Hot Aisle vs. Cold Aisle Containment

By John Niemann

White Paper #135

by Schneider Electric

Executive Summary

Both hot and cold air containment can significantly improve the predictability and efficiency of data center cooling systems. While both approaches eliminate the mixing of hot and cold air, there are practical differences in implementation that have significant consequences. This paper examines both methodologies and highlights the reasons why hot aisle containment emerges as the preferred best practice.

Introduction

High energy costs and spiking data center energy consumption rates have forced data center professionals to rethink their data center cooling strategies. Although traditional cooling approaches (such as perimeter cooling through a raised floor plenum) are still quite prevalent, new approaches such as hot aisle and cold aisle containment are making significant inroads. According to Bruce Myatt of EYP Mission Critical, the separation of hot and cool air "is one of the most promising energy-efficiency measures available to new and legacy data centers today." (Mission Critical Magazine, Fall 2007).

Much confusion exists as to how these new approaches work and how hot and cold aisle containment are different from or similar to traditional data center cooling methods. This paper reviews the differences between hot and cold aisle containment systems (CACS) and highlights hot aisle containment system (HACS) efficiency, operational flexibility, and reliability benefits. CACS emerges as a solution that offers improvements over traditional methods while still inheriting some of the flaws of the traditional approach.

***Figure 1** – Traditional cooling floods the entire space and mixes hot and cold air*

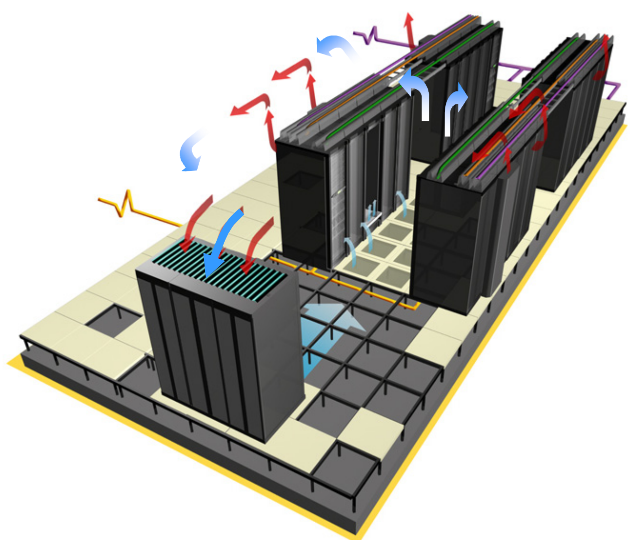

Historical Data Center Room Cooling Methods

Traditional cooling approaches share the following characteristics (see **Figure 1**):

- Perimeter cooling (CRAC units are placed on the outer perimeter of the rows of racks)
- Raised floor (cold air is delivered to the rows of racks via a plenum under the raised floor)
- The hot exhaust air from IT equipment mixes with cold air as the hot air makes its way back to the CRAC unit return vents.
- Rows of racks are not set up in a consistent hot aisle / cold aisle arrangement.
- Oversized power and cooling components reduce the data center efficiency.

The traditional environment developed and became accepted business practice because, for many years, rack densities were low (below 2 kW per rack), energy costs were negligible (in fact the fuel bill was invisible to most IT organizations), and a culture of “oversizing” hardware became a commonplace practice for reducing the risk of capacity shortage or downtime.

Now most data center professionals recognize that these past practices are inefficient, costly to the organization, and wasteful from a carbon footprint point of view. Two recent technology breakthroughs have helped to remedy this situation: row-based cooling and separation of hot and cold air streams. Row-based cooling brings the cooling source in very close proximity to the load (by being imbedded into the server rows) so that energy is not wasted forcing cold air across long distances, under a congested raised floor, to the load.

