

Realtime  
publishers

*The Shortcut Guide<sup>™</sup> To*



# Data Center Energy Efficiency

*sponsored by*

**APC**<sup>®</sup>

by Schneider Electric

*David Chernicoff*

# Only APC InRow® cooling actively adjusts to eliminate hot spots.

Now, cool at the row level with the first-ever variable speed fan technology.

## Today's data centers are really heating up.

Racks are packed with more and more equipment, driving the highest-ever rack power densities. The result: unprecedented heat levels, row by row. Meanwhile, virtualization is everywhere, leading to more dynamic loads and shifting hot spots. Tackling this challenge with traditional raised floors and perimeter cooling alone poses a real struggle: How can you locate enough cooling exactly where it's required? Too often, the result is inefficiency, worsened by soaring energy costs. What's the efficient and effective solution? InRow® cooling from APC® by Schneider Electric.

## Variable-speed fans target heat and improve efficiency.

Rack-mounted sensors monitor the temperature, giving you real-time information on where heat is hiding. As heat loads shift around the room, unique variable-speed fans automatically adjust to meet the demand. By closely matching cooling with the heat load, you use the cooling that's required in the right place at the right time, reducing waste by preventing hot and cold air mixing and eliminating hot spots. You improve efficiency and avoid overcooling.

## Modular design delivers maximum flexibility.

Scalable, modular InRow cooling units can be easily deployed as the foundation of your entire cooling architecture or in addition to current perimeter cooling for a high-density zone within an existing data center. With this kind of hybrid environment, there is no need to start over, and installation is quick and easy.


So go ahead: Pack the racks without fear of hot spots or inefficiency. Intelligent, efficient, InRow cooling handles high-density heat at the source.



APC offers the most efficient, comprehensive line of cooling solutions for any IT environment.




Room-level cooling:  
InRoom Chilled Water, InRoom Direct Expansion, NetworkAIR® PA



Row-level cooling:  
InRow RC, InRow RD, InRow RP, InRow SC



Rack-level cooling:  
RackAIR® Removal Unit SX, RackAIR® Distribution Unit, Rack Side Air Distribution, Rack Fan Tray




Download a **FREE** copy of APC White Paper #137: "Energy Efficient Cooling for Data Centers: A Close-Coupled Row Solution".

Visit [www.apc.com/promo](http://www.apc.com/promo) Key Code **i204w** • Call **888-289-APCC x6031** • Fax **401-788-2797**



InfraStruXure®  
DATA CENTERS ON DEMAND

Achieve greater efficiency with InRow® cooling.



APC is proud to be a member of the green grid.

**APC**®  
by Schneider Electric

---

Chapter 3: Free, or Nearly Free, Things to Do in an Operating Data Center.....	35
The Traditional Cooling Model.....	35
Auditing .....	36
Find Out What’s Really Going On.....	36
Checking Capacity.....	37
Checking Hardware.....	38
Systems Testing.....	40
Understanding the Impact of Temperature on Servers.....	41
Finding the Sweet Spot for Minimizing Energy Consumption.....	41
How Managing Airflow Reduces Costs.....	41
Understanding the Use of Blanking Panels.....	41
Minimal Cost / Maximum Gain .....	42
Proper Placement of Vents and Tiles .....	43
Containment .....	44
Keeping the Cold Air, or Hot Air, Where You Need It.....	44
Evaluating Cooling Architectures .....	45
Conclusion .....	49

## Copyright Statement

© 2009 Realtime Publishers. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtime Publishers (the “Materials”) and this site and any such Materials are protected by international copyright and trademark laws.

THE MATERIALS ARE PROVIDED “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. The Materials are subject to change without notice and do not represent a commitment on the part of Realtime Publishers or its web site sponsors. In no event shall Realtime Publishers or its web site sponsors be held liable for technical or editorial errors or omissions contained in the Materials, including without limitation, for any direct, indirect, incidental, special, exemplary or consequential damages whatsoever resulting from the use of any information contained in the Materials.

The Materials (including but not limited to the text, images, audio, and/or video) may not be copied, reproduced, republished, uploaded, posted, transmitted, or distributed in any way, in whole or in part, except that one copy may be downloaded for your personal, non-commercial use on a single computer. In connection with such use, you may not modify or obscure any copyright or other proprietary notice.

The Materials may contain trademarks, services marks and logos that are the property of third parties. You are not permitted to use these trademarks, services marks or logos without prior written consent of such third parties.

Realtime Publishers and the Realtime Publishers logo are registered in the US Patent & Trademark Office. All other product or service names are the property of their respective owners.

