategy would be to derive a function for the predicted CWST using a tower simulation model⁶. Guessing wrong on the predicted CWST is not really a big deal because the WSE will simply be disabled if it is not providing benefit. The WSE is disabled if it has run for a minimum on time and the chilled water leaving the WSE is not at least 0.5°F colder than the chilled water entering the WSE.

Controlling the Economizer

The WSE is not necessarily sized for the full design CHW flow rate. If valve V-1 (see Figures 1 and 2) is fully closed and CHW flow is near design then there is a chance that the chilled water pumps will not be able to achieve the differential pressure setpoint. Therefore, when the WSE is enabled valve V-1 is modulated to maintain CHW differential pressure across the HX at a fixed setpoint (e.g. 1.5 times the design CHW pressure drop across the HX). V-1 is fully open when the WSE is disabled and the WSE is bypassed.

Tower Speed Control

When the WSE is disabled tower speed is controlled as if the WSE did not exist. When the WSE is enabled and a chiller is still enabled then one option is to run the tower fans at 100% to make the condenser water as cold as possible and maximize the WSE output. With variable speed drives on the tower fans, however, the last bit of tower speed (say from 90% to 100% speed) does little to lower the condenser water temperature but increases tower fan energy significantly. A better sequence when the WSE and chiller are enabled is to control the tower speed to maintain the CWST at or just above the predicted CWST, as described above. If the WSE is operating alone then the tower speed can be modulated to maintain the CHWST at setpoint. At very low loads and/or very low wet-bulbs the tower speed may be low enough to stage off one or more tower cells.

Enabling/Disabling the Chillers

When the WSE is enabled and the CHWRT downstream of the WSE (entering the chillers) approaches the CHWST setpoint then the WSE is carrying the entire load and the chillers can be disabled. Conversely, when the WSE can no longer hold the CHWST at setpoint the lead chiller must be enabled. A deadband should be included to prevent short cycling chillers.

Chilled Water Supply Temperature Setpoint Reset

Both the chilled water supply temperature setpoint and the chilled water loop differential pressure (DP) setpoint should be reset based on coil control valve positions. A single reset control point should be used to control both setpoints. When the plant includes a WSE, the reset should “lead” with pressure to keep the water temperature as high as possible to maximize WSE operation (see Figure 5). The reset sequence must be slow acting. A trim and respond sequence such as the following is easier to tune than a PID loop: “Every 2 minutes, decrease (“trim”) the CHW plant reset point by 1% and if there are more than 3 requests then increase (“respond”) the reset point by 2% for each cooling request, up to a maximum of 5% increase every 2 minutes. A cooling request is generated when any chilled water control valve position exceeds 90% open. All values shall be adjustable for tuning. Note: only valves whose CHW Request Enable/Disable software switch is in the enable position may generate a request.” As the CHW Plant Reset Point increases from 0% to 50% the DP setpoint is first reset from minimum (e.g. 5 psi) to maximum (e.g. 15 psi). From 50% to 100% the CHWST setpoint is reset from maximum (e.g. 55°F) to minimum (e.g. 42°F).

CRAH Unit Control

The CRAH units should be controlled to maximize the return water temperature by minimizing fan speed and maximizing supply air temperature (SAT) setpoint. It is important to understand that unlike a typical comfort cooling application the

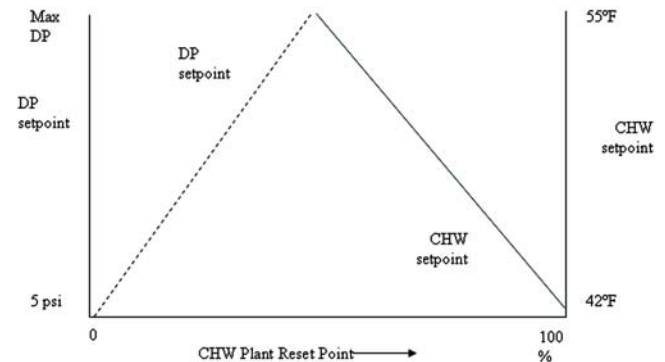


Figure 5 CHWST and DP reset.

only space temperature that matters in a data center is inlet temperature at the servers. ASHRAE recommends inlet temperatures in the range of 68 to 77°F. Air temperatures elsewhere in the data center (e.g. CRAH return temperature) can be over 90°F without any problems. A representative sample of server inlet temperatures should be monitored and maintained at setpoint by modulating the fan speed and SAT setpoint (the CHW valve maintains the supply air temperature at setpoint). One possible sequence for modulating fans speed and SAT setpoint is shown in Figure 6. Note that the 75% fan speed “plateau” in this figure corresponds to the design airflow rate for the data center (with all redundant CRAH units operating).

WET-BULB SENSORS

The wet-bulb sensor is a critical component of a waterside economizer control system. If the sensor reads high then it may never enable the economizer. If it reads low then the economizer will short cycle. Unfortunately commercial grade wet-bulb sensors are generally not very accurate and quickly go out of calibration if not well maintained. An alternative to a wet-bulb sensor is to have the energy management system automatically download the wet-bulb data from a nearby NOAA weather station via the internet every few minutes. At least one of the major HVAC control system vendors has a standard plug-in that accomplishes this task.