Figure 2 *– Cold aisle containment (CACS) deployed with a perimeter cooling approach*

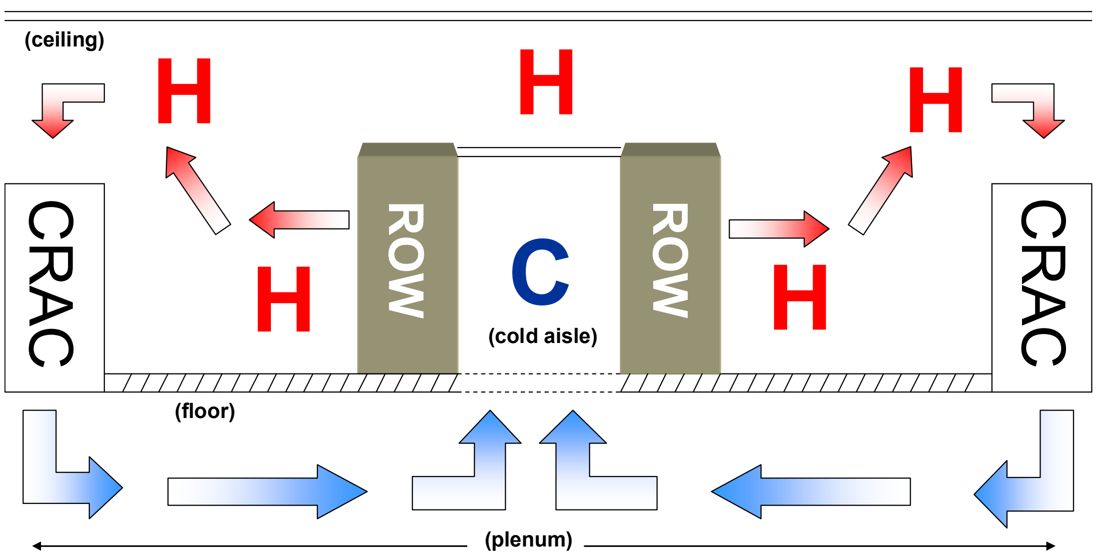

Containment, both hot aisle and cold aisle, offers the following benefits:

- **Cooling systems can be set to a higher temperature (thereby saving energy) and still supply the load with safe operating temperatures.** The temperature of traditional cooling systems are set much lower than required by IT equipment (i.e. approx 55°F / 13°C). This is done to compensate for heat picked up by the cold air as it makes its way from the CRAC unit to the front of the racks.

- **Reduction of humidification / dehumidification costs -** Typically when hot IT equipment exhaust air is captured and sent directly back to the CRAH unit, no humidity is removed from the air. If no humidity is removed, than adding humidity is not required which saves energy and water.

- **Better overall physical infrastructure utilization which enables rightsizing, which, in turn, results in equipment running at higher efficiencies -** Larger oversized equipment experiences larger fixed losses than right-sized equipment. However, over-sizing is necessary for traditional cooling because the extra fan power is required to both overcome under the floor obstructions and to pressurize the raised floor plenum.

Cold Aisle Containment

The Cold Aisle Containment System (CACS) is typically deployed in traditional perimeter-based cooling environments. Traditional cooling environments use the entire room as a hot air return plenum and use deliver cold air via the raised floor plenum to the cold aisles. The CACS encloses the cold aisle allowing the rest of the data center to become a large hot air return plenum. By containing the cold aisle, the hot / cold air streams within the data center are separated.

Figure 2 shows a fundamental illustration of how CACS works. Some home grown solutions are being deployed where data center operators are taking various types of plastic curtain material suspended from the building structure to enclose the cold aisle (see **Figure 3**). Some vendors are now beginning to offer ceiling panels and end doors that mount to adjoining racks for data centers to help separate the cold aisles from the rest of the warm air circulating in the data center space.