If you have any questions about these terms, or if you would like information about licensing materials from Realtime Publishers, please contact us via e-mail at [info@realtimepublishers.com](mailto:info@realtimepublishers.com).

[**Editor's Note:** This eBook was downloaded from Realtime Nexus—The Digital Library for IT Professionals. All leading technology eBooks and guides from Realtime Publishers can be found at <http://nexus.realtimepublishers.com>.]

## Chapter 3: Free, or Nearly Free, Things to Do in an Operating Data Center

In an average data center, several steps can be taken to improve the cooling and operational efficiency of the data center. In the past, the data center was considered a single entity. This meant that cooling and energy delivery systems treated the data center as if conditions were equivalent throughout the space. The reality is that in most data centers, a number of different types of activities are being performed by the equipment that results in very specific sets of conditions. By grouping these different sets of activities by their need for power and cooling, it is possible to more efficiently deliver both services to the locations and systems within the data center that most need the support.

By understanding the conditions within your data center, it is possible to find efficiencies in the delivery of power and cooling that can be achieved using the existing infrastructure and making organizational or minor physical changes to the data center itself. Doing so, you will be able to optimize the delivery of power and cooling services to devices within, on a limited budget.

### The Traditional Cooling Model

In the traditional data center cooling model, cold air is delivered to the room in a manner designed to circulate the air across the devices that require cooling. This can be done via a raised-floor delivery system (as shown in Figure 3.1) or, in smaller data centers, with placement of computer room air conditioners (CRACs) and returns as appropriate.




Figure 3.1: Raised-floor cooling in a traditional data center.

Oftentimes, servers are mixed into the data center in no particular order, with high-load and light-load servers running side by side; there is no consistent methodology as to placement of racks and components, though ideally they are placed in locations that don't impede the air circulation within the facility.

Due to the minimal expense in providing power and cooling services to the data center, older data centers often just threw oversized cooling systems at the potential problem; a solution that is no longer cost effective with the rising energy costs.

## Auditing

Auditing is the process of identifying potential problem areas for cooling in the data center. At its most basic level, the purpose of the cooling audit is to make sure that cool air has an unobstructed path from the CRAC to the air inlet on the server and that hot air has an equally unobstructed return path.

Failure to provide sufficient cooling in the right places can result in damaged equipment or temperature-related electronic equipment failure. By auditing the cooling of the data center, these potential problems can be avoided.

The first step in the auditing process is to perform a baseline audit; by doing so, you will be able to measure the success of any corrective actions that you determine are necessary to deliver proper cooling. If this is a new data center, you will be able to use this data to establish a baseline that is easily used to determine the impact of future changes. In an existing data center, this baseline will give you the starting point for establishing the cooling efficiency of your facility. When you complete your changes and have resolved that your data center cooling model is operating at peak efficiency, you will have established a new baseline model that future changes can be evaluated against.

### Find Out What's Really Going On

Understanding how your cooling is actually working within your data center is critical for maintaining the long-term efficiency of not just the cooling equipment but also the hardware that is being cooled. Once you have established your data center cooling baseline, you will be able to maintain the cooling efficiency of your environment. The next section of this chapter will provide basic guidelines for performing a cooling audit in your facility. We will look at the fundamental tasks necessary to properly establish the baseline information required for future cooling audits.

### Checking Capacity

The first step in performing your cooling audit is the capacity check. This test determines whether you have sufficient cooling in place for your existing and planned environment. Remember that 1 W of power requires 1 W of cooling.

#### Note

Although there are many types of air conditioners, for the purposes of this chapter, we will use the generic term CRAC to refer to all types of computer room air conditioners.

Remember that our goal is to match the cooling capacity to the cooling need. In the previous chapters, we discussed the issues involved in oversizing the cooling capacity of your data center and pointed out the problems and inefficiencies that result from this condition.

By doing some research on the model of each CRAC unit, you should be able to find manufacturer's specifications that will establish the cooling capacity of each unit. The manufacturer will specify the optimal operating range of the cooling equipment and its efficiency based on the entering air temperature and the level of humidity. Remember that the air-conditioning equipment is not solely the cooling component but also the external heat rejection equipment. In smaller environments, both air conditioning and heat rejection may be contained within the same unit. In larger environments, this will not be the case, though it's possible that both components may be acquired from the same vendor. In these cases, the cooling and heat rejection capabilities will almost always be matched. However, if you are using multiple vendors, these components may not be equally matched. In this case, use the lowest-rated component for capacity calculations.

