BICSI news MAGAZINE

PLUS

- Industrial-grade Infrastructure and Equipment
- I Say Bonding, You Say Grounding
- Data Center Containment

The largest savings opportunity from containment derives from the higher temperature ranges that provide access to more free cooling hours.

EMBRACING THE BENEFITS of Data Center Containment

by Ian Seaton

Ian Seaton has more than 30 years of mechanical and electro-mechanical product and application development experience. For the past 13 years, he has spearheaded Chatsworth Products' data center thermal management initiatives. He is the working group leader for the rack and cabinet section of BICSI-002-2010, *Data Center Design and Implementation Best Practices,* and also served on the mechanical working group for that standard. He can be reached at iseaton@chatsworth.com.

Data center airflow containment is the most cost-effective and least disruptive path toward maximizing efficiency gains and reducing data center operating costs. However, there is one slight paradigm shift that, while not necessary, can significantly help with embracing data center containment—the data center can be viewed, not as a room full of computers, but as an industrial process. As such a controlled industrial process, the containment data center provides the basis for seven significant benefits:

- Elimination of hot spots
- Support for significantly higher cabinet power and heat density
- Elimination of surplus supply air volume with resultant fan energy savings
- Improvement of cooling unit coil efficiency for water cooled systems
- Improvement of chiller plant efficiency
- Increase in available free-cooling hours
- Greater flexibility for architectural design elements

Historically, we have cooled our data centers in essentially the same manner we have cooled our homes and offices. This has delayed our perception of data centers as an industrial process rather than as a room full of computers. At home, we are familiar with rooms in the immediate vicinity of the thermostat that are relatively comfortable, and other rooms that may be a little warmer than comfortable in the summer or cooler than comfortable in the winter. In some cases, people also keep heaters running under their desks in both winter and summer. While most of us accept

Electronic reprint with permission from *BICSI News Magazine*-July/August 2012 Issue

these inefficiencies on a day-to-day basis and adapt our behavior to accommodate them, the stakes take on a whole different proportion in data centers.

For example, a five-ton air conditioner that easily keeps 185.8 square meter (m^2 [2000 square-foot (ft^2)]) home comfortable (and even over-cooled) in central Texas or Florida, is inadequate to cool one data center cabinet with three or four blade servers. When that same philosophy is applied to a data center, with set points typically around 40° C (72° F), there are two undesirable results—the temperature variations in the room are even more extreme than what we see in our homes and offices, and we end up with hot spots. Combined, that can lead to one third of a data center's entire operating budget being spent to cool air that is likely already cool enough to cool a contained data center for free.

For the purposes of this article, the following efficiency claims will be based on supplying air to information technology (IT) equipment within the ASHRAE TC9.9 recommended environmental limits of 18° C (64.4° F) to 27° C (80.6° F).[1]

Goodbye Hot Spots, Hello Higher Heat Densities

The most obvious and important result of implementing a containment system that separates supply and return air in the data center is the elimination of hot spots. Hot spots, or computing equipment that is ingesting air at a temperature above a pre-determined threshold, are simply caused by the presence of warmer air either directly re-circulating to IT equipment intake fans, or indirectly raising the ambient temperature of the room in areas where an insufficient volume of cool air is delivered to meet the needs of the IT equipment. With containment, by definition, that warm air does not have access to IT equipment. No containment is perfect, so there will always be some leakage between the supply and return sides of containment. Nevertheless, as previously noted, with a good containment system the supply temperature variation should not be more than a 3° F range. Even in high densities above 25 kilowatts (kW) per cabinet, that range should not exceed 2.8° C (5° F). With a 23.9° C (75° F) supply set temperature, no servers would exceed the ASHRAE recommended upper threshold.

Containment eliminates the conditions that cause hot spots and supports much higher per-cabinet heat densities than had previously been considered possible to cool with perimeter air cooling. Conventional wisdom has traditionally located the density threshold somewhere between 6 and 10 kW per cabinet, and close-coupled cooling marketing propaganda has placed that threshold as low as 2 kW per cabinet.[2] This threshold is based on the proposition that enough cold air needs to be delivered through a perforated floor tile directly in front of each particular cabinet with an adequate air flow volume to not only supply the requirement of the cabinet IT load, but to also provide enough excess bypass to shield the front of that cabinet from the effects of any re-circulating return (warm) air that may be responding to the low pressure created by the server fans.

While the industry trend is toward designing servers with lower cubic meters per hour (CMH)-to-kW ratios, in the worst case scenario, we can still see servers pulling in up to 272 CMH (160 cubic feet per minute [CFM]) per kW. At this ratio, a 5 kW cabinet would be pulling in 1359 CMH (800 CFM) of air. While 1359 CMH (800 CFM) is realistically deliverable through almost any style perforated floor tile, assuring consistent airflow throughout the data center from every single perforated floor tile becomes problematic and therefore establishes the basis for pegging an air-cooling threshold somewhere around 6 to 10 kW per cabinet, or even lower.