COMMISSIONING

The key to a successful WSE installation is commissioning. This includes regularly calibrating the wet-bulb sensor, functionally testing all control sequences and reviewing trend graphs of all HVAC systems to insure the controls are stable and operating as expected. In particular it is critical that the CHW plant reset sequences are tuned and that rogue zones are identified and remediated. The key trim and respond parameters include the number of requests that are ignored and the rate of change. The sequence above calls for 3 requests to be ignored before the reset point is increased. In a large system there will invariably be at least one “rogue” zone so a minimum number of “ignores” is necessary.

The rate of change is the combination of the time between steps and the step size. Both the trim and respond rates need to be slow so that the valve control loops have time to react to changes in temperature and pressure. If the rate is too quick then the reset will cycle excessively and the plant load will also cycle excessively. If the reset is too slow then zones will be starved, the plant will “loose” the load and then see a load spike when it “catches the load”. The respond rate can be a little quicker than the trim rate (to avoid starving zones) but they both need to be slow.

A rogue zone is one that is rarely or never satisfied and therefore is consistently inhibiting the reset. There are many potential reasons for a rogue zone including incorrectly sized equipment, faulty actuators, plugged coils, and unreasonably low supply air temperature or zone temperature setpoints. The control system must have graphics and reports that easily

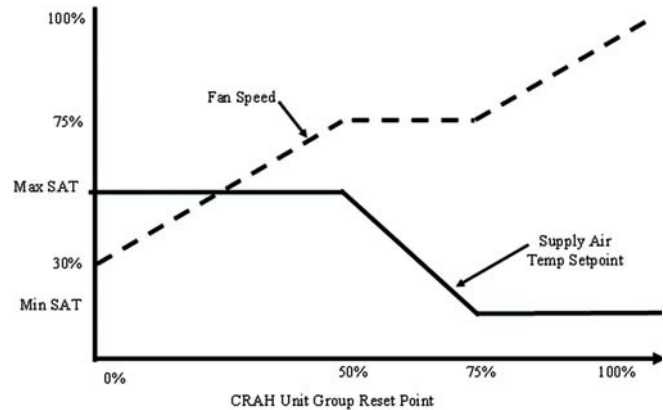


Figure 6 CHWST and DP reset.

allow the commissioning agent and building operators to see which valves are currently and consistently driving the reset sequence.

Figure 7 shows trend data for a typical day for the chilled water plant with a WSE that serves the 500,000 ft² office building and 2,000 ton data center mentioned above. The number of “ignores” is set to 3. At night the CHW DP is reset down to 5 psi and the CHWST is reset up to 55°F. The WSE carries the entire 250 ton load from about 3 A.M. to about 10 A.M. At 8:45 A.M. the number of requests causes the CHW DP to begin resetting. By around 9:30 A.M. the DP is at maximum and the CHWST starts resetting down, reaching 42°F by about 9:45 A.M. By 10:30 A.M. the valves are almost all satisfied and the CHWST begins resetting upwards. The WSE is reenabled around 4 P.M. and operates in integrated mode with one chiller also operating. You can see from the slope of the CHWST setpoint and CHW DP setpoint trend lines that the trim rate is slower than the respond rate.

CONCLUSION

Airside economizers are generally more efficient and less expensive than waterside economizers. When an airside economizer is not practical or desirable a waterside economizer can be a very cost effective way to improve data center efficiency without compromising reliability. To achieve the potential energy savings of a waterside economizer it should be integrated as opposed to non-integrated. It is also important that the entire CHW system be designed and operated to achieve high chilled water return temperatures. For example, CHW coils should be selected for high ΔT (e.g. 20°F); hot or cold aisle containment should be used and CRAH unit fan speed and supply air temperature setpoint should be reset to maintain server inlet temperatures at setpoint. The system must also be thoroughly commissioned. In particular the chilled water supply temperature reset sequence must be tuned and any rogue zones identified and remediated.




Figure 7 WSE trend graph.

REFERENCES

- ¹Sorell, V. 2007. "OA economizers for data centers." *ASHRAE Journal* December 2007.
- ²Scofield, M. and Weaver, T., 2008. "Data Center Cooling: Using Wet-Bulb Economizers" *ASHRAE Journal* August 2008.
- ³Tschudi, W. 2007. "Data Center Economizer Contamination and Humidity Study." LBNL/Pacific Gas & Electric.
- ⁴http://hightech.lbl.gov/documents/DATA_CENTERS/EconomizerDemoReportMarch13.pdf.
- ⁵<http://www.datacenterknowledge.com/archives/2008/09/22/new-from-microsoft-data-centers-in-tents/>.
- ⁶<http://www.datacenterknowledge.com/archives/2008/09/18/intel-servers-do-fine-with-outside-air/>.
- ⁶Hydeman, M. and Zhou, G., 2007. "Optimizing Chilled Water Plant Control", *ASHRAE Journal* June 2007.