

Data Center Containment 101

SUBZERO / WHITE PAPER

By Gordon Johnson



INTRODUCTION

Regardless of if we're entering a data center for the first time or have been doing so for years, most data centers have something in common. As you walk through rows of racks, you'll alternate between cold and hot aisles. You'll hear expressions like "CRACs", "PUE", "White Space", "Cold Aisle Containment", "Hot Aisle Containment", and many more. The purpose of this White Paper is to assist those new to the data center and those assigned with making key decisions to get the most out of existing "legacy" and newly designed data centers.

Since energy efficiency and data reliability are key goals for anyone managing or associated with data centers, how can we achieve both in the shortest amount of time while getting the quickest ROI (*Return of Investment*)? When is it more appropriate to use one type of containment instead of another type? Which saves more money? This paper will answer these and other questions while discussing the following topics:

- Data Center Terminology
- Typical Data Center Cooling Methods
- Bypass Air, Recirculation Air, and Once-Through Cooling
- PUE (*Power Usage Effectiveness*)
- Containment Basics
- Cold Aisle Containment Versus Hot Aisle Containment
- Full Containment Versus Partial Containment
- Savings: Cold Aisle Containment Versus Hot Aisle Containment

DATA CENTER TERMINOLOGY

A complete list of all the terminology used in, and to describe, data centers is beyond the scope of this paper, but there are several basic expressions that are good to be familiar with. For example, when describing the type of data center, what is the difference between a brownfield and a greenfield site? When inside, what is the difference between white space and gray space?

A brownfield deployment is the installation and configuration of new hardware or software that must coexist with existing or legacy IT equipment. A greenfield deployment, in contrast, is the installation and configuration of hardware or software that a company has not used before and is not dependent upon existing or legacy technology. A greenfield data center is built from the ground up and is not subject to the constraints posed by the existing building or data center. With data centers, brownfield deployments can be cost-effective because the infrastructure to support the new installation likely already exists. However, brownfield deployments can be complicated by the need to rectify dependencies between the new and current installations.¹

Some data centers are described as “Green” or simply called green data centers. A green data center is a facility that is entirely built, managed, and operated on green principles. It provides the same features and capabilities of a typical data center, but uses less energy and space, and its design and operation are environmentally friendly. The entire infrastructure is installed with the lowest power and carbon footprint.²

The inside of a data center has some unique expressions as well. For example, every data center has both white space and gray space. White space in data centers is the space allocated for IT equipment or the space where IT equipment is placed. It includes servers, storage, network gear, racks, air conditioning units, and power distribution systems. Gray space in data centers is the area where back-end equipment is located. This includes switch gear, UPS units, transformers, chillers, and generators.³

TYPICAL DATA CENTER COOLING METHODS

It takes an enormous amount of energy to cool a data center. According to Global Market Insights, on average, cooling systems represent 40% of total data center energy consumption, with data centers utilizing approximately 3% of the global electrical supply.⁴ In order to increase energy efficiency and reduce the carbon footprint, it's essential not to waste energy while cooling the data center. Before discussing methods to reduce energy consumption, let's first review some typical data center cooling methods. This is by no means a complete list, just some of the more common methods being used in the industry:

1. **CRAC (DX Cooled):** The computer room air conditioner (CRAC) provides cooling by blowing air over a cooling coil filled with refrigerant. A CRAC has a direct expansion (DX) refrigeration cycle built into the unit, meaning that the compressors required to power the refrigeration cycle are also located within the unit. CRACs are typically constant volume meaning they're either ON or OFF. CRACs use an air-cooled condenser located on the building roof to reject heat into the outside atmosphere.
2. **CRAC (Glycol Cooled):** Unlike air-cooled DX systems, this system uses flowing glycol (*mixture of water and ethylene glycol, similar to automobile anti-freeze*) and has a dry cooler on the building roof to reject heat into the outside atmosphere.
3. **CRAH (Water Cooled):** The computer room air handler (CRAH) provides cooling by blowing air over a cooling coil filled with chilled water. Chilled water is typically supplied to the CRAH by a chilled water plant (*i.e. chiller and cooling tower*) and the heat is rejected to the outside atmosphere via the cooling tower. CRAHs typically have Variable Frequency Drives (VFDs) that modulate fan speed to maintain a desired static pressure beneath the raised floor or in the overhead ducts, resulting in efficiency savings since fan speeds can be reduced while still supplying sufficient cold supply air to the IT equipment.
4. **IRC:** In-Row Cooling (IRC) is where the cooling unit(s) are placed between the server racks in a row for supplying cold air to the IT equipment. These units can be DX, water, or glycol and typically have VFD fans to vary the airflow volume. IRCs use a horizontal airflow pattern that supports the conventional hot aisle/cold aisle layout. IRCs can be used in data centers with or without raised floors.
5. **Free Cooling:** Air-side free cooling is where outside air is brought into the data center directly through filters or indirectly through heat exchangers. Water-side economizers and adiabatic cooling are also referred to as free cooling.

