

Hot Air Isolation Cools High-Density Data Centers

By: Ian Seaton, Technology Marketing Manager, Chatsworth Products, Inc.

Contrary to some beliefs, air is quite capable of dissipating heat loads in excess of 30 kW per rack, and can do so quite economically. This is great news to the 92% of data center managers who have not yet tried liquid cooling and the 65% who said they would never use liquid cooling in their data centers, according to a SearchDataCenter.com survey.¹

The key principle in effectively using air to cool high-density data centers is to achieve as complete as possible isolation between the chilled supply air and the heated return air. This guiding principle is the reason for arranging data centers in a hot aisle/cold aisle layout, which is supported by all reputable standards and is cited as a best practice in recent studies by the Lawrence Berkeley National Laboratory² and Intel³.

Hot air isolation can be accomplished using various practices, including cabinets with ducted exhausts directing heated air into a suspended ceiling return plenum, cabinets direct-ducted back to the cooling unit, enclosed hot aisles with exhaust air ducted out of the data center space (direct-ducted to the AC, ducted into a suspended ceiling return plenum or ducted out of the building), as well as the partial, although more complete than legacy hot aisle-cold aisle. In this method isolation is accomplished by evacuating hot aisles into a suspended ceiling return plenum through ceiling grates or routing exhaust air through cabinet chimney ducts into the space above the equipment in a room with much higher than normal slab-to-slab height. In both of these latter partial isolation strategies, the effectiveness of the isolation can be improved on an access floor by extending the perforated tiles beyond the end of the row of cabinets to help block return air re-circulation around the ends of the rows.



Example of a cabinet with a ducted exhaust directing heated air into the ceiling plenum.

The cabinets in these isolation topologies actually serve as that barrier in the room that secure the separation between supply air and return air, and therefore need to include all elements of best practices, such as blanking panels in all unused rack-mount spaces to seal off re-circulation around the perimeters of the equipment mounting areas, and sealing off cable access tile cut-outs. Fans are sometimes used to route the return air through and out of this barrier system, but their use can be avoided, along with the extra power consumption, heat, noise and potential failure points, by engineering spaces that take advantage of the basic physics of the Bernoulli equation. In particular, the Venturi Effect can create high velocity, low pressure areas in the rear of the cabinet and the Vena Contracta Effect can produce low pressure areas at the apex of exhaust

ducts. Both of these “Effects” can therefore work together to effectively remove heat from cabinets without additional fans.

High densities can be achieved with this degree of isolation because it frees the individual cabinet heat load from its dependency relationship on the immediately adjacent perforated floor tile delivery of chilled air. The basis of the conventional wisdom about the limits to air cooling (typically in the 6-8 kW per cabinet range) reside in the relationship between airflow, temperature rise and heat dissipation. This relationship can be described by the equation $CFM = 3.1W/\Delta T$ where:

CFM = cumulative air consumption of the equipment in the cabinet

W = heat load of the cabinet in watts

ΔT = difference between the equipment in-take temperature and exhaust temperature

Because there is some fixed limit on how much air can be pushed through a perforated floor tile or tile grate, the CFM part of this relationship defines the ceiling on air cooling. For example, if we were able to get 800 CFM through a tile grate and we had equipment in the cabinet that operated with a 20°F temperature rise, we could figure we could cool about 5.2 kW in that cabinet.

$$\begin{aligned}800 \text{ CFM} &= 3.1W/20^\circ\text{F } \Delta T \\ W &= 5161.3\end{aligned}$$

However, by removing the return air from the room, the dependency on chilled air being delivered from a near enough source to diminish the effects of mixing with re-circulated return air is broken. The delivery points for the chilled air no longer matters – chilled air can enter the room anywhere and as long as the room remains pressurized, there will be, by definition, adequate air volume in the room to cool any heat load any place in the room.

This isolation between supply air and return air also means that supply air temperatures can be run significantly higher. Typically supply air temperatures are set around 55°F in order to meet the ASHRAE TC9.9 recommendation for air delivered to equipment to fall in the 68 - 77°F range. The lower supply air temperature is required to compensate for the natural mixing that takes place in the data center between supply air and re-circulated return air, and typically results in significant temperature stratification from the equipment lower in the cabinet to equipment higher in the cabinet. Without any return air in the room, there is no longer any such stratification and therefore no longer a need to deliver supply air below the required temperature for the equipment. In a room with such isolation, it is safe to raise that supply air up to the 68 - 77°F range.