However, if containment results in no warm air in the data center, that targeted point of delivery for the cool air is no longer an issue. If there is a great variation in static pressure zones under the raised floor (e.g., 2039 CMH [1200 CFM] coming through some tiles and 850 CMH [500 CFM] coming through others), the variation will no longer create hot spots. Rather than cabinets being underserved by low flow from proximate floor tiles, they will pull excess supply air from the rest of the room. As long as the room itself is adequately supplied, all servers within will pull in its required cool air volume.

With the maximum density threshold no longer based on the obstacles of delivering the required volume and temperature air to each specific point of need, the threshold becomes defined by the ability to remove the heated air without adversely affecting server fan performance. A well-designed cabinet containment (e.g., closed rear and vertical exhaust duct) system can passively evacuate up

Example of Fan Law Energy Costs for Data Centers			
	Year 1	Year 5	Year 10
Legacy Data Center with 44 CRAHs	$1,270,765	$2,433,823	$3,887,645
Containment with 44 CRAHs	$1,022,075	$1,187,743	$1,398,116
Containment with 24 CRAHs	$621,417	$1,187,086	$1,894,172
Annual CRAH fan energy cost, plus capital cost for 20 extra CRAH units			

Figure 1: Fan Law Energy Cost Example

to 5097 CMH (3000 CFM), which translates to anywhere from 18.75 kW up to 50 kW per cabinet depending on server efficiency (CMH:kW ratio) to significantly higher thresholds in hot aisle containment (HAC) or cold aisle containment (CAC). Increasing density is not merely a shadow trend to be easily avoided by taking up more space and spreading the load—it is both a response to industry trends in computing efficiency as well as a path to total cost of ownership (TCO) reductions by as much as 40 percent.[3]

Out with Waste, In with Savings

In addition to the substantial cost avoidance of compensatory close-coupled cooling solutions and the TCO benefits of increasing density, containment benefits data center efficiency in several important ways, including:

- Minimizing cooling unit fan energy by eliminating wasteful bypass airflow.
- Eliminating the need to over-produce cooling to compensate for recirculation of warm air.

With containment, to avoid waste all chilled air by definition must pass through an IT heat load and conduct some degree of heat transfer before it finds its way back to the cooling source. Since a data center with a 2.0 power usage effectiveness (PUE) likely has anywhere from 200 to 250 percent the cooling capacity required by the IT heat load demand, and a data center with a PUE of 3.0 is likely deploying up to 400 percent of the cooling capacity actually required by the IT heat load, the opportunity for energy savings by eliminating or greatly reducing waste is sizeable.[4] Furthermore, because of fan energy laws, those savings are not directly linear—fan energy savings can be several times greater than fan output savings. This relationship can be described by the following equations:

- $(N_1/N_2) = (Q_1/Q_2)$—indicating a directly proportional relationship between revolutions per minute (RPM) and flow
- $(N_1/N_2)^3 = (P_1/P_2)$—indicating a cubed relationship between RPM and power

In other words, any savings in airflow produces a resultant energy use that is the cube of the airflow. For example, a fan operating at 80 percent capacity would use only 51.2 percent ($.80^3 = .512$) of the energy than at full speed.

The impact of fan energy savings can be illustrated through a two-megawatt case study data center operating at an N+2 redundancy level and requiring 44 computer room air handler (CRAH) units to sufficiently cool a legacy hot aisle/cold aisle space. With variable air volume (VAV) CRAH fans (could be variable frequency drive [VFD] or electronically commutated [EC]), the air handlers were operating at 94 percent capacity, which at $0.10 per kilowatt hour (kWH) is a cost of $290,765 to operate annually. Conversely, with full containment and the resultant elimination for the need to overproduce, those 44 CRAHs were able to operate at 49 percent fan speed and only consume $42,075 in energy annually for an 85.5 percent savings.

Interestingly, if the data center had been designed correctly for containment and the cooling capacity deployed accordingly, 24 CRAH units would have been sufficient. However, those 24 units, running at 90 percent fan capacity, would have drawn $141,417 in energy, or 236 percent more than the 44 CRAH units at 49 percent. While the example illustrates the consequential significance of fan energy savings, it is not intended to suggest the efficacy of spending half a million dollars for 20 unnecessary CRAH units. Nevertheless, Figure 1 indicates that within five years the capital investment would be absorbed, and within 10 years not only would the "unnecessary" CRAH units be paid for, but there would be a $500,000 positive return on investment (ROI). Granted, the example does not take into consideration lost floor space, service contract costs and extra water expenditures, but it still provides dramatic evidence for the value of eliminating surplus supply air production and at least suggests a potential economic path for justifying 2N cooling redundancy.

