BICSI news MAGAZINE

may/june 2010

OM4

Meet the Next Generation of Fiber

- BICSI Data Center Standard
- Airflow Containment
- Managing Structured Cabling

BICSI®

design & deployment

Airflow CONTAINMENT:

The Road to Data Center Efficiency

by Chris Jaggers and Travis North

Ducting the exhaust of the equipment is a much simpler way to capture the return air compared with enclosing the entire row. It also is regarded as the most contained of all the solutions.

Chris Jaggers is the director of design engineering for Chatsworth Products, Inc. (CPI) and has more than 14 years of experience leading the development of electromechanical solutions. Chris can be reached at cjaggers@chatsworth.com.

Travis North is an expert in the field of single-phase and two-phase heat transfer and in the design of electromechanical systems. Travis is currently the thermal design manager for Chatsworth Products, Inc. (CPI) and can be reached at tnorth@chatsworth.com.

As compute capability and power density continue to rise with each generation of server products, the requirement for traditional data centers to efficiently support increased rack density becomes critical. To place this in perspective, a study by Intel[1] illustrated that the same compute capability in 2002 that required 25 racks and consumed 128 kilowatts (kW) of power (5.1 kW/rack) could be replaced in 2008 by a single rack of servers requiring only 21 kW of power (21 kW/rack). For the same compute capability, over 127 kW of power could be saved (around $10/operating hr), along with a significant reduction in physical infrastructure.

Factors contributing to this increasing compute density phenomenon are the rapid adoption of virtualization, the development of multi-core CPU architecture, larger system memory capacities and speeds and an increasing trend toward general-purpose computation on graphics processing units (GPGPUs). Virtualization, while not directly increasing the overall density requirement of the rack, is altering the usage profile where the individual rack will have higher utilization, in essence running hotter for longer periods of time. Memory capacities and speeds continue to increase with central processing unit (CPU) core count, and the latest dynamic random access memory (DRAM) technology trends continue to increase memory power utilization. As engineering and science applications continue to take greater advantage of the graphics processor for floating point operations, GPGPU compute is driving the traditional individual server peripheral component interconnect (PCI) cooling from 25 W to 300 W and beyond.

While power requirements for individual components continue to rise, the system performance is outpacing the required power, thus providing significant performance-per-watt savings compared with legacy technology. From these examples, the case can be made that if a data center can effectively cool high-compute rack densities, it is possible to glean power savings and reduce data center infrastructure, in essence improving the overall efficiency of the data center.

If a data center can effectively cool high-performance rack densities, significant energy savings can be obtained, which leads to the question of what density level is required. ASHRAE[2] has projected that the required rack density to support 1U and blade compute servers will require 37 kW of cooling by 2014, while the typical data center is currently at an average density of 6.5 kW. This 5X density gap will require innovative solutions, with the advantages and disadvantages of each described below.

CONTAINMENT

In a poorly designed legacy data center, it has been stated that 60 percent of the supply air is lost due to inefficient airflow management. Legacy open loop hot aisle-cold aisle data centers are challenged by two primary airflow management concerns—recirculation and bypass. Both reduce the overall available airflow for cooling, but do so by two differing mechanisms. Recirculation is the mixing of the hot exhaust with the cold intake of the IT equipment, while bypass short circuits the cooling path by going around the IT equipment. Recirculation can also lead to temperature variations along the height of the cabinet, robbing the data center of cooling capacity. The primary method for minimizing the impact of both bypass and recirculation is to separate the supply and return air streams through containment.

A multiyear case study from Intel[3] showed that by applying smart air stream containment methodologies to minimize both bypass and recirculation, Intel was able to increase rack densities from 310 watts per square foot (W/ft^2) to over 781 W/ft^2 (15 kW to 40 kW per cabinet). Additional studies by Dell[4] and HP[5] proposed similar magnitude increases in rack densities by various hot aisle and cold aisle containment strategies. With containment, the rack density is only limited by the capability of the cooling system to deliver the proper airflow required by the IT equipment. A wealth of technical information is available on containment methodologies, and we will discuss each of the various solutions.

HOT AISLE VS. COLD AISLE

Hot aisle containment (HAC) encloses and captures the hot IT exhaust and ducts the hot air directly back to the computer room air conditioners/handlers (CRAC/Hs). Figure 1 is an illustration of the flow path through a typical fully contained HAC system. While HAC captures the return, cold aisle containment (CAC), as the name implies, contains the cold air supply. Both methodologies have the benefit of isolation of the hot air return from the supply air, allowing for increased CRAC/H efficiencies gained from higher return cooling coil temperatures. Additionally, both will reduce the need for humidification and dehumidification as the air is delivered directly back to the CRAC/H without mixing. The two differ in three areas—scalability, thermal mass and operator comfort.