Higher supply air temperatures, in conjunction with cooling unit economizers, offer a significant economic benefit. Historically, the ROI for water side economizers in data center applications has been a hard sell in most geographic areas. Chiller plants typically need to keep the water at 42°F to deliver 55°F supply air, and assuming a 5°F approach temperature, the economizer would only take over for the chiller plant when the temperature dropped below 37°F. However, with the supply air temperature allowed to creep into the 70's °F, then the condenser water can run around 55°F and with a 5°F approach temperature, the economizer can kick in whenever the temperature

drops below 50°F. At that temperature, data center managers in most areas of the country would have access to significantly more “free” days of cooling. In fact, Intel reported a \$144,000 energy savings for a medium sized data center in the Pacific Northwest that depended on effective return air isolation, which allowed for a higher supply air temperature (see Garday, 2007). In addition, McKinstry Co. has analyzed the economic benefits of deploying chilled water air handlers with evaporative air economizers and absolute separation of supply air and return air in the data center. Five megawatt data center tenants, in eight different geographic areas found significant cooling cost savings versus a benchmark facility with standard open hot aisles and cold aisles, using chilled water CRAC units with water economizers.⁴

HVAC Operating Costs

Location	Hot Aisle/Cold Aisle	Supply-Return Isolation
Atlanta	\$318,722	\$122,712
Boston	\$413,673	\$137,217
Dallas	\$1,042,479	\$289,010
Denver	\$449,557	\$120,636
Minneapolis	\$655,527	\$216,306
San Jose, CA	\$742,564	\$185,496
Seattle	\$514,168	\$138,806
Wenatchee, WA	\$225,010	\$71,137

Table 1: Comparison of HVAC electrical, HVAC water and maintenance costs between a standard hot aisle/cold aisle data center with chilled water CRACs and water side economizer, versus separated hot and cold aisles with chilled water air handlers with evaporative air economizer. Five megawatt data center. Data from Sloan, 2006.

The McKinstry study also looked at various options for improving the total efficiency of the standard data center open hot aisle/cold aisle layout compared to the fully isolated arrangement at various levels of loading, and the segregated supply air and return air models delivered lower operating costs than all other combinations.

Chart 1, HVAC Options Annual Cost Comparison, displays the averages of those results for the eight geographic areas of the McKinstry Co. study. While the variations between regions can be extreme, the chart of averages nevertheless provides an accurate view of the relative operating economies of the different systems. The “Base” data center employs chilled water CRAC units with water side economizer, with the room arranged according to hot aisle/cold aisle best practices. Option A employs chilled water air handlers with air side economizer and stream humidifier, again with hot aisle/cold aisle best practices. Option B employs chilled water air handlers with evaporative air economizer and hot aisle/cold aisle best practices. Option C employs the same cooling plant as Option B, but assumes absolute separation between the supply air and return air.