Figure 3 – *Example of a home grown cold aisle containment system*

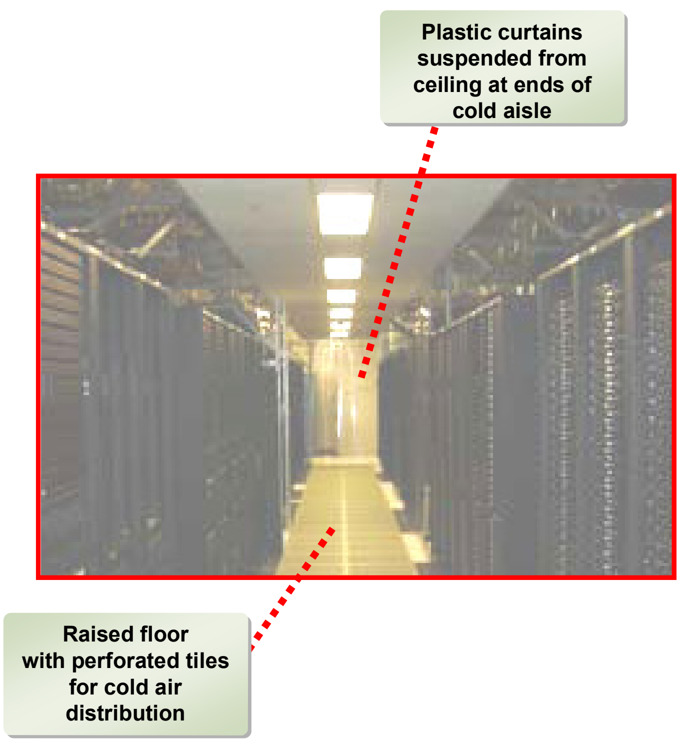

CACS efficiency limitations when deployed in a room-based approach

Although a CACS does offer efficiency improvements over traditional cooling, it does exhibit some drawbacks when deployed in a room-based, perimeter cooling environment:

- **Inefficiencies resulting from distances and pressures required for adequate air distribution -** The single largest contributor to inefficiency in a room-based approach is the requirement to move cold air from a perimeter CRAC unit to a distant load. A row-based cooling approach brings the source of the cooling in close proximity to the load. As a result, much less energy is required to deliver the cold air to its destination. This is not the case if the data center owner chooses to deploy CACS with row-based cooling. See APC White Paper # 130, "The Advantages of Row and Rack-Oriented Cooling Architectures for Data Centers" for a detailed comparison of row and room based cooling approaches.

- **Density limitations of using cold air distribution through raised floor -** The practical density limit when using a CACS approach is approximately 6 kW per rack. See APC White Paper # 46, "Cooling Strategies for Ultra High Density Racks and Blade Servers" for details regarding the reasons for this limitation. Higher densities can be achieved only if an investment is made in a customized design. To address some of the limitations of the raised floor plenum, some CACS solutions are offered with fan powered floor tiles. This improves airflow for higher density racks. The use of additional fan assisted devices further reduces the efficiency of the CACS. The extra fans contribute to the overall power consumption and add heat to the supplied cold air. The efficiency gains achieved by the CACS are therefore diminished by the added requirement for floor-based fans. This density limitation can be avoided if the data center owner chooses to deploy CACS with row-based cooling.

