White Paper

Data Center Containment Cooling Strategies

WHITE PAPER – EC9001

Geist
Updated August 2010

Abstract

Deployment of high density IT equipment into data center infrastructure is now a common occurrence yet many data centers are not adequately equipped to handle the additional cooling requirements demanded by this high density IT equipment. This results in non-desirable conditions such as recirculation or mixing of hot and cool air, poorly controlled temperature and humidity, and costly cool air over-provisioning. Many systems claim to provide efficient and effective answers to these problems by physically separating cold supply and hot return air and when physical separation is managed appropriately separation will result in a stable thermal environment for the IT equipment while improving the performance

of the cooling infrastructure. This paper will present multiple containment cooling strategies, outline limitations of these strategies and characterize them for performance and their ability to reduce operational costs.

IT Equipment Challenges

The average total rack load density for many data centers has increased significantly due to the many benefits of deploying 1U and 2U rack mount and blade servers. A typical 42U rack fully loaded with 1U servers would yield a total rack load of approximately 10 kW. Many facilities have to limit the number of servers deployed in a rack to keep total load below 6-8 kW because of power and cooling limitations. Some organizations have created high density areas within their existing data centers with added power and cooling to handle the higher per rack load density.

Obtaining an accurate industry estimate for average rack power loads would normally be difficult; however, based on a recent Data Center Decisions poll, shown in Figure 1, rack loads have increased to where 30% are now indicating that the average rack load in their facility is 6-10 kW. Any cooling and cool air distribution system for a high-density data center retrofit or for facilities currently on the drawing board, should aim to provide an environment that is within ASHRAE Class 1 limits[1] for temperature and humidity.

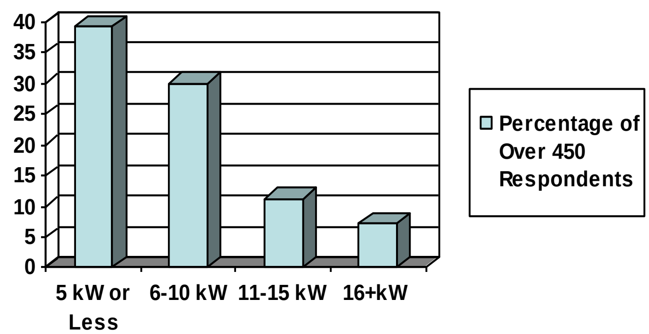

Figure 1: Average rack loads based on a Data Center Decisions poll from October 2007.

ASHRAE Class 1 Standards for Mission Critical IT Equipment

Data center facilities with mission critical operations and tightly controlled environmental parameters (dew point, temperature, and relative humidity) typically contain computer servers, networking and storage devices. The recommended conditions for ASHRAE Class 1 limits[1] are;

68-77 °F supply air temp
40-55% RH

Figure 2 shows locations for measuring the environmental conditions of ASHRAE Class 1.

Manufacturers of IT equipment are supportive of the ASHRAE design standards and support for this standard is not limited to IT equipment that is currently shipping from the equipment manufacturer. Data center facility designers and engineers need to know that the IT equipment providers will stand behind these standards in the future for data centers that are being planned today.

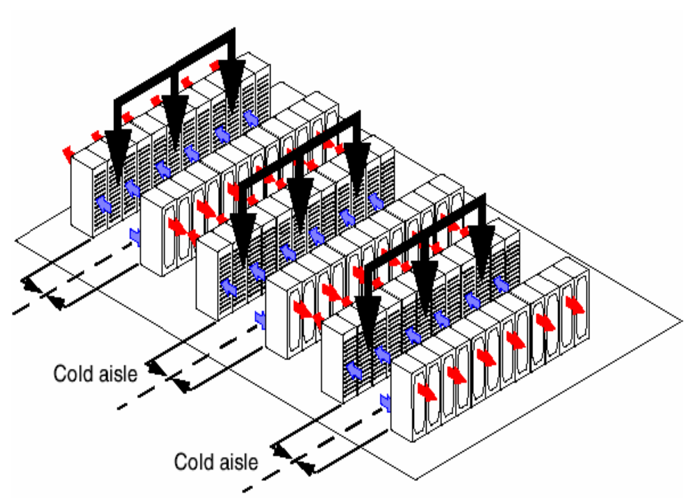

Figure 2: ASHRAE specifies 3 measurement locations each 59” from floor in cold aisle.

Conventional Hot/Cold Aisle Cooling Limitations

Conventional hot-aisle/cold-aisle cooling methods for higher density equipment loads may provide an adequate volume of air to the IT space, but does the cool air get to the equipment before being inappropriately contaminated by the IT equipment waste heat?

A common approach to prevent waste heat from contaminating the cool supply air is to over-cool and over-supply the cold-aisle. Trying to eliminate hot spots at the front of the IT equipment rack by supplying significantly more air than is required by the IT equipment is inefficient and will produce unpredictable and non-uniform results. In addition, any satisfactory balance one could achieve with the distribution of over-provisioned cool air will be affected when new applications are deployed onto the data center floor. Adequate separation between the cool supply and hot return airstreams is necessary for efficient and predictable results.