Efficiency Gains, from Cooling Unit to Chiller Plant

Another benefit derived from containment is the efficiency gains on cooling units from the higher return air temperatures that are contained and

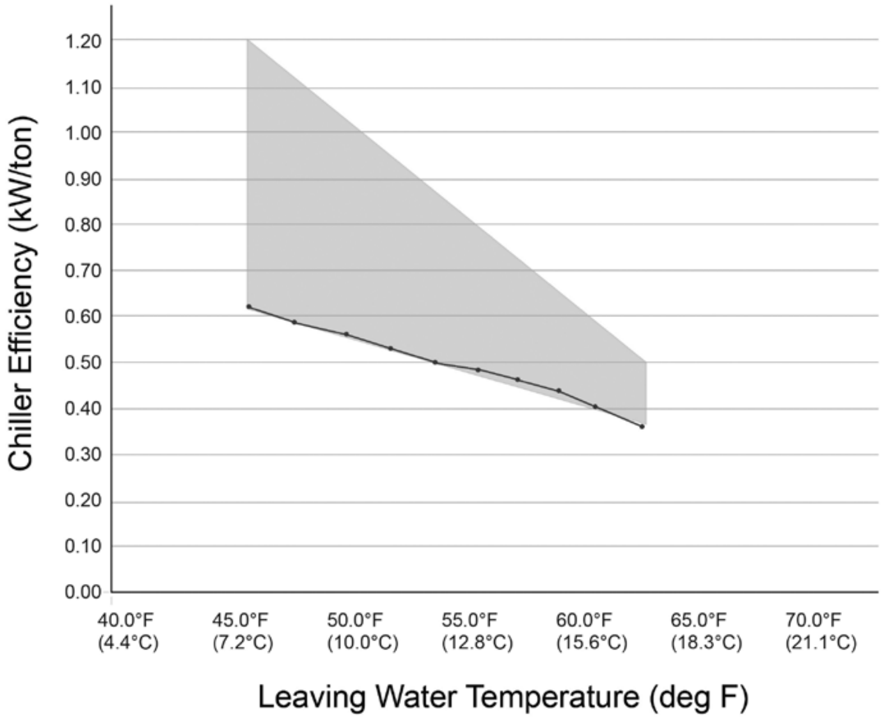

Figure 2: Water Chiller Efficiency

delivered. When cooling is triggered by return air temperature exceeding some set point, the return air is typically dropped by 10° C (18° F) when it is pulled across the cooling coils. That temperature drop typically defines the rated sensible cooling capacity of these cooling units. In other words, a 20-ton cooling unit that has dropped N CMH (N CFM) from 40° C (72° F) to 30° C (54° F) has removed the amount of heat that would have been required to melt about 20 tons of ice in 24 hours. However, with a supply set point, the temperature drop across the coils will vary based on the load.

There is a wide variation in the mechanical efficiency of IT equipment, with airflow requirements ranging from 102 CMH (60 CFM) per kW up to 272 CMH (160 CFM) per kW. This causes inlet versus outlet temperature differentials (ΔT) that range from 10.6° C (19° F) to over 27.8° C (50° F). Note that all these ΔTs are higher than the ΔT rating that establishes the nominal sensible cooling capacities of cooling units. Most blade servers have ΔTs in the 16.7° to 19.4° C (30° to 35° F) range. If that complete ΔT can be captured and returned to the cooling unit, it will result in an approximate 50 percent increase in the sensible cooling capacity of that cooling unit.

There is typically not a significant economic gain to this efficiency improvement, unless it is coupled with any form of economization that delivers partial free cooling whenever the data center return air exceeds the ambient condition, such as series waterside economization, air-to-air heat exchange or indirect evaporative cooling. With these technologies, the higher the return air temperature, the more hours the system will operate without reverting to full refrigerant cooling.

Note that cooling unit efficiency gains from higher ΔTs are restricted to water-cooled air handlers and that containment-captured IT equipment ΔTs will not only be harmful to direct expansion computer room air conditioning (CRAC) units, it will be counter-productive to the degree that higher return air temperatures result in higher supply air temperatures at typically around an 80 percent ratio. Another significant benefit resulting from containment is the chiller plant efficiency gain resulting from the higher supply temperatures supported by containment. While a legacy hot aisle/cold aisle data center is going to cool 40° C (72° F) air down to 30° C (54° F) (plus or minus a degree or two), and therefore require that 23.3 to 25° C (42 to 45° F) water be delivered from the chiller plant, a containment data center can operate with 41.7° C (75° F) supply air and an associated 35.6 to 36.1° C (64 to 65° F) water temperature. The energy savings for the chiller plant can range from 1.5 to 4 percent energy reduction per degree of temperature increase.