Figure 4 *– Raised floors become congested as data center requirements change*

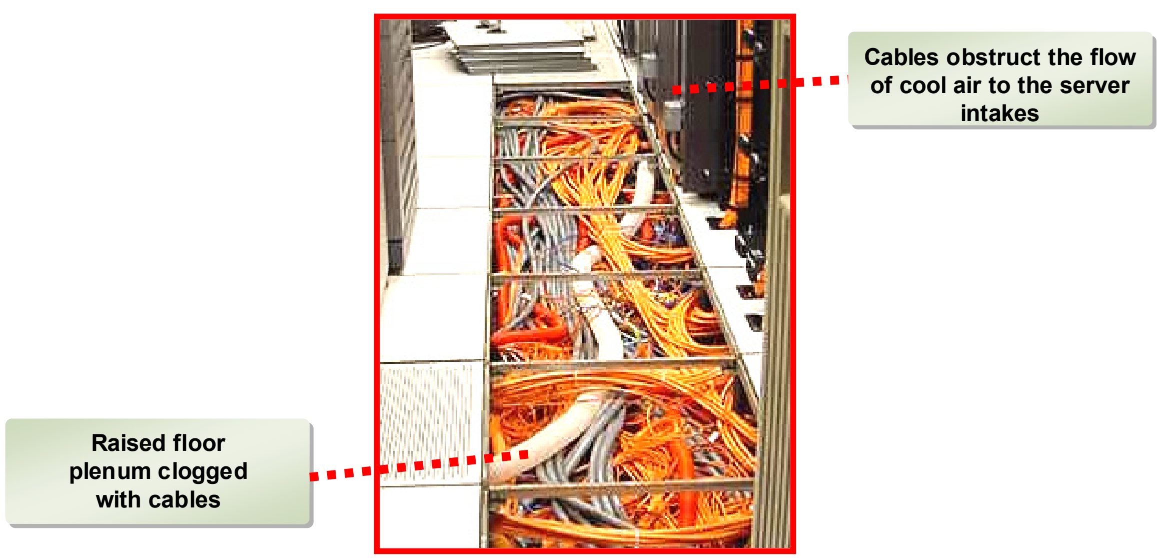

- **Predictability of the raised floor -** Cold aisle containment helps improve predictability through the elimination of hot and cold air mixing. However, it does not eliminate the variable of the raised floor. Cabling, piping, and other obstructions are added below the raised floor as the data center evolves. These obstructions limit the delivery of sufficient cool air to the IT equipment. **Figure 4** provides an example of how air dams in the plenum under the raised floor can hinder the predictable delivery of cold air to the cold aisle. This is not the case if the data center owner chooses to deploy CACS with row-based cooling as the need for the raised floor might not exist.

CACS limitations when deployed in a row-based cooling approach

A row-base deployment of CACS is more advantageous, from an efficiency perspective than a traditional room-based approach. However some limitations still exist:

- **Availability of cold air during a loss of power / cooling -** Containing the cold aisle minimizes the overall pool of cold air available to the servers, should a loss of power and / or cooling occur. The reduced volume of cold air results in more rapid temperature increases in the event of a failure. **Figure 5** depicts a sample data center and compares the volume of air in a contained cold aisle to the volume of air in an uncontained cold aisle. The **uncontained** cold aisle shows a volume of cold air that is 17 times greater than that of the cold air volume found in the **contained** cold aisle. This reduced air volume shortens the amount of time (seconds instead of minutes) it would take for the servers to overheat if a failure were to occur.

***Figure 5** – Comparison of cold air volumes in contained and uncontained cold aisles*

Cold Air Volume Sample Calculation

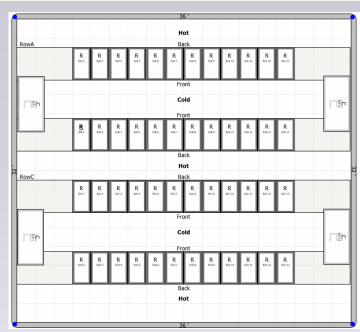

Room Dimensions:

- 36 ft(11m) x 31 ft (9.4M) x 10(3m) ft
- Cold Aisle Width: 4 ft (1.2m)
- Hot Aisle Width: 3 ft (0.9m)
- Rack Height: 42U – 6.5 ft (1.99M)
- Rack Width: 1.97 ft (0.6M)

Contained Cold Aisle Volume =

4ft x (1.97 ft x (12 racks per row)) x 6.5 ft = 614.6 ft^3 (17.2 M^3)

Room Volume (without hot aisles) =

(36 ft x 31ft x 10 ft) – (3 ft x (1.97 x12) x 10) = 10,450.8 ft^3 (3,185.4 M^3)