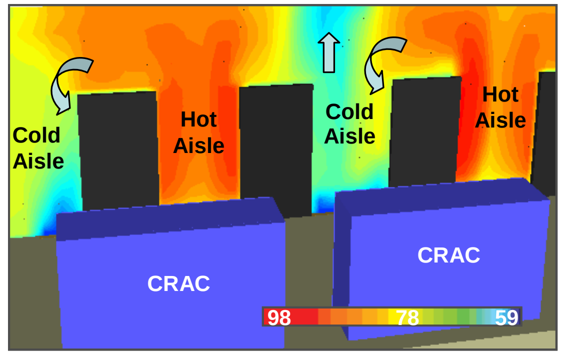

Figure 3: Rack inlet temperature is 76 Deg F with conventional cooling and 50% excess cool air supply.

Figures 3 and 4 demonstrate the effects of hot air recirculation for two different operating scenarios using Computational Fluid Dynamic (CFD) models. Figure 3 demonstrates the hot air recirculation at the face of the IT equipment rack when 50% excess cool air is supplied than is required by the IT equipment. Figure 4 demonstrates the hot air recirculation at the face of the IT equipment rack when 75% excess cool air is supplied than is required by the IT equipment. In both examples the supply air is delivered at 59° F (represented by the darker blue color) and the exhaust air exiting the IT equipment load is 98° F (represented by the darker red color).

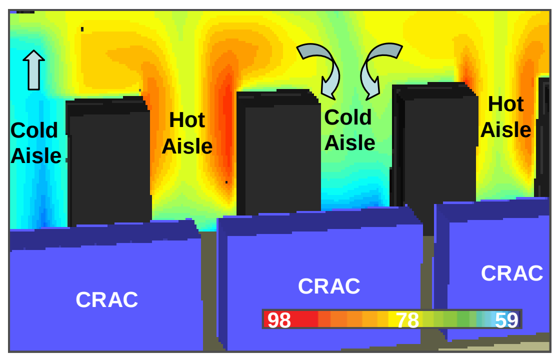

Figure 4: Rack inlet temperature is 73 Deg F with conventional cooling and 75% excess cool air supply.

A recent study demonstrated that 2.6 times more cool air than is required by the IT equipment was being supplied to the data center floor[2]. This study looked at 19 facilities totaling 204,000 sq-ft of raised floor space.

The excessive delivery of cool air also results in a significant amount of cool air mixing with heated exhaust air before returning to the cooling unit. The return air temperature to a cooling unit coil in this scenario is typically only 10-15 °F warmer than the cool supply air. Data from chilled water cooling unit manufacturers indicate the optimal cooling tonnage and effectiveness is at higher return air temperatures of 90-105 °F. This will be covered in greater detail in the characterizing performance and efficiency section.

As shown in Figure 5, for 1000 kW IT equipment load, we demonstrate the cost of distributing cool air, excluding improvements that can be made with chiller and cooling tower performance. To deliver the required cool air for this load, approximately 68,000 CFM (cubic feet per minute) is required. At electricity rates of $0.055 per kW/hr, the cost for delivering the required cool air is $18,500 annually. The cost for supplying 2.6 times more cooling than is required for the IT equipment load is an additional $28,200 annually[3].

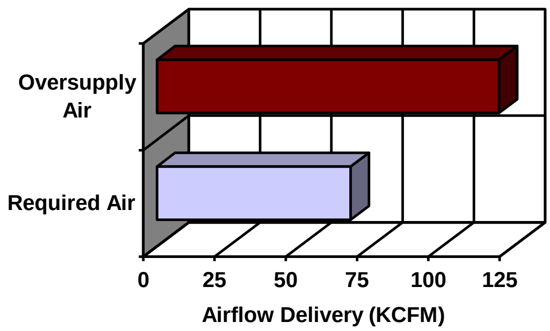

Figure 5: Based on a 1000 kW load at $0.055 per kW/ hr, the actual cost to deliver only the required cool air is $18,500 annually. 2.6 times oversupply costs at additional $28,200 annually.

With energy, personnel and real estate at a premium, the market will continue to demand high-density equipment operating in a predictable and energy efficient environment. Whether the amount of cool air that is short circuiting back to the cooling unit is 30, 50 or 70% for a facility, a significant volume of cool air is being generated, poorly distributed and wasted. When up to 50% of the operating costs of an enterprise data center facility can come from electricity alone, reducing waste is imperative.

High Density Cooling Strategies

Supplying significantly more cool air to the IT equipment than is required may be the most common high density cooling strategy in practice today. This is not a very good strategy for cost savings or for carbon footprint reduction but, the demand for more applications can come at a rapid pace and cramming the data floor with these applications and additional cooling units is a common approach. Another practice to increase cooling density is to line up cooling units, tightly spaced end to end, down each opposing long wall on the data center floor or just outside the data floor in a separate galley. This practice, along with extensions for the return air to a higher point in the ceiling has been very common for newly designed and built data centers.