Once you have the theoretical maximum cooling capability of your data center cooling equipment, you can use the worksheet shown in Table 3.1 to calculate the heat output of the equipment in the data center. Remember that our goal is to have the cooling capacity and the heat output match, which results in the most efficient operation for the data center. However, there are a number of factors that prevent the cooling equivalent from operating at maximum efficiency and achieving the maximum theoretical cooling capacity.

Item	Data required	Heat output calculation	Heat output subtotal
IT Equipment	Total IT load power in watts	Same as total IT load power in watts	_____ watts
UPS with Battery	Power system rated power in watts	$(0.04 \times \text{Power system rating}) + (0.06 \times \text{Total IT load power})$	_____ watts
Power Distribution	Power system rated power in watts	$(0.02 \times \text{Power system rating}) + (0.02 \times \text{Total IT load power})$	_____ watts
Lighting	Floor area in square feet, or Floor area in square meters	$2.0 \times \text{floor area (sq ft)}$ , or $21.53 \times \text{floor area (sq m)}$	_____ watts
People	Max # of personnel in data center	$100 \times \text{Max \# of personnel}$	_____ watts
Total	Subtotals from above	Sum of heat output subtotals	_____ watts

**Table 3.1: Estimated heat output calculation worksheet.**

### Checking Hardware


CRAC units generally have four modes of operation: cooling, heating, humidification, and dehumidification. We are looking for these units to operate in a coordinated mode. Although multiple conditions may exist simultaneously, for example cooling and dehumidification, it is important that all systems within the defined area be operating in the same mode. This defined area may be a rack, a series of racks, a row, or even a series of rows or a contained area within the data center.

This allows us to prevent the condition known as “demand fighting.” When demand fighting is occurring, for example, one unit is humidifying while the unit next to it is dehumidifying; this situation not only increases the operating expenses by increasing energy use but also reduces the cooling efficiency and capacity of your data center.

Studies have shown that failing to address the demand fighting problem can result in a 20 to 30% reduction in efficiency. The simplest way to avoid this condition is to ensure that the set points for temperature and humidity are consistent on all the CRAC units within the data center. These settings can be more narrowly focused based on the techniques being used to cool the data center, such as row and rack cooling. But in all cases, units within the same grouping should be set to match.

To properly check hardware and to test the performance of the cooling system, it is necessary to measure both the return and supply temperatures. As Figure 3.2 shows, there are three points at which temperature should be monitored.





**Figure 3.2: Monitoring points.**

The temperature at the monitoring points starts at the supply air temperature. The temperature required at the server air inlet is the target temperature for the supply air temperature. The most common problem in delivering the appropriate temperature is short cycling.

Short cycling is when the cool air supply leaves the CRAC and rather than flowing through the IT equipment at which it is targeted, it instead bypasses that equipment and flows directly into the air return duct for the cooling unit. The good news is there are inexpensive solutions to addressing air flow performance.

Another very common problem that is easily dealt with is the issue of dirty filters in the air conditioning equipment. Simply making sure that clean air filters are in place goes a long way towards ensuring efficient operation of the cooling equipment.

### Systems Testing

Testing cooling equipment for proper operation and to ensure that it's operating within its optimal design parameters requires specific knowledge of cooling equipment. This level of maintenance knowledge is unlikely to be present within an IT department. It is, however, available through a facilities department, a maintenance company, or HVAC contractor. There are a number of features that should be checked by the appropriate technician, including:

- Current status of the equipment
- Chilled water cooling circuit
- Condenser water circuit
- Air cooled refrigerant piping

Although some aspects of the air conditioning equipment require specialized skills to test, the basic task of measuring temperature within the data center at specific locations, such as the aisles between the equipment racks, requires little more than temperature measuring equipment and a solid plan. Standards exist for the positioning of temperature sensors to allow for optimum testing. These standards are propagated by the American Society of Heating, Refrigeration, and Air-conditioning Engineers, in Standard TC9..

#### Resource

More details on the ASHRAE standards and their implementations can be found had to be [www.ashrae.org](http://www.ashrae.org).

In general, recommended inlet temperature will be in the 64.4 to 80.6°F range (also part of the TC9.9 standard). Temperatures that exceed this range can result in a reduction in system performance, reduced lifespan, and unexpected shutdowns of the equipment. For general temperature measurement, the temperature should be measured at a point 5 feet off the floor, and the measurements should be carried out over a 48-hour period recording maximum and minimum temperature levels within that timeframe.

Temperature should also be measured at the top, middle, and bottom inlet points of each rack, as very often air circulation is not equivalent at each point and equipment at the bottom of the rack is generating heat that is rising and can potentially cook equipment mounted above it.