Uncontained cold air volume is **17 times** greater than cold air in contained cold aisle scenario [614.6 ft^3 (17.2 M^3) x 17 = 10,450.8 ft^3 (3,185.4 M^3)]

- **All of the cold aisles in the entire data center must be contained in order to realize benefits -** Containing only some of the cold aisles in the data center will yield little benefit because any other cold air that is allowed to mix with hot air will diminish any expected savings. Mixing will cause the cooling system to operate in a less efficient manner (a smaller difference between return "hot" air and cooling coil temperatures). To minimize mixing and to maximize cooling system efficiency, all cold aisles must be contained. Only then will hot return air temperatures reach maximum potential thereby allowing the cooling equipment to operate at much higher efficiency levels.

- **Overall perception and operation of a hot data center –** ASHRAE Standard TC9.9 recommends that server inlet temperatures range from 68-77° F (20-25° C). When cold aisles are contained, the air in the rest of the room becomes hotter (well above 80° F / 27° C and in some cases as high as 100° F / 38° C), and anyone entering the data center is exposed to unusually high temperatures. People are generally alarmed when entering such hot conditions, and tours become impractical. People's expectations need to be adjusted so that they understand that the higher temperatures are "normal" and not a sign of impending system breakdown. This cultural change can be challenging for workers not accustomed to entering a data center operating at higher temperatures.

  When operating a data center at elevated temperatures, special provisions need to be made for non-racked IT equipment. This is equipment that cannot be integrated into a CACS. Since, with a CACS system, the room is a reservoir for hot air, miscellaneous devices (such as tape libraries and standalone servers) will need to have unique ducting in order to enable them to pull cold air from the contained cold aisles. In addition, electric outlets, lighting, fire suppression, and other systems within the room will need to be evaluated for suitability of operations at elevated temperatures.

Hot Aisle Containment

The Hot Aisle Containment System (HACS) encloses a hot aisle to collect IT equipments hot exhaust air and cools it to make it available for IT equipment air intakes. This creates a self-contained system capable of supporting high density IT loads.

Mixing of hot and cold air streams in the data center lowers availability of IT equipment. Returning the warmest possible air to the computer room air conditioners increases the efficiency and capacity of the system. The HACS ensures proper air distribution by completely separating supply and return air paths.

The design of HACS assimilates many of the advantages of the CACS and avoids many of the pitfalls. When upgrading a data center to be more efficient and less costly to operate, any move away from the traditional perimeter cooling approach is a step in the right direction. While CACS is a "better" scenario compared to traditional approach, the "best" scenario is embodied in a HACS system.

A HACS system consists of doors at either end of the row, a roof for containing hot air, in-row cooling with variable speed fans, and a temperature controlled air supply to the cold aisle (see **Figure 6**).

Figure 6 *– Hot aisle containment system (HACS) operating as an independent zone*

Hot aisle containment system (HACS) efficiency benefits

Efficiency – The efficiency of the HACS will be higher because the hot aisle is capable of maintaining higher temperatures. In a typical high density server environment the temperature difference between the server exhaust air and the room temperature is typically around 30° F / 17° C). If the room is maintained within ASHRAE TC9.9 standards at 72° F / 22° C, a 30° F / 17° C temperature difference, would yield a server exhaust air temperature of 102° F / 39°F.

In a typical HACS environment the cooling units tend to move slightly more air than the servers and draw in a small amount of room air into the hot aisle. The effect of this can cause a slight reduction (about 2° F) in the return temperature, yielding a return temperature to cooling units of 100° F / 38°C. The net effect of this elevated return temperature (i.e., 100° F / 38° C) to the cooling unit enables better heat exchange across the cooling coil, better utilization of the cooling equipment, and overall higher efficiency. **Figure 7** gives an example of the effect of elevated return temperatures on sensible cooling capacities (the ability of an air conditioning system to remove heat from the air).

The effect of increasing return temperatures on cooling unit capacity holds true for virtually all air conditioning equipment. Some equipment may have limits as to the maximum return temperature they can handle, but, in general, all cooling systems will yield higher capacities with warmer return air.