Figure 6, which illustrates the air distribution challenges for conventional cooling, will be used as a baseline for comparison as we explore air distribution for available high density cooling strategies.

A high density cooling strategy must include a method to provide good separation of the cool supply air from the heated exhaust air. Systems are available which provide either a degree of separation between the cool supply and hot return airstreams or provide a physical barrier for separation of the cool and hot return airstreams. For example; water-cooled enclosures, not pictured here, contain a heat exchanger in the base of the cabinet and provide a physical barrier, which is the vertical stack of the IT equipment, separating the two airstreams.

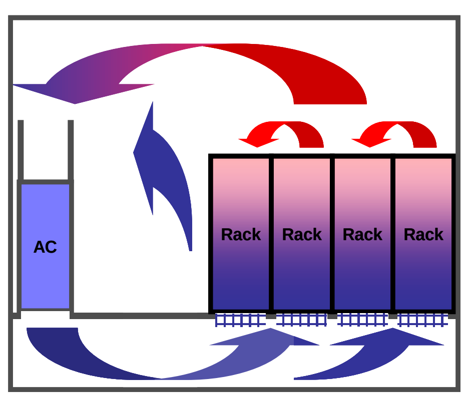

Figure 6: Oversupply of cool air for a high density application can be marginally effective but will be costly to operate.

Although these systems are referred to as "water cooled", the benefit comes from relocating the cooling element closer to the load for improved air distribution to and from the IT equipment.

Above-cabinet cooling units, see Figure 7, and cooling cabinets placed in the row, see Figure 8, are positioned to collect and cool the heat from the rear of the IT equipment cabinet. The cool air is then discharged to the cold aisle and the intakes of the IT equipment.

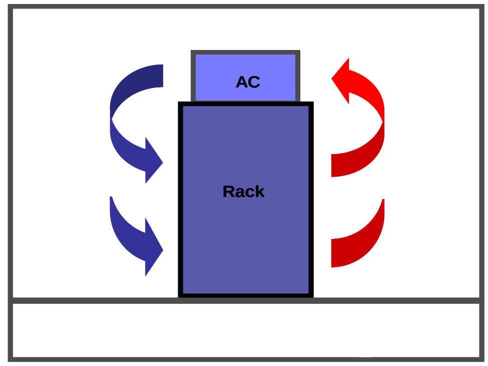

Figure 7: Airflow distribution patterns with above-cabinet cooling units as supplemental cooling to conventional cooling.

These strategies locate the cooling element closer to the IT equipment exhaust heat and do a much better job at collecting the waste heat than conventional cooling methods. Above-cabinet cooling units still require conventional CRAC/H systems for a significant portion of the required cool supply air.

Oversupply of cool air and therefore bypass of cool air to the return of the cooling units still exists with the methods demonstrated in Figures 7 and 8. The level of oversupply required for effective performance should be modeled using CFD under various load conditions before choosing this option. If the desired approach is to completely reduce oversupply and the resulting recirculation, a physical barrier separating the hot and cold airstreams will be necessary.

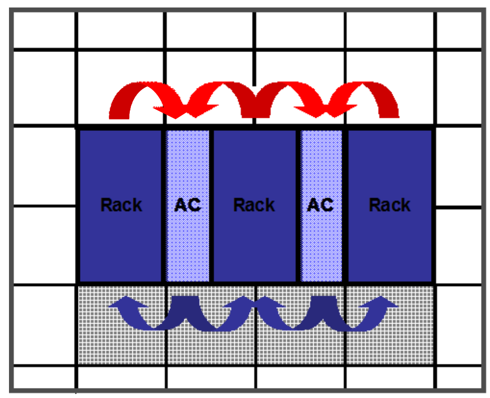

Figure 8: Airflow distribution patterns with cooling units placed in the row - view looking at top of cabinet row.

The heat containment method shown in Figure 9 utilizes a physical barrier to contain the IT equipment heat and provide a predictable pathway for its return to the data center cooling units. The cool air is delivered to the IT equipment through the raised floor perforated tiles, as in a conventional system. Heated exhaust air from the IT equipment can be removed from the rear of the rack using high flow low power fans. The air is then ducted to the raised ceiling plenum for return to the cooling units. In this arrangement, all heated exhaust air is contained. The maximum amount of cool air will be available to the IT equipment load when all waste heat is contained and returned to the cooling units.

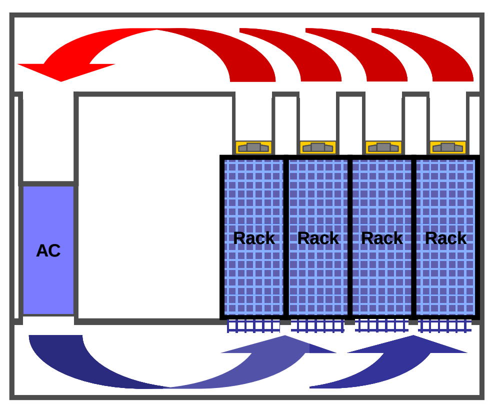

Figure 9: Airflow distribution patterns with HDHC method of cooling - view looking at front of cabinet row.