## Understanding the Impact of Temperature on Servers

The long-held belief on data center server temperature has been that energy demands will decrease as server inlet temperatures increase. The trick, however, is finding the temperature that allows optimal server performance with minimal energy consumption. As erring to the wrong side (too warm) may cause hardware issues, it is important to properly balance server performance, energy demands, and cooling.

### Finding the Sweet Spot for Minimizing Energy Consumption


The key to utilizing less energy is to look for techniques that will enable cool air to be delivered to the server air inlet with as little reduction in the temperature of the provided air as possible. Normally, the CRAC will deliver air at 55 degrees in order for the air to reach the server at 70 degrees. The temperature increase comes from air mixing between the CRAC and the server. Thus, making changes to the data center environment that reduce air mixing, implementing containment systems, and managing airflow will allow the source air to be delivered at a higher temperature and still be delivered to the server at the optimal temperature, reducing the energy costs associated with cooling the data center.

## How Managing Airflow Reduces Costs

Unless the cooling capabilities are seriously undersized, there are a number of inexpensive technologies and tricks that can be used to aid in airflow management. By properly managing the airflow within the data center, there are significant potential costs savings that can be achieved by allowing equipment to operate in its optimal temperature range.

### Understanding the Use of Blanking Panels

Rack mounting is the most common way that servers are deployed within the data center. Blanking panels are the pieces that are inserted into empty spaces in those racks. These inexpensive pieces are a critical component in managing airflow in the data center (see Figure 3.3).



**Figure 3.3: Airflow with and without blanking panels.**

As the figure illustrates, without the blanking panel in place, warm air from the server rack is drawn back into the server, requiring additional cooling to keep the incoming air at a temperature that the server will be happy with. The simple expedient of installing the blanking panels prevents this from happening.

#### Minimal Cost / Maximum Gain

So if blanking panels make such a significant difference with so little effort, why aren't they always used? The answer is a simple one: In most cases, data centers are used to having bolt-on blanking panels on hand. This means that each blanking panel requires four screws and nuts to bolt it into place, which is a minor annoyance with a single panel, but a major effort with tens or hundreds of blank spaces to fill in the data center. Additionally, some blanking panels come in multiple sizes; the rack might need 4 U of space filled and the only panel available might be 3 U in size, making it a bit of a jigsaw puzzle to match up blanking panels with the spaces that need to be filled.

To address this issue, third-party vendors have started introducing standard-size snap-in modular blanking panels. Instead of needing to be screwed onto the rack, these modular panels snap-in and are quickly and easily installed. They also only come in a single size (1U) so that the only combination required is having sufficient 1U panels available to fill the open spaces. Modular blanking panels provide a simple and inexpensive solution to this problem.

Although blanking panels address the most common issues caused by improperly circulating air in your rack mounts, Table 3.2 highlights other common problems and offers simple and inexpensive solutions to make your server racks more effectively cooled.

Contributing Factor	Consequence	Control / Checklist
Unused vertical rack space	Open area permits hot exhaust air to return to equipment air intake resulting in overheating	Use airflow management blanking panels in all unused rack locations
Rack rails inset from side of enclosure	Open side area permits hot exhaust air to return to equipment air intake resulting in overheating	Do not use 23 in (584 mm) racks with rails set to 19 in (483 mm) Use racks which do not have open space between rail and side of enclosure
Monitors on shelves	Open space around monitor permits hot exhaust air to return to equipment air intake resulting in overheating	Use thin flip-top LCD monitors Obtain rack mount bezels for CRT monitors
Tower servers on shelves	Open space around servers permits hot exhaust air to return to equipment air intake resulting in overheating	Use rack mount servers Note: the power density of tower servers in the rack is very low which reduces the magnitude of this problem
Vertical rack space used to pass cables from front to back of rack	Open space around cables permits hot exhaust air to return to equipment air intake resulting in overheating	Use airflow management blanking panels equipped with a flexible brush or shield that allow the cables to pass through and reduce air leakage
Front or rear doors on the rack with restricted ventilation	Airflow resistance of the doors creates a pressure gradient which amplifies all of the above effects	Use fully perforated front and rear doors Do not use doors with glass, or doors with limited perforation
Space between racks	Open area permits hot exhaust air to return to equipment air intake resulting in overheating	Bay racks together wherever possible

**Table 3.2: Common cooling issues with rack-mounted servers.**

### Proper Placement of Vents and Tiles

In a traditional raised-floor data center, the placement of the floor tiles and vents is critical to maintaining proper cooling. In many older data centers, these tiles are not reconfigured when new or additional server racks are installed, resulting in a reduction in cooling efficiency due to the unplanned modification of the existing airflow paths. Properly reconfiguring the floor tiles and vents can address the changes made by the additional server racks and reduce the thermal load increase.