In the case of the HACS, hot aisles operating at 100° F / 38° C with high density servers is typical. Contrast this with a cold aisle contained room where the entire room space would have to be maintained at 100° F / 38° C in order to achieve the same level of efficiency. While CACS would enable higher return temperatures, the typical data center operator will not operate the entire data center room at 100 F / 38° C in order to achieve the same efficiency as a HACS.

***Figure 7** – Effect of increased return temperature on sensible cooling capacity**

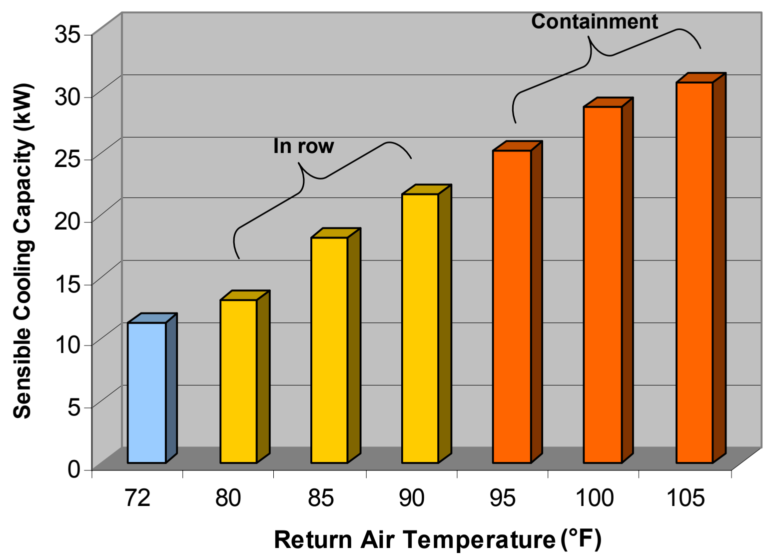

* APC In-row RC model air conditioner

Improved Flexibility – Unlike CACS, a HACS does not impact the temperature of the surrounding room. A HACS is, in effect, a room neutral solution. For example, if the temperature of the data center is set for 75° F (24° C), and a CACS system is implemented, the room temperatures outside of the cold aisle will rise because hot air will mix in with the air outside of the cold aisle on its way to the intake of the cooling system. The hot air inside of the HACS is contained from the rest of the room. The HACS does not deliver any hot air to the outside room; therefore the existing cooling system is not rendered less efficient.

A HACS can be "dropped in" to the data center without requiring any changes to the existing data center cooling architecture. When utilizing a row-based cooling approach (as opposed to a room-based approach), no need exists for the installation of specialized duct work and no adjustments need to be made to the existing HVAC systems to handle elevated return temperatures.

Higher Availability – The "Cold Air Volume Sample Calculation" in **Figure 5** demonstrates the differences in cold air volume when comparing CACS volume to room volume (uncontained cold air volume is 17 times greater than cold air in a contained cold aisle). This difference has a significant impact on the ability of the systems to support a cooling failure (i.e., runtime). A runtime that could be minutes in an uncontained room scenario might only be seconds if a CACS approach is deployed. With HACS only the hot air is contained, leaving the rest of the data center environment cool. Therefore, the servers will draw air from a larger pool of cool air outside the contained hot aisle, thereby extending available runtime.