Heat containment has another advantage; the IT equipment is able to draw cool air from the room's total volume of air. A high-density rack can be positioned next to a low-density rack without negative consequences. Uneven under floor pressure distribution and the resulting disproportionate airflow rates from the floor tiles are not of concern when the hot air is contained.

Characterizing Performance and Efficiency

Total cooling output should be matched to the total cooling requirements for the data center to be most effective at optimizing energy costs. Sacrificing some energy savings for operational flexibility may be an important trade-off. Anyone who has taken the time to evaluate and characterize high density cooling strategies can appreciate the difficulty of determining performance and efficiency benefits. When comparing design or engineering strategies, an often used and very useful tool is a decision matrix. A decision matrix can be invaluable when evaluating and comparing high density cooling strategies for the following reasons;

- Each cooling strategy will have its strengths and weaknesses. One cooling strategy will not meet all the requirements of the facility, IT and networking community.
- There are many skills and different departments involved in the operation of a data center. Having a common tool will help to facilitate input, evaluation and be a method for documenting assumptions.
- The process of defining and rating the various levels of critical needs will allow teams to deliberate based on fact and not predetermined bias.

An example of a data center cooling decision matrix is shown as Figure 10[4]. This is one example and the actual criteria, value for each criteria and score you apply, is dependent on the business strategy for the data center and the team responsible to execute on the strategy. Additional criteria, such as data center appearance or data center facility conveys reliability, could be added if necessary for the business strategy. For the decision matrix shown in Figure 10, each section has been normalized to account for the number of criteria in that section. For example; In the Initial Cost section the score for each criterion was multiplied by 1/5 or 0.2 based on there being 5 criteria in that section versus 9 criteria in the User Reliability & Risk section. In this latter section, each criterion was multiplied by $1/9^{th}$. Additionally, criteria such as Chiller Performance High or CRAC/H Performance High would require separate evaluations to understand the effects of a cooling strategy on these systems.

Data Center II Cooling Strategy Decision Matrix									
								Rating: 5 Highest Value	
		Conventional Cooling - High Return		High Density Heat Containment		Conventional Cooling Plus Supplimental In-row Cooling		Conventional Cooling Plus Supplimental Over-cabinet Cooling	
Initial Cost - CAPEX	**Category Value**	**Rating**	**Value**	**Rating**	**Value**	**Rating**	**Value**	**Rating**	**Value**
			9.6		9.6		9.6		9.6
Procurement Cost Low	4	3	2.4	3	2.4	3	2.4	3	2.4
Hardware Installation Cost Low	4	3	2.4	3	2.4	3	2.4	3	2.4
Initial Training Time and Cost is Low	3	3	1.8	3	1.8	3	1.8	3	1.8
Electrical, Network and Sensor Installation Low	3	3	1.8	3	1.8	3	1.8	3	1.8
Initial Start-up Time is Quick and Limits 'Experts'	2	3	1.2	3	1.2	3	1.2	3	1.2
Continuing Cost - OPEX	**Category Value**	**Rating**	**Value**	**Rating**	**Value**	**Rating**	**Value**	**Rating**	**Value**
			11.6		11.6		11.6		11.6
Service Life High	5	3	2.1	3	2.1	3	2.1	3	2.1
Free Cooling Hours High*	5	3	2.1	3	2.1	3	2.1	3	2.1
Maintenance Costs Low	4	3	1.7	3	1.7	3	1.7	3	1.7
Cooling Tower Performance High*	4	3	1.7	3	1.7	3	1.7	3	1.7
Chiller Performance High*	3	3	1.3	3	1.3	3	1.3	3	1.3
CRAC/H Performance High*	3	3	1.3	3	1.3	3	1.3	3	1.3
Humidification Performance High*	3	3	1.3	3	1.3	3	1.3	3	1.3
User Reliability & Risk	**Category Value**	**Rating**	**Value**	**Rating**	**Value**	**Rating**	**Value**	**Rating**	**Value**
			10.5		10.5		10.5		10.5
Fewest Components & Interconnects	5	3	1.5	3	1.5	3	1.5	3	1.5
Reduces Human Interaction / Easily Maintained	5	3	1.5	3	1.5	3	1.5	3	1.5
Single Failure/Repair Not Effecting DC Production	5	3	1.5	3	1.5	3	1.5	3	1.5
Reliability Data is Available for Cooling Strategy	4	3	1.2	3	1.2	3	1.2	3	1.2
Provides Early Alarm Conditions	4	3	1.2	3	1.2	3	1.2	3	1.2
Proven Technology and Control Systems	4	3	1.2	3	1.2	3	1.2	3	1.2
Same System Used in Existing and New Facilities	3	3	0.9	3	0.9	3	0.9	3	0.9
Installation Does Not Cause Production Interuption	3	3	0.9	3	0.9	3	0.9	3	0.9
Systems Report Available Capacity	2	3	0.6	3	0.6	3	0.6	3	0.6

Figure 10: Example of a high density cooling strategy decision matrix.