6. **Raised Floor:** Cooling for this environment is delivered via downflow CRAC/CRAH units directly under the data center floor. The cold supply air is pushed up from under the floor through perforated tiles typically placed in front of each rack or row of racks.
7. **Slab Floor:** Some data centers are designed with concrete slab floors, where the cooling is delivered via CRAC/CRAH units and forced down from above via supply ducts and diffusers to the front of racks, or delivered horizontally via IRCs within the cold aisle(s). In addition, XDO (*Overhead Cooling Module*) and XDV (*Vertical Top Cooling Module*) units typically use slab floors and deliver cold supply air from above.

BYPASS AIR, RECIRCULATION AIR, AND ONCE-THROUGH COOLING

In order to have an efficient and reliable data center, it's important to understand bypass air and recirculation air and the negative impact they have on the data center.

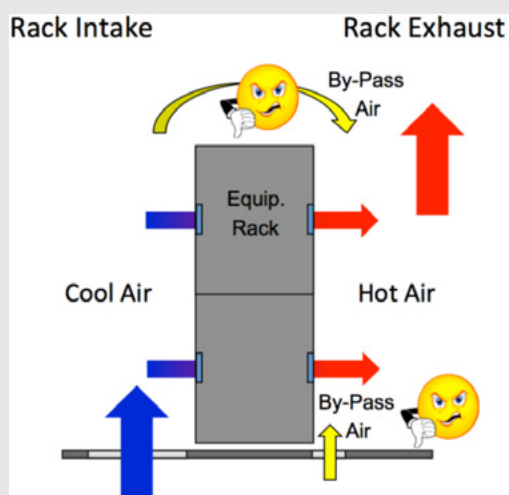


FIGURE 1: BYPASS AIR
Graphics courtesy of DCEP Training Program⁶

Bypass air is defined as air that does not participate in cooling the IT equipment in the rack (see Figure 1). As its name suggests, the cold supply air bypasses the server racks and blows back to the cooling units. Bypass air is wasted air which is an expensive problem because it costs money to cool and blow cold air. Bypass air typically occurs when CRACs are providing an excess of cold supply air to the servers (the acronym CRAC is used throughout this paper regardless of referring to a Computer Room Air Conditioner or a Computer Room Air Handler). Failure to follow airflow management “best practices” will also contribute to this problem, such as misplaced perforated tiles, missing blanking panels, cable cutout floor leakage, etc.

A good analogy is running the air conditioner for your car but leaving all of the windows and doors open at the same time. Bypass air equals wasted money.⁵

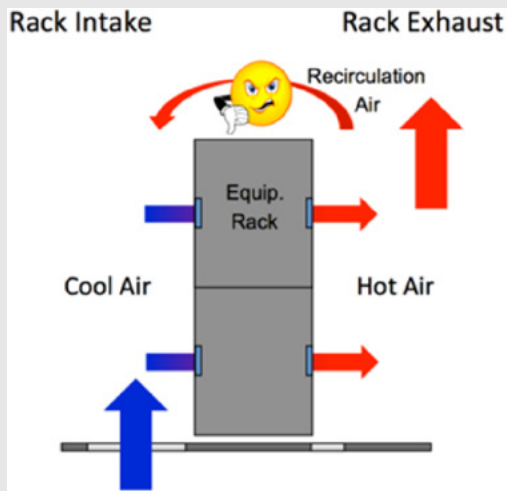


FIGURE 2: RECIRCULATION AIR
 Graphics courtesy of DCEP Training Program⁶

Recirculation air is when there's not enough supply air to satisfy the servers, so the hot exhaust air recirculates above the racks and around the end of the aisles into the server inlets, resulting in overheating of the IT equipment (see Figure 2). Recirculation air can also be caused by missing blanking panels, gaps between racks, etc.