Table 1 *– Cold aisle containment vs. hot aisle containment summary*

Characteristic	Cold Aisle Containment	Hot Aisle Containment	Comment
Efficiency improvements	Yes	Yes	HACS is more efficient than CACS because HACS typically will operate at higher return temperatures due to isolation of the hot air from the rest of the room.
Ability to increase cold air supply set point without impacting entire data center	No	Yes	With HACS cooling set points can be set higher while still maintaining a comfortable work environment. Increasing CACS cooling set points results in uncomfortably high data center temperatures.
Leverages maximum number of potential free cooling days	No	Yes	By increasing cooling set point containment systems allow for increased free cooling. However, Increasing the set point of CACS results in increased room temperatures which is undesirable from a free cooling days perspective.
Room neutral solution	No	Yes	A HACS deployment is a "drop in" solution. CACS impacts the surrounding data center infrastructure.
Ease of deployment with room cooling	Yes	No	CACS is preferred when using room level cooling with a free return system which draws its return air from the room. A HACS without in-row cooling would require special return duct work or ceiling plenum.
Ability to scale for high density	No	Yes	CACS is often implemented with raised floor and inefficient fan assisted floor tiles are needed in order to achieve higher density.
Room neutral design	No	Yes	HACS is room neutral – it doesn't impact the outside room temperature in any way. CACS makes the air outside of the contained rows hotter.
Adverse temperature impact on non-racked equipment	Yes	No	With CACS, because the cold aisles are contained, the rest of the data center is allowed to become hot. Equipment outside of contained areas would have to be evaluated for operation at elevated temperatures.

Fire Suppression Considerations

Depending upon the location of the data center, fire detection and / or fire suppression may be required inside the enclosed area of the HACS or CACS. The primary suppression mechanism is usually sprinklers, which are heat activated. Gaseous agents are usually a secondary system which can be initiated by smoke detectors. The National Fire Protection Association standard NFPA 75 does not state an opinion as to whether sprinklers or gaseous agents should be provided in a HACS or CACS. However, NFPA 75 documents the following two requirements that could be applied to HACS / CACS:

- Automated information storage system (AISS) units containing combustible media with an aggregate storage capacity of more than 0.76m^3 shall be protected within each unit by an automatic sprinkler system or a gaseous agent extinguishing system with extended discharge. (Note: This information is significant because it sets a precedent for fire detection and suppression in an enclosed space within a data center).
- Automatic sprinkler systems protecting ITE rooms or ITE areas shall be maintained in accordance with NFPA 25, *Standard for the Inspection, Testing, and Maintenance of Water-Based Fire Protection Systems*

In practice, HACS have been successfully installed and approved with sprinklers and gaseous agent suppression in many sites. The AHJ should be contacted for specific requirements in a given location.

Conclusion

Prevention of hot and cold air mixing is a key to all efficient data center cooling strategies. Hot aisle containment (HACS) is a more efficient approach than cold aisle containment because the HACS methodology allows for the channeling of the hottest air directly into the coolers. Cooling set points can also be set higher while still maintaining a comfortable work environment.

Since it does not affect its outside environment, a HACS solution can also be deployed anywhere within the room. With a HACS solution, equipment inlets are uncontained and therefore can draw cooling airflow from the room in event of a cooling failure. This allows more time for a switch to generator or for a graceful shutdown of servers.

Both HACS and cold aisle containment (CACS) offer superior power density and efficiency when compared with traditional cooling approaches. CACS can offer some improvement in a traditional room-based perimeter cooling layout. However, HACS with a row-based cooling architecture is more efficient, more flexible, provides better ride through capability, and can more easily address the higher IT density requirements without increasing the temperature of the entire data center. For most users the additional heating of the uncontained operator space is an unacceptable condition which eliminates CACS as an option. For the above reasons, most high efficiency, high density data center projects, both for new designs and retrofits, incorporate some form of hot aisle containment.

About the author

John Niemann is the Product Line Manager for Row and Small Systems Cooling Products at American Power Conversion (APC), and is responsible for planning, support, and marketing for these product lines. John has led product management for all of APC's InRow Cooling products since 2004. He has 12 years experience in HVAC. His career began in the commercial and industrial HVAC market where he focused on custom air handling and refrigeration systems, with expertise focused on energy recovery and filtration for critical environments. His HVAC experience spans applications engineering, development, product management, and technical sales. John is a member of ASHRAE and the Green Grid, and holds a degree in Mechanical Engineering from Washington University in St. Louis, MO.