Scalability of CAC will be limited by the floor space to deliver airflow and the ability to manage airflow across multiple cold aisles. While the argument can be made that this is truly only limited to the number of CRAC/Hs available to deliver and manage the required IT volume of airflow, expressed in cubic feet per minute (CFM), there will be a practical pressure drop limit associated with delivering the airflow through the raised floor tiles, and in the case of in-row coolers, the amount of row space required. HAC on the other hand utilizes the entire data center operating space as a thermal capacitor, which typically is an order of magnitude

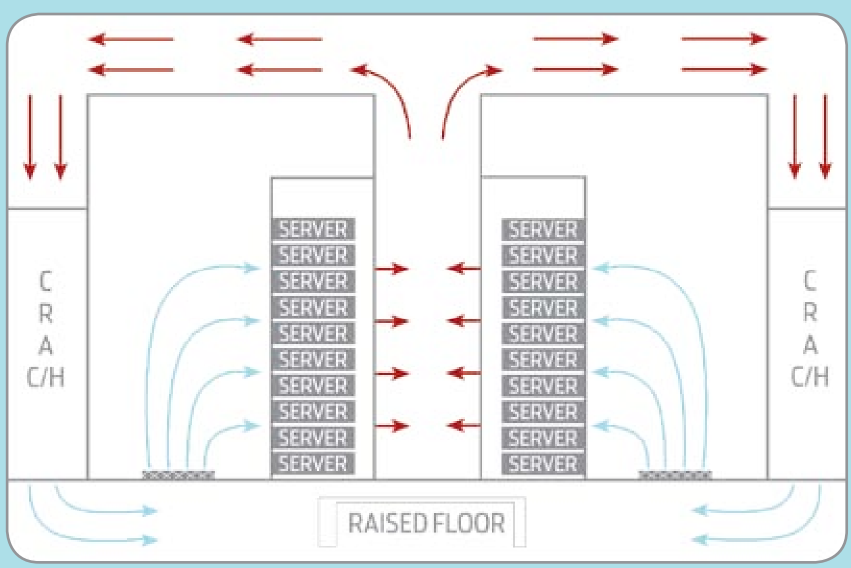

Figure 1: Typical hot aisle containment

higher than for CAC. In the event of a cooling failure, this will provide additional operating time to take action. Finally, as server cooling efficiencies continue to improve, the required cooling CFM per kW has significantly reduced, increasing the temperature rise across the server with exhaust temperatures now reaching the maximum server case specification of 55 °C (131 °F). With CAC, the user and any noncontained equipment will be exposed to this high temperature condition, which can be detrimental to both the operator and equipment performance. When properly implemented, HAC provides a few advantages over CAC containment, but both offer the advantages of increased rack density if properly integrated.

FULLY CONTAINED VS. PARTIALLY CONTAINED

We can conclude from the previous study that HAC offers a few select advantages over CAC. Building on this information, it would now be prudent to focus on the various HAC implementation strategies. This can be broken down into two classifications of designs—fully contained and partially contained systems. In the fully contained environment, nearly all of the return air is ducted directly into the cooling equipment. Chimney cabinets and rigid wall stationary systems are two examples of fully contained systems. In the partially contained environment, space or gaps exist that reduce the overall supply air reaching the IT equipment. Common partially contained solutions include rigid and soft wall curtain solutions.

While fully contained solutions will allow the greatest improvement in supported rack density, it is not always possible due to budget constraints or existing facility infrastructure. If this is the case, the simplest solution is to use hanging partitions. Both hard wall and soft wall hanging partitions, while not fully eliminating recirculation and bypass, will contain a significant portion of the hot return air. The challenge of hot spot management may still exist, and special attention should be made for fire prevention by using fusible links or other fire prevention tactics. It is safe to assume that a significant improvement in rack density could be obtained over the legacy data center by using a partial solution. However, to reach the 2014 density goal of 37 kW per rack, it is very likely that a fully contained strategy will be required.

FULLY CONTAINED IN-ROW AND OVERHEAD SOLUTIONS

For a data center adopting a fully contained strategy, one available option is in-row cooling. In-row cooling solutions require a stationary, fully encapsulated hard paneled structure to separate the hot return from the cold supply. The advantages of this system are that a raised floor or drop ceiling is not required, and that the cooling coil is located closer to the IT equipment, reducing the airflow path.

The challenge with in-row cooling is that you lose the energy efficiencies gained by using large air handling systems. For example, Table 1 illustrates the airflow delivery capabilities of four various fan sizes. The arbitrary operating pressure of 0.6 inches of water (inH_20) was intersected with the pressure versus flow curve of each fan to find the CFM delivered for a given power of the fan. This analysis demonstrates how larger fans deliver significantly greater CFM per watt than the smaller fans. Additional challenges associated with the in-row cooling include lack of redundancy and air side economization, reduction of rack space, and the rare potential of coolant leaks close to the IT equipment. In-row cooling is a viable means of full containment, but the data center operator should analyze the challenges of the air delivery system to the in-row approach.