When all critical areas have been discussed and the decision matrix process has run its course, it is not out of the question that a data center cooling strategy with a lower score is chosen. In this scenario, it's likely that the assumptions and associated risks could not be well defined and participants choose a path that has a prior track record and is more familiar to the team. If this happens, a conclusion should be documented by the team explaining the reason for choosing an alternative path to what the decision matrix tool had indicated.

Fan Power Considerations

Fan power efficiency for a cooling system is one of the considerations for choosing a cooling strategy. Fan power and airflow do not have a linear relationship. The cubic fan power law has a significant effect on power consumption as shown in Figure 11. With CRAC/H fans delivering 75% of the rated airflow capacity, the power consumption by the fans are approximately 50% of the full rated power.

One redundant CRAC/H for each specified area within the data center should provide the added capacity required when another CRAC/H unit fails or undergoes repair.

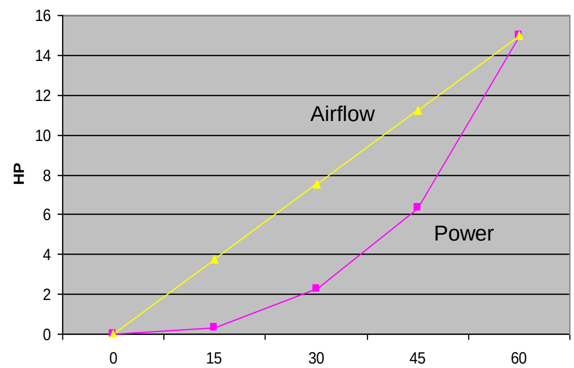

Figure 11: Cubic fan laws demonstrate significant power savings at slightly reduce airflow rates.

Referring to Figure 11 again, there is significant power savings if the fans are running at a slower speed when compared to fewer CRAC/H units all running at full speeds for the same volume supply requirement. Other benefits for controlling fan speeds in CRAC/H units are;

- Redundant CRAC is already running
- Increase Mean Time between Failures running at slower speeds
- Soft start

The CRAC/H manufacturer's standard is to use chilled water valve position to control the VFD speed. This produces a variable under floor pressure, which would vary the delivery of airflow to the data center. A CRAC/H internal VFD control signal can be divorced and the BMS system can be used to control CRAC/H fan speed. CRAC/H fan speeds can be controlled based on under floor pressure or other data representing the total volume of supply air required.

Supply Air Temperature to Maintain ASHRAE Class 1 Standards

Utilizing a conventional cooling strategy and return air temperature control will not maintain a steady supply air temperature. Supply air leaving the floor grate at 59-65 ºF, which is well below the ASHRAE accepted range, will typically result in return air temperatures of 70-75 ºF. This is due to significant mixing of airstreams. The CRAC/H unit is set up to measure return air temperature and control the cooling valve position to maintain a steady return temperature. The return air control strategy is driving the supply air temperature and relative humidity outside ASHRAE standards. This strategy will require more energy to provide proper relative humidity for the IT equipment.

By using supply air temperature control, one could set the delivery of air at 70-72 degrees and 45% RH. Moving temperature control and RH sensors under the raised floor into the CRAC supply air stream will provide the critical information needed to maintain ASHRAE temperature and humidity standards. However, the CRAC/H will perform very little de-humidification in this mode. This strategy would also require;

- Containment of exhaust heat, thereby eliminating mixing of heated air with the cool supply air.
- Humidity control regulated by other means such as a central air handler, which can also be responsible for the facilities make-up air.

CRAC/H performance considerations

A CRAC/H in a heat containment cooling strategy will operate at greater efficiency than a CRAC/H installed in a convention cooling system. The return air will be much drier and the temperature will be elevated. A CRAC/H manufacturer has supplied the following data for operating in these conditions.

Return Dry Bulb	Total KBTU/h	Sensible KBTU/h	Enter Fluid Temp	Leave Fluid Temp	GPM	Supply Dry Bulb Out	Supply RH
72	438	364	45.0	58.5	70	51.1	91.0
80	561	493	45.0	62.0	70	51.4	89.1
90	716	643	45.0	66.5	70	52.1	88.0
100	871	779	45.0	71.0	70	53.2	87.1

Table 1: 45-degree entering chilled water temperature with control valve full open.

Referring to Table 1, the top line is fairly close to a conventionally cooled data center with return temperature controls. The CRAC is capable of cooling from a dry bulb 72 to 51.1 ºF. Here the supply air conditions are well outside of the ASHRAE Class 1 standard. Also, notice that the sensible cooling is 364 KBTU/h, quite a bit lower than the total cooling of 438 KBTU/h. The 364 KBTU/h provides only 30 tons of sensible cooling for a 40 ton CRAC/H unit[5].