In Figure 2, which is derived from the American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) handbook, *Best Practices for Datacom Facility Energy Efficiency*, the bottom edge of the shaded area shows the watts per ton to operate a highly efficient centrifugal chiller at different leaving water temperature (LWT) settings.[5] Based on the ASHRAE chart, a 1 megawatt (MW) data center would save $160,000 per year by raising the LWT from 25° C (45° F) to 36.1° C (65° F), at $0.10 per kWh. Most data center operators are going to be operating somewhere in the shaded area and will not likely get down to less than 400 watts per ton at 36.1° C (65° F). However, their baseline condition is likely to be closer to a kW per ton than 620 watts per ton and therefore their resultant savings from this LWT change will likely be even greater.

More Cooling Hours, for Free

The largest savings opportunity from containment derives from the higher temperature ranges that provide access to more free cooling hours. It does not really matter if the point of demarcation is a 36.1° C (65° F) chiller LWT or a 41.7° C (75° F) data center supply air temperature—the

City	Legacy Hot Aisle/ Cold Aisle	Containment Economization: Waterside	Pumped Refrigerant	Wheel	Air-side	Indirect Evaporative
Los Angeles	0	88	90	96	97	99
Seattle	8	96	95	97	98	100
Washington, D.C.	18	68	71	80	82	86
Phoenix	1	76	41	47	48	82
Denver	28	97	82	87	88	100
San Francisco	0	97	98	99	99	99
Atlanta	6	58	71	74	77	77

Figure 3: Percentage of Annual Hours of Free Cooling Using Containment Combined with Economization

containment-dependent temperatures will support more hours of free cooling. This also applies for waterside economization or any of the various different economization systems benchmarked on air temperature, whether that is dry bulb-based or wet bulb-based.

Figure 3 displays the percent of hours of the year for a sampling of cities for legacy hot aisle/cold aisle data center operating temperatures versus the percent hours of the year for free cooling for five different common types of economization.

Other than going into a 100 percent virtualization implementation that results in disposing of 75 percent of the no longer needed excess servers, there is really no more significant impact on the operating budget of the data center than turning off the chiller for 80 percent or more of the year. It should be further noted that four of the five technologies listed on this chart do not require letting the outside environment into the data center.

Flexibility Without Compromise

Finally, containment makes the data center less dependent on the targeted air delivery for which raised floors have been so instrumental over the years. Because hot air is removed from the data center, it is no longer necessary to deliver enough volume of cold air to the point of use to fend off the incursion of re-circulated warm air. Therefore, cold air can be delivered into the data center from a more flexible variety of sources.

For example, if roof-mounted air side economization or indirect evaporative cooling is being considered, that cool air can be dumped directly down into the data center without incurring the fan energy penalty from pushing that air through ducts to get it under the floor. Another example could be using a fan wall to minimize the pressure head the air movers are pushing against from a wheel cooling cell. With containment, the architecture of the data center can focus on supporting the overall functionality of the space rather than making compromises to accommodate a specific air delivery topology.

Closing Thoughts

Air containment in the data center, whether cold aisle containment or hot aisle containment, or cabinet containment enabled by a vertical exhaust duct and captured return air path, represents an easily executable path to effectively cooling a data center and efficiently minimizing the energy cost for delivering that cooling.[6] By definition, containment prevents hot spots and supports higher densities than conventional wisdom had deemed possible with air cooling. In addition, containment saves air handler fan energy, improves cooling unit efficiency, reduces chiller operating costs and provides access to significantly more free cooling hours. ■

References

1 ASHRAE TC9.9, "2008 ASHRAE Environmental Guidelines for Datacom Equipment – Expanding the Recommended Envelope," 2008

2 Dunlap, Kevin and Rasmussen, Neil, "The Advantages of Row and Rack-Oriented Cooling Architectures for Data Centers," APC White Paper #130, 2006

3 Patterson, M.K., Costello, D.G., Grimm, P.F. and Loeffler, M., "Data Center TCO: A Comparison of High-Density and Low-Density Spaces," Intel White Paper submitted to Thermes 2007 for publication, Jan. 2007.

4 Malone, Chris, "Efficiency at Any Scale: Facility Best Practices to Drive High Efficiency," keynote address at 2011 Uptime Symposium, May, 2011

5 ASHRAE, "Best Practices for Datacom Facility Energy Efficiency," second edition, 2009, p.41.

6 Weidmann, Rainer and Leberecht, Markus, "Data Center 2020: Hot-Aisle and Cold-Aisle Containment Efficiencies Reveal No Significant Differences," DataCenter 2020 whitepaper, 2011.