A good analogy is trying to cool your car by turning on the vent allowing outside air in to the car but not turning on the air conditioning unit itself. The result is simply recirculating the hot air. Recirculation air puts the IT equipment at risk of heat related failures.⁵

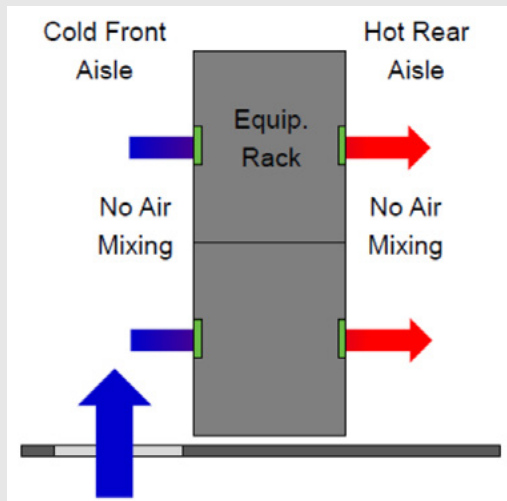


FIGURE 3: ONCE-THROUGH COOLING
 Graphics courtesy of DCEP Training Program⁶

The goal is to get the cold supply air as close to the IT equipment as possible without mixing with the hot exhaust air. The IT equipment is designed to pull in the cold supply air at the front of the rack (*air intake*), use it to cool the IT equipment, then evacuate the hot air (*air exhaust*) from the rear of the rack and move it back to the CRAC where it is cooled again and the process starts over.⁵ This is referred to as one-through cooling (see Figure 3).

PUE (POWER USAGE EFFECTIVENESS)

The solution to managing and eliminating bypass and recirculation air starts with containment. As we'll see in the next section, containment isolates the cold supply and hot exhaust from each other, resulting in increased efficiency. Without containment, it's not uncommon for data centers to supply 2X or more the required airflow to the IT equipment to negate the effects of recirculation air (*server hot spots*). Containment, and following airflow management "best practices", results in most data centers reducing airflow so that the demand airflow is approximately 85-90% of the supply airflow required to cool the IT equipment.

PUE (*Power Usage Effectiveness*) is a popular metric used to determine the energy efficiency of the data center. PUE measures the efficiency of energy use by showing how much energy is actually consumed by the IT equipment versus the data center support equipment (*cooling, airflow, power distribution, lighting, security, etc.*). After containment, and depending on the reduction of supply airflow and how high the supply air temperature to the IT equipment is increased, there can be a significant reduction to the PUE.

$$\text{PUE} = \frac{\text{Total Facility Energy}}{\text{IT Equipment Energy}}$$

FIGURE 4

PUE is determined by dividing the amount of power entering a data center by the power used to run the computer infrastructure within it.⁷ PUE is expressed as a ratio, with overall efficiency improving as the PUE decreases closer to 1.0 (see Figure 4). A data center with a PUE of 1.6 or less is considered

efficient, and a PUE of 1.2 or less is considered very efficient.

With the exception of the IT equipment, the dominant part of data center power utilization has to do with cooling (*includes CRACs, IRCs, chillers, cooling towers, etc.*). Some estimate that for every kW of IT equipment you put into a room you burn another kW trying to get the heat back out of the room, so a PUE of 2.0 is not uncommon in data centers without containment.

CONTAINMENT BASICS




FIGURE 5: AISLEFRAME CONTAINMENT SYSTEM

Why should either cold aisle containment (CAC) or hot aisle containment (HAC) be installed in both existing and all new data centers? Containment makes existing cooling and power infrastructures more effective because it fully separates the cold supply and hot exhaust air (see Figure 5).