Consider the other rows of data with the elevated return dry bulb air temperatures. In a heat containment scenario, we can expect any of these return temperatures to occur. The CRAC/H is capable of increased tonnage as the return air temp is elevated. In fact, the CRAC/H almost doubles its capacity if the return dry bulb air is 100 degrees. The return air temperature goes up dramatically, but the supply air temperature is fairly consistent, it only goes up a few degrees. Greater temperature differential from chilled water and return air, improves coil performance. This results in better efficiency from the CRAC/H unit.

Meeting ASHRAE Class 1 standards will require the CRAC/H controls to maintain temperature by throttling the chilled water control valve. Table 2 contains CRAC/H manufacture supplied data to maintain 68 degrees supply dry bulb and 45% RH.

Return Dry Bulb	Total KBTU/h	Sensible KBTU/h	Enter Fluid Temp	Leave Fluid Temp	GPM	Supply Dry Bulb Out	Supply RH
72	46.6	46.6	45.0	65.6	7	68	45
80	181	181	45.0	69.2	17	68	45
90	356	356	45.0	78.2	23	68	45
100	518	518	45.0	83.8	28	68	45

Table 2: 45-degree entering chilled water temperature with control valve throttled.

The coil performance data indicates that the CRAC/H requires a lower cooling water flow rate. The low flow rate creates laminar effect in the coil, which should be avoided. In fact, the cooling performance of the CRAC/H has increased to the point where it will be most efficient to dial back some cooling capacity and let the chillers run at their most efficient operating parameters.

One method to decrease cooling would be to let the CRAC/H regulate flow as required; however, this will erode the control valve if the valve is operated in this condition for extended period of time. For both reliability and efficiency, reducing CRAC/H cooling capacity to meet the on-going and changing needs of the facility, should be done by raising the chilled water temperature[5].

Chilled Water Plant Performance Considerations

Figure 12 represents manufactures data demonstrating power usage for two types of chillers at 45 and 50 degrees entering chilled water temperature.

Both refrigerant type chillers run more efficiently and give additional capacity if the chilled water temperature is raised. The R134-A high pressure chiller creates a 9% capacity increase and 6% energy savings and the R123 low pressure VFD chiller creates 17% capacity increase and 12% energy savings just by increasing chilled water temperature 5 degrees. CRAC/H units as well as Air Handler Units (AHU) receive the same temperature chilled water. Chilled water temperature reset must be driven by the air handler discharge air temperature to maintain control of the de-humidification process. Splitting the chilled water loop or adding supplemental cooling to the AHU so the data floor chilled water temperature can be maximized would be a prudent approach[5].

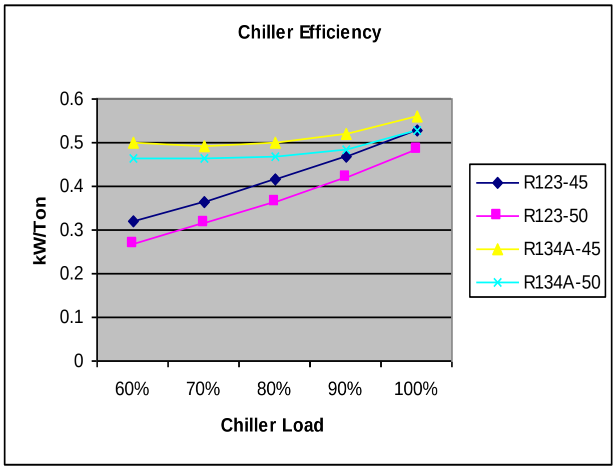

Figure 12: Chiller Efficiency as a function of refrigerant, supply water temperature and chilled water temperature.

Maximizing chilled water temperature will provide increased hours for available waterside economizer operation, to the point where it becomes economically feasible even in warmer climates. Using the above as reference, and considering a more effective use of a water-side economizer to gain more free hours of cooling, raising the supply air temperature to 70 ºF would require approximately 55 ºF chiller condenser water. In comparison, a 59 ºF supply air temperature would require approximately a 45 ºF condenser water temperature. With a 5 ºF approach temperature, water-side economizers could be utilized at up to 50 ºF outdoor air temperature for a 70 ºF supply air temperature versus an up to 40 ºF outdoor air temperature if supply air is left at 59 ºF.

CRAC/H Performance at Elevated Chilled Water Temperatures

CRAC/H units operating at higher chilled water temperatures should be looked at carefully. The CRAC/H must continue to have sufficient cooling capacity for all conditions.

Figure 13 indicates that the CRAC/H operating on elevated chilled water can produce ASHRAE standard 70 ºF supply air as long as the return air temperature is at or below 105 ºF. If the return air exceeds 105 ºF the chilled water will be required to be reset back to design conditions or let the supply air exceed the 70 ºF set point.