FULLY CONTAINED DIRECT EXHAUST DUCT AND IN-RACK SOLUTIONS

Solutions that do not require the entire row space to be segregated are commonly known as rack level solutions. These solutions are typically an extension of the existing rack. The most common form of

Part Number	Fan Speed (RPM)	Fan Size (mm)	Operating Pressure (inH_2O)	Power (Watt)	CFM Delivered	CFM/Watt
AFB0912GHE	5400	92x92x38	0.6	22.2	27	1.22
AHB1348EHE	5100	127x127x38	0.6	29.28	90	3.07
FFB1424GHE	5250	140x140x38	0.6	71.52	245	3.43
AHB1548GHG	5200	172x150x50.8	0.6	72.96	300	4.11

Table 1: Fan performance data for various fan sizes

Figure 2: Vertical exhaust ducts or "chimney cabinets" routing IT exhaust into the return plenum

rack level containment is the vertical exhaust duct, or better known as the chimney cabinet. The vertical exhaust duct features a passive duct that routes the IT exhaust vertically through the back of the enclosed cabinet and into a return plenum. Chimney cabinets offer a different set of advantages and challenges compared with the in-row solution. Ducting the exhaust of the equipment is a much simpler way to capture the return air compared with enclosing the entire row. It also is regarded as the most contained of all the solutions.

Because the ducting solution is not tied into the row cooling system, the implementation can be completed in a piece part approach, allowing for expansion on an as-needed basis. Additionally, because the supply is detached from the row, the data center can take advantage of the performance efficiencies exemplified by Table 1 by using large air handling systems. The supply can be delivered by raised floor and floor tiles (the most common methodology) or directly through the cold aisle, which both allow for an effective path to air side economization.

The first challenge with the passive vertical exhaust duct solution is that a return plenum is required, using either ceiling return or ducted return. This requires additional floor-to-ceiling space, which is not necessarily required by the in-row solution. The second challenge is minimizing the external pressure imparted from the enclosure onto the IT equipment. The external pressure through the cabinet and duct will be a function of the CFM through the cabinet and the cross-sectional area at the rear of the cabinet and in the duct, as well as the overall flow length of the cabinet and duct. Care should be taken in the design of the duct and cabinet to ensure that the vertical exhaust duct does not impart excessive external pressure on the IT equipment. For most conditions this is minimal, and a passive solution is all that is required.

For the extreme case where the IT equipment is excessively long, or the rear cross-sectional area is very tight, an active chimney solution may be required. An active solution adds an additional system fan into the chimney, which acts like a booster fan to make up for the impedance of the small exhaust cavity. Adding an additional active component into the system adds a layer of inefficiency and complexity not required by the passive chimney and should only be used where necessary.

The remaining in-rack solution is a liquid cooled door, typically located in the rear of the rack. For this solution, the IT return air is passed through the cooler coils at the rear of the chassis, cooling the air to the required condition. A variation of this solution is one in which the return air is ducted back into the intake of the IT supply, allowing for a complete isolated loop. The advantages of this system are similar to the in-row cooling solution in that the cooling coil is brought close to the IT load and the return room temperature approaches room neutral. The challenge to this approach is minimizing the external pressure loaded onto the IT equipment due to the presence of the coil, the expense for plumbing and running cooling fluid to the cabinet, and similar to the in-row cooler, the rare chance of leaking close to the IT equipment. The three fully contained in-rack solutions each offer unique opportunities that optimize rack density, depending on the room and cooling conditions required.

CONCLUSION

HAC and CAC both represent opportunities to increase rack densities by improving airflow management and maximizing return temperatures, providing an efficient path to support the latest generation of IT systems. While unique features are offered by both solutions, HAC provides additional value in thermal capacitance, scalability and room comfort. Two viable solutions for the data center operator to consider include unique HAC solutions like in-row cooling that has advantages if floor-to-ceiling height is constrained and in-rack cooling solutions like vertical exhaust duct that offers energy efficiency if a drop ceiling or return duct is available. ■

REFERENCES

1. Patterson, M.K. and Fenwick, D. The State of Data Center Cooling, Intel White Paper, 2008.
2. ASHRAE. Datacom Equipment Power Trends and Cooling Applications, 2005
3. Garday, D. and Costello, D. Air-Cooled High-Performance Data Center: Case Studies and Best Methods, Intel White Paper, 2006.
4. Moss, D.L. Universal Application and Benefits of Air Containment: A Practical Guide, 2009.
5. Hilss, R. and Lyer, S. Improved Data Center Efficiency – Incorporating Air Stream Containment, HP White Paper, 2009.
6. Delta Electronics. http://www.delta.com.tw/product/cp/dcfans/dcfans_main.asp,2010.