With containment, the data center makes increasingly efficient use of the same or less cooling, reducing the cooling portion of the total energy bill. Data centers can often power down CRACs, saving utility and maintenance costs. Containment allows for lowering CRAH fan speeds, higher chilled water

temperatures, decommissioning of redundant cooling units, and increased use of free cooling. A robust containment solution can reduce fan energy consumption by up to 25% and deliver 20% energy savings at the cold water chiller, according to the U.S. EPA. It can save a data center approximately 30% of its annual utility bill without additional CapEx.⁸

COLD AISLE CONTAINMENT VERSUS HOT AISLE CONTAINMENT

Containment is a key strategy implemented in data centers to improve efficiency and lower PUE. The goal is to raise server inlet temperatures, but to stay within the recommended temperature range. Regardless of the Server Class, ASHRAE (*American Society of Heating, Refrigerating, and Air-Conditioning Engineers*) recommends that the temperature at the server inlets be no higher than 80.6°F (27°C).

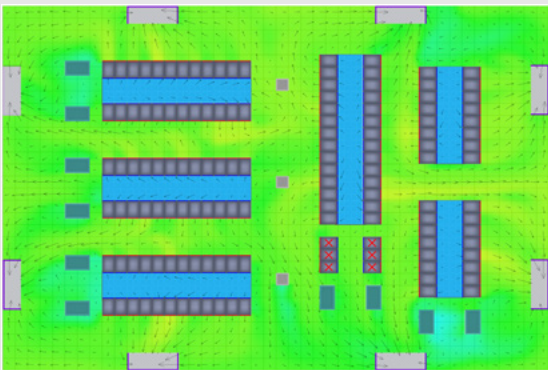


FIGURE 6: 6' HORIZONTAL XY PLANE WITH CAC

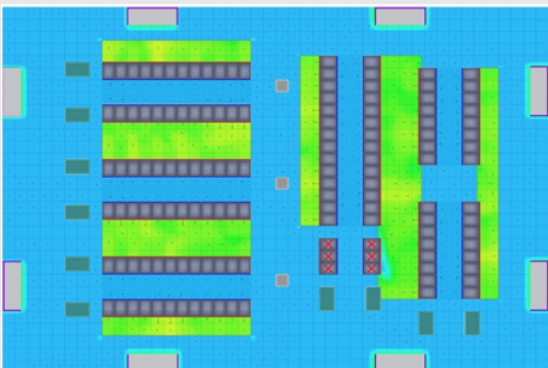


FIGURE 7: 6' HORIZONTAL XY PLANE WITH HAC

There are 2 types of containment. The first, cold aisle containment, encloses the cold aisle, ensuring that only cold supply air flows into the air intakes of the IT equipment, thus separating the cold supply and hot exhaust air. (see Figure 6) When you contain the cold aisles, the rest of the data center is basically one large hot aisle.

The second, hot aisle containment, encloses the hot aisle to collect the IT equipment's hot exhaust air, ensuring that only hot exhaust air returns to the CRACs, thus separating the cold supply and hot exhaust air. (see Figure 7) When you contain the hot aisles, the rest of the data center is basically one large cold aisle.

Which type of containment should be used usually depends on the layout and configuration of the data center and any constraints that may exist. For existing data centers, an assessment of the facility will be needed to choose the right containment solution. For example, here are some typical data center assessments and solutions. (In all examples, the data center already supports the conventional hot aisle/cold aisle layout.)

ASSESSMENT #1: The data center raised floor is being used to supply air through perforated tiles in front of the IT racks. There is no drop/return air plenum ceiling in place, with CRACs located around the perimeter of the data center. There are also overhead obstructions (*cable trays, etc.*).

SOLUTION #1: CAC should be used.

ASSESSMENT #2: Same as #1, only there are several stand-alone racks located on the data center floor.

SOLUTION #2: CAC should still be used due the lack of a drop/return air plenum ceiling. However, thought should be given on how to cool stand-alone equipment, such as using a stand-alone containment booth with enclosed perforated tiles.