All the CRAC/H units must continue to produce at least the rated cooling capacity to maintain cooling redundancy. In this case, each CRAC/H must produce at least 40 tons total cooling. Referring to Figure 14, with elevated chilled water temperatures, the CRAC/H can produce at least 40 tons total cooling for all conditions except one. It will not make rated capacity if the return air drops below 79 ºF[5]. This unlikely condition will not cause a cooling disruption if a CRAC would fail due to the CRAC not being heavily loaded.

Efficient operation of the CRAC/H units at elevated chilled water temperatures can be achieved for most conditions.

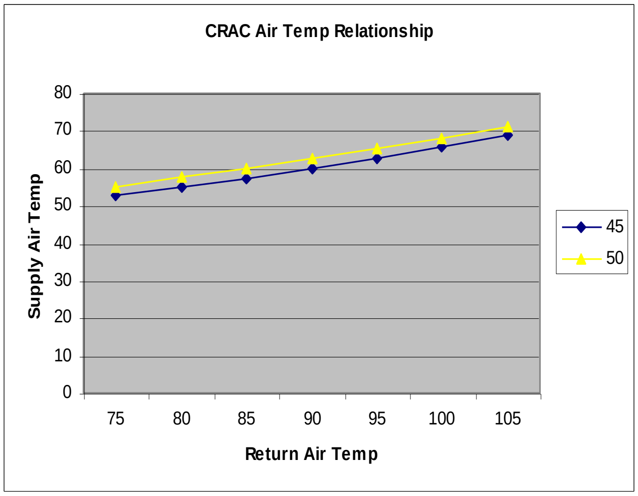

Figure 13: Minimum CRAC supply temperature as a function of return air temperature and chilled water supply temperature.

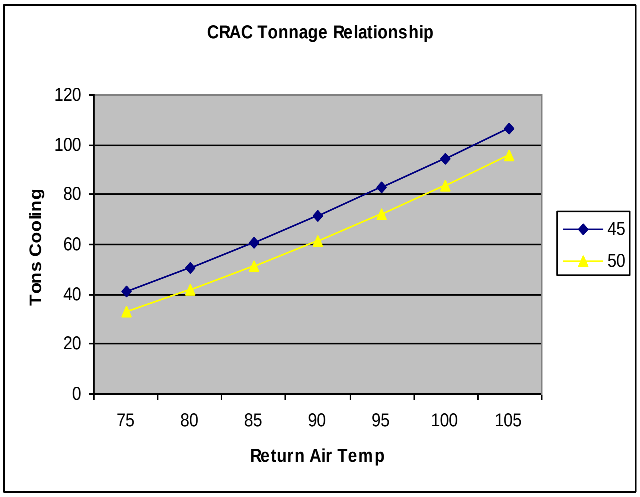

Figure 14: CRAC/H cooling capability as a function of return air temperature.

Moisture Control Strategy

With CRAC/H units only cooling to 70 ºF, the air will not be de-humidified and therefore another method of de-humidification will be required. The use of outside air handlers will deliver pre-conditioned air into the building to maintain a positive pressure of air inside the IT equipment environment.

Air handlers can also be responsible for de-humidifying and humidifying the environment. The air handler supply air temperature can be reset driven by an average calculation block of all the relative humidity of the CRAC/H supply air. Only in an extreme high relative humidity condition would the CRAC/H need to operate in a de-humidification mode. CRAC/H dehumidification must be implemented in stages, utilizing areas with coldest return air first to minimize the CRAC/H unit load to de-humidify. This de-humidification control method ensures controls do not allow simultaneous dehumidification and humidification, significantly effecting energy saving performance.

Rack Airflow and Thermal Management

For maximum efficiency, each rack in a containment system should return close to the same airflow volume as the IT equipment within the rack. The air distribution patterns, as shown in Figure 15, include cool air supply, return duct exhaust and CRAC/H return air.

When return duct exhaust fans are oversized and not controlled, more air volume is being returned and a slight negative pressure will exist in the rack. Although this is an improved condition for the operation of the IT equipment fans, cool supply air will be pulled into the rack through the many small openings, thereby bypassing the IT equipment and returning to the CRAC/H unit unused.

When return duct exhaust fans are undersized for the IT equipment airflow, the rack will be pressurized and waste air will leak out many of the small openings in the enclosure. Upon close observation one can see that these openings can be substantial. Openings are also located at the interface between the IT equipment and at the front mounting rails of the equipment rack.

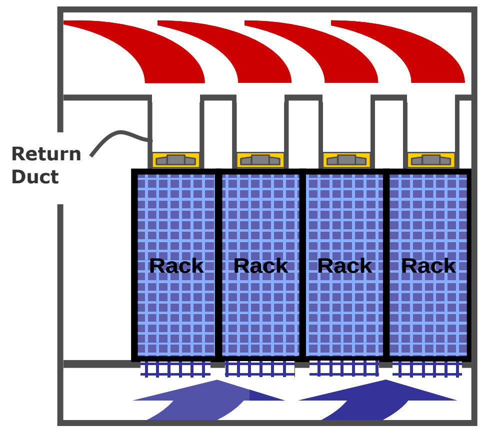

Figure 15: Return duct exhaust within a heat containment distribution system.