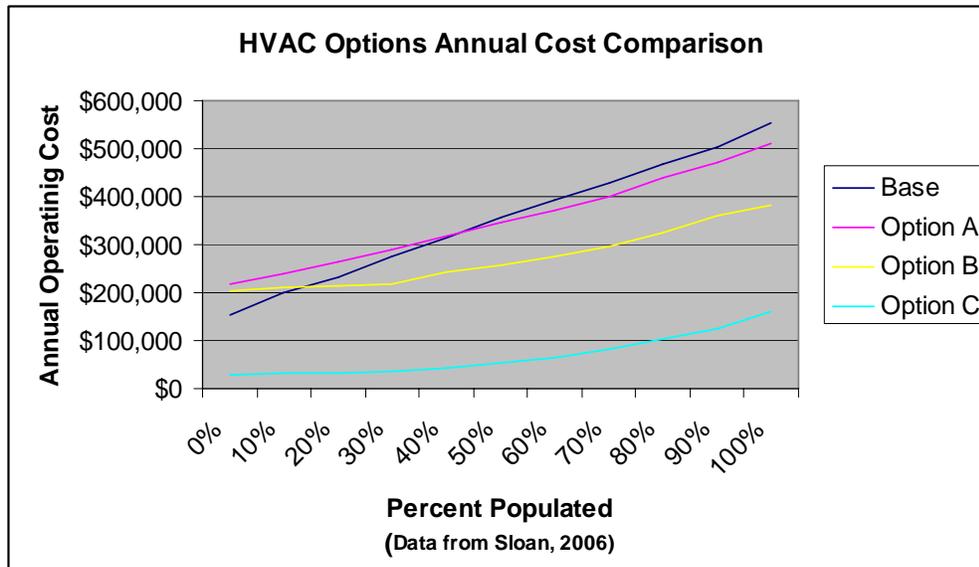


Chart 1

While the evidence clearly favors the model of fully isolating supply and return air over a standard hot aisle/cold aisle layout with open space return air paths for both lower operating costs and higher densities, there are also some implications regarding the claims of proponents for closely coupled liquid cooled solutions. For example, it is claimed that closely coupled liquid cooled solutions are more economical to implement incrementally, thus saving the cost of operating a full blown data center cooling plant when a room is not fully built out. The McKinstry Co. data suggest otherwise, however, with the extremely low operating costs for Option C when the space is less than 50% populated. In addition, lower operating costs in general have been central to the value proposition offered for closely coupled liquid cooled solutions, particularly as a rebuttal to objections over their significantly higher acquisition costs. In fact, the keynote address at this spring's AFCOM Data Center World explained the PUE (power usage effectiveness) metric for measuring data center power and cooling efficiency (total data center power consumption divided by IT power consumption) and suggested a PUE of 1.6 was a measure of efficiency toward which the industry could aspire with closely coupled liquid cooling solutions. Compare this to today's typical data center with a PUE around 3.0 and perhaps a 2.4, which can be achieved through implementing best practices in hot aisle/cold aisle layout and maybe a computational fluid dynamics (CFD)-enabled 2.0 PUE.⁵

The previously cited Intel white paper on wet side economizers uses the HVAC effectiveness ratio as a metric to demonstrate the efficiency of their solution model (Garday, 2007), but that 6.19 ratio on an IT load of 5830 kW, equates to a PUE of approximately 1.3. In addition, the eight data center models in the McKinstry study (Sloan, 2006) showed cooling operating costs (HVAC electricity + water + maintenance) ranging from only 5.1% up to 9.1% of the total data center operating costs (IT + UPS + transformers). Proponents of close coupled cooling solutions also claim to make more floor space available for IT; however, the McKinstry study shows an average of less than 2% of the data center floor space taken up with indoor HVAC equipment. Therefore, with recognized higher acquisition costs and apparently without the touted operating cost advantages, close coupled solutions do not appear to offer the user benefits that overcome concerns about redundancy and fail-over complexities.

While the message is clear regarding the operational advantages of air cooling where maximum isolation is maintained between supply air and return air, particularly in conjunction with centralized chilled water air handlers and evaporative economizers, the data center manager can still reap benefits in a space equipped with conventional chilled water CRAC units and water side economizers. The primary benefit derives from increased cooling unit efficiency because of higher return air temperatures resulting from eliminating the cooling effect of mixing with bypass air in the data center. These efficiency gains can be quite remarkable, as shown in Table 2. Employing the strategy of isolating supply air from return air can provide an economical path for making incremental density increases.

Cooling Unit	Supply Air Temp.	Return Air Temp.	Cooling Capacity
Standard 10 ton CRAC	60°F	70°F	7.8 tons
	60°F	90°F	15.5 tons
	60°F	105°F	20.7 tons
Standard 30 ton CRAC	60°F	70°F	23.0 tons
	60°F	90°F	46.0 tons
	60°F	105°F	61.3 tons

Table 2: Effect of Return Air Temperatures on CRAC Performance Rating (Source: ANSYS Corporation, CoolSim cooling unit vendor library)

I don't want to leave you underestimating the complexity involved with implementing a well engineered air cooled data center with effective isolation between cool supply air and heated return air. However, building this style of data center is not really any more complex than any data center design and construction project, whether that is a conventional hot aisle/cold aisle layout or a close coupled liquid cooled variation. It just requires making a bit of a paradigm shift up front and you will find that most HVAC engineers will understand these values quite clearly. The rewards are higher operating efficiencies, lower energy costs and the peace of mind that comes with knowing you have created an environment that minimizes the kinds of complexity that jeopardizes uptime availability.

¹ Stansberry, Matt and Schlack, Mark. July 13, 2007. "Data Center Infrastructure: Too Soon for Liquid Cooling," Search DataCenter.com.

² Tschudi, William; Mills, Evan; Greenberg, Steve; and Rumsey, Peter. "Measuring and Managing Data Center Energy Use," HPAC Engineering

³ Garday, Doug. May 2007. "Reducing Data Center Energy Consumption with Wet Side Economizers," Intel white paper

⁴ Sloan, Jeff. October 26, 2006. Unpublished study. For information, "Contact Us" tab on <http://www.mckinstry.com/>

⁵ Belady, Christian. March 26, 2007. "Getting the Most Out of Your Data Center: Why Does It Matter?" Keynote presentation at Spring 2007 AFCOM Data Center World.