ASSESSMENT #3: The data center raised floor is being used to supply air through perforated tiles in front of the IT racks. There is a drop/return air plenum in place that is already, or could be, used as a return back to the CRACs (*if currently unused, verify that the drop ceiling is rated for air distribution*). There are several stand-alone racks located on the data center floor.

SOLUTION #3: HAC should be used.

ASSESSMENT #4: There is no raised floor, the drop ceiling is currently being used as a return air plenum, and the CRACs are blowing supply air throughout the data center.

SOLUTION #4: HAC should be used.

ASSESSMENT #5: There may or may not be a raised floor, there may or may not be a drop/return plenum ceiling, and In-Row Coolers (*IRCs*) are placed within the rows of racks to supply cold air to the IT equipment.

SOLUTION #5: Either CAC or HAC can be used, although HAC is often preferred for this environment.

As can be seen from these sample assessments, typically the layout and configuration of the data center will determine which type of containment is used. But from a thermodynamics standpoint, it really doesn't matter; either type of containment will have similar results because they're ultimately both doing the same thing, preventing the cold supply and hot exhaust air from mixing in the data center. Both will improve efficiency and lower data center operating costs.

For cold aisle containment, it's important to have sufficient cold supply air within the contained PODs to ensure positive pressure across the racks. There should be slightly higher

supply air (CFM) than demand air (CFM) for the IT equipment for each separate contained POD to ensure proper pressure differential between the cold and hot aisles. This will result in a minimal temperature delta across the server racks (*top to bottom*) within the PODs.

FULL CONTAINMENT VERSUS PARTIAL CONTAINMENT

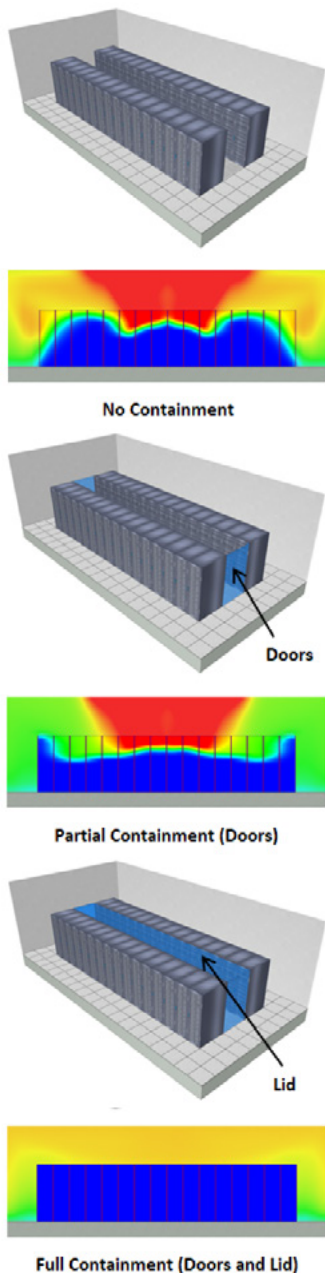


FIGURE 8: 3D AIRFLOW CONTAINMENT COMPARISON

What about employing partial aisle containment in a data center to improve efficiency? For example, is just adding doors to the end of aisles sufficient versus additionally adding a lid across the top of the cold aisle or running containment to the top of the ceiling?

Partial containment such as end of aisle doors will prevent hot air from recirculating around the end of the aisles, but will not have a large impact on improving any existing issues of hot exhaust air recirculating over the racks and back into the cold aisles. (see Figure 8) If possible, full containment should be installed since this is the best way to eliminate “hot spots” and improve efficiency in the data center.

What if overhead obstructions just below the ceiling, or even the fire suppression system, make it difficult or impossible to install full containment over the racks or up to the ceiling? Even containment 18” below the ceiling will have a substantial impact in preventing the mixing of the cold supply and hot exhaust air, thus improving efficiency.

As a side note, both full cold aisle containment and hot aisle containment have been successfully installed and approved with sprinklers and gaseous-agent suppression systems. However, it is always recommended that the Authority Having Jurisdiction (AHJ) be contacted for specific requirements well before any containment project is started.

SAVINGS: COLD AISLE CONTAINMENT VERSUS HOT AISLE CONTAINMENT

We know that containment increases efficiency and saves money, but is one type of containment better than the other? Which is more efficient, cold aisle containment or hot aisle containment?