Table 3 shows measured data from one heat containment installation[5]. In this particular example, the return duct exhaust fan is oversized and not controlled to match the IT equipment airflow rate. As can be seen here, the differential airflow is 240 CFM of cool supply air that is bypassing the IT equipment and returning to the CRAC/H unit, thereby wasting the cool air and reducing the temperature differential across the CRAC/H coil. Also shown in Table 3 is measured data when the return duct fan is turned off. In this configuration, each rack would allow 360 CFM of waste heat to leak out into the environment.

IT Equipment Airflow, CFM	Return Duct Fan Status	Exhaust Airflow, CFM	Airflow Rate Differential, CFM	Rack Pressure, inches WC
1400	Full-on	1640	240 (bypass)	-0.05
1400	Off	1040	360 (leakage)	0.16

Table 3: Airflow rate versus rack pressures are compared for two modes of a return duct system.

Rack plenum pressures and the resulting bypass or leakage air will vary based on total airflow rates and system impedance (resistance to flow). The challenge with any heat containment system is to eliminate bypass of the cool supply air and to eliminate contamination of the cool supply air with waste heat. To achieve optimal efficiency, the return duct exhaust fans should operate at the same flow rate as the IT equipment. Equipment within a rack operating with a zero or a slightly negative back pressure will have additional efficiency improvements for the entire data center and actively balancing return duct airflow to IT equipment airflow based on rack plenum pressure should be a goal for future heat containment systems.

Summary

Many high density data centers are not operating effectively due to significant mixing of cool supply and hot return airstreams. Overprovision of the cool supply air will not allow ASHRAE Class 1 conditions to be maintained because supply temperatures will be well below the lower limit of the ASHRAE Class 1 range in an attempt to eliminate hot spots. CFD models as well as published studies have shown that twice as much air, at colder than required temperature, is being delivered to maintain the upper limit of the ASHRAE standard. There is a significant cost associated with oversupply as well as missed opportunity to efficiently operate the CRAC/H and chiller plant to further reduce operational costs.

With thirty percent of the data center community now indicating they have deployed 6-10 kW per rack averages and with energy costs continuing to rise, an efficient and reliable approach to cooling is critical. A strategy to physically separate, not just close-coupling the cooling element to the IT equipment load, is required to eliminate cool air oversupply and the resulting bypass.

Heat containment systems provide the necessary environmental conditions for higher density IT equipment while significantly improving energy efficiency. The total cooling delivered should be matched to total cooling requirements for the data center in order to be most effective at optimizing energy costs, however, sacrificing some energy savings for operational flexibility may be an important decision that has to be made and a data center cooling decision matrix is a useful tool for this purpose. This tool should include; initial capital costs, operational efficiency and maintainability metrics. Reducing oversupply and providing higher supply air temperature allows significant energy saving improvements to CRAC/H and chiller operation and will also make available additional hours of free cooling for a water-side or air-side economizer, justifying the capital expense of the economizer system, even in lower regions within the US.

Equipment familiarity for any system, where operational data and training are already in place, is a very important consideration. Familiar and time-tested components; where knowledge, operational data, training and maintenance programs are already in place, will prove to be the most reliable cooling systems.

With on-going electrical power costs outpacing equipment costs, the drive for the highest efficiency makes good business and financial sense. Additionally, companies that deploy energy efficient data center systems to reduce their carbon footprint will have a lot to talk about.

References:

[1] ASHRAE Datacom Series, Design Considerations for Datacom Equipment Centers, 2005, Thermal Guidelines for Data Centers and Other Data Processing Environment, 2004 and New Guideline for Data Center Cooling, *Don Beaty and Davidson*, ASHRAE Journal. 2003.

[2] Uptime Institute's Data Center Efficiency EPA Conference January 2005

[3] Data Center Dynamics, Energy Efficiency Opportunities in Data Centers, *Mukesh K. Khattar, Ph.D., PE*, San Francisco, July 2006,

[4] Opengate Data Systems, Data Center Cooling Decision Matrix (www.opengatedata.com), *Mark Germagian*, November 2007

[5] Uptime Institute, Orlando and Afcom Data Center World, Las Vegas, 2006, Oracle High Density Data Center, *Mitch Martin and Mark Germagian*

About the Author: Mark Germagian is currently serving as President for Opengate Data Systems, a division of PCE, with responsibility to lead the firm into new technology areas relating to effective and efficient data center operation. Mark has directed technology development and was a design engineer for 22 years with experience designing products for telecom and technical work environments as well as cooling distribution systems for computer data centers. Mark holds multiple U.S. and international patents and is a contributing author

for ASHRAE TC9.9 datacom series publications. (PCE is the parent company of Opengate Data Systems and Geist.)