As stated earlier, from a thermodynamics standpoint, both will improve efficiency and lower data center operating costs. Recent energy efficiency studies also show that the savings between the two are negligible. If given the choice, such as during a new data center design, a close look at some of the advantages and disadvantages between the two will help make the choice easier.

The cold aisle containment system separates the cold supply and hot exhaust air from each other, often with simple modifications to the room. The down side is that employees must work in the open warm space as cold supply air only cools the IT equipment.

The hot aisle containment system also separates the cold supply and hot exhaust air from each other, but one advantage to this type of system is for data center personnel. Hot aisle containment creates a pleasant working atmosphere for personnel as they walk into a cool room. Some feel they can increase the supply temperatures higher than they could with cold aisle containment and thus see larger energy savings.⁹ However, when taking the human element out of the equation, both systems are equally effective at providing optimized and well-managed airflow to the IT equipment.

CONCLUSION

Energy efficiency through the separation of cold supply and hot exhaust air can easily be attained by using containment in the data center. The use of cold aisle containment or hot aisle containment is typically determined by accessing the facility constraints and reviewing potential solutions. Regardless of which type of containment is used, large energy savings can be achieved by raising the supply air temperature and optimizing the airflow rate in the data center. Other benefits include higher energy densities per rack, longer Mean Time Between Failures (*MTBF*) of the IT equipment, lower PUE, and more. An ROI between 6 and 18 months is typical for most projects regardless if cold aisle containment or hot aisle containment is deployed.

Whenever possible, existing data centers should be retrofitted with either cold or hot aisle containment, with containment always automatically included as part of any new data center design.

ABOUT THE AUTHOR

Gordon Johnson is the Senior CFD Engineer at Subzero Engineering, and is responsible for planning and managing all CFD related jobs in the U.S. and worldwide. He has over 25 years of experience in the data center industry which includes data center energy efficiency assessments, CFD modeling, and disaster recovery. He is a certified U.S. Department of Energy Data Center Energy Practitioner (*DCEP*), a certified Data Centre Design Professional (*CDCDP*), and holds a Bachelor of Science in Electrical Engineering from New Jersey Institute of Technology. Gordon also brings his knowledge and ability to teach the fundamentals of data center energy efficiency to numerous public speaking events annually.

REFERENCES

1. Definition: Brownfield (brownfield deployment, brownfield site), Posted by Margaret Rouse; <https://searchdatacenter.techtarget.com/definition/brownfield-site>
2. Technopedia: Green Data Center; <https://www.techopedia.com/definition/14754/green-data-center>
3. What is White Space and Gray Space in data centers?, Posted on 6 Jul, 2016 by Storage Servers; <https://storageservers.wordpress.com/2016/07/06/what-is-white-space-and-gray-space-in-data-centers/>
4. Meeting Cooling Challenges in the Mission Critical Facility, Posted on 21 Aug, 2019 by Andrew Schaap; <https://www.missioncriticalmagazine.com/articles/92509-meeting-cooling-challenges-in-the-mission-critical-facility>
5. Data Center Efficiency: The Benefits of RCI & RTI (Part 2 of 3); <https://www.rfcode.com/data-driven-data-center/bid/233382/Data-Center-Efficiency-The-Benefits-of-RCI-RTI-Part-2-of-3>
6. Graphics Courtesy of the DCEP Training Program; <http://datacenters.lbl.gov/DCEP>
7. Definition: Power Usage Effectiveness (PUE), Posted by Margaret Rouse; <https://searchdatacenter.techtarget.com/definition/power-usage-effectiveness-PUE>
8. Data Center Knowledge: Approaches to Data Center Containment, Posted on 8 Nov, 2012 by Kevin Normandeau; <https://www.datacenterknowledge.com/archives/2012/11/08/approaches-to-data-center-containment>
9. DataCenter 2020 White Paper: hot-aisle and cold-aisle containment efficiencies reveal no significant differences, Posted by Intel; <https://www.intel.com/content/dam/www/public/us/en/documents/white-papers/data-center-efficiency-datacenter-2020-paper.pdf>