## Containment

As data center professionals have been driven to reduce their operational costs, the energy management practices of an earlier generation have been identified as particularly wasteful, especially with the drive to green IT. A number of technologies have been developed to deal with more efficient cooling practices. Chapter 2 talked about one of the most successful—row-based cooling. In this chapter, we are looking at techniques that let us maximize energy efficiency in our existing environment. For that purpose, we need to consider the issue of containment.

### Keeping the Cold Air, or Hot Air, Where You Need It

As we have previously discussed, one of the primary issues in data center cooling is the mixing of hot and cold air inappropriately, reducing the overall efficiency of the cooling process and requiring an increase in energy expenditure to provide the proper air temperature at the server inlet. This brings us to the concept of containment: keeping the hot and cold air separate, contained in an area that allows greater control over the environmental variables that most affect IT loads (temperature and humidity).

Containment falls into two categories:


- Cold air containment—Deployed in existing data centers, cold air containment systems (CACS) make use of the existing perimeter-based CRACs and a containment system that delivers cold air to the computer rows and uses the bulk of the room as the hot air return plenum.
- Hot air containment—A hot air containment system (HACS) makes use of in-row cooling and an enclosed aisle to build a self-contained system that can support high-density IT loads.

Although the systems are very different in operation and IT requirements as well as configuration and ease of implementation, there are a number of benefits that both types of containment systems have in common:

- Right-sized physical infrastructure—By matching the cooling capability to the cooling demand, data centers are able to avoid the problems associated with oversized solutions. Both containment systems allow IT to deliver the cooling capacity where it is needed.
- Better energy efficiency—As we discussed earlier, the ability to operate cooling equipment at higher temperatures, yet deliver the same level of cooling to the server hardware, reduces the overall energy expenditure and reduces costs.
- Reduced humidity control costs—In closed systems, such as CACS and HACS, little to no humidity changes are caused when the air is circulated through the system, reducing the need to expend energy to modify the data center humidity to maintain optimal environmental conditions.

### Evaluating Cooling Architectures

Organizations with traditional raised-floor computer rooms or those data centers simply using perimeter CRACs to maintain temperature and humidity can implement a CACS approach without a significant investment in new equipment.

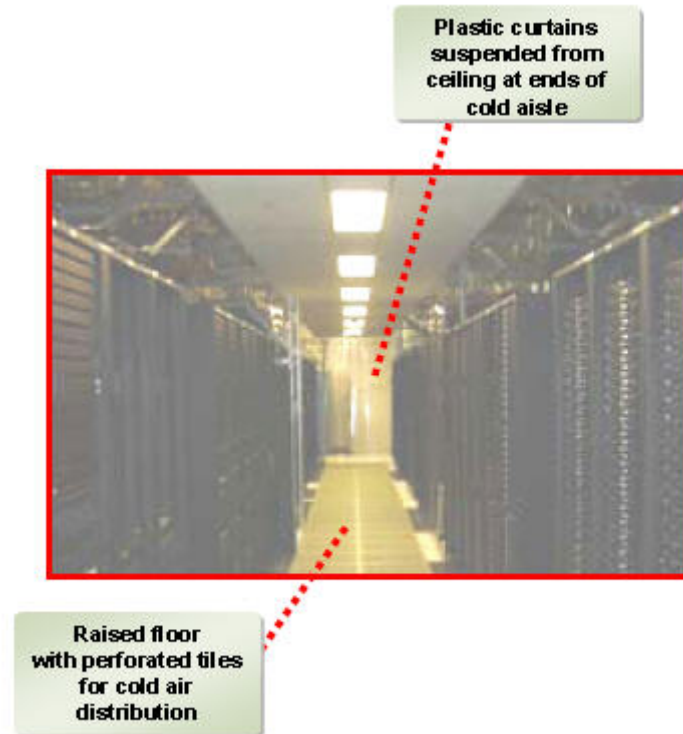


**Figure 3.4: Air flow diagram for a cold air containment system.**

As Figure 3.4 shows, the CACS delivery system doesn't look much different from a standard raised-floor data center. The difference is that the cooled air has been directed to a row (or rows) that have some form of containment system installed. This isn't a system that can be used for a small percentage of the data center. If this model is used, each row of rack servers must be contained within its own CACS as the overall temperature of the data center server room outside of the containment area will be well above the recommended 80.6-degree maximum inlet temperature for server operation.

Other equipment within the data center that is not contained within the rows of servers, such as tape libraries and other supporting devices, will need to have arrangements made to deliver cooling. Air temperatures within the data center outside of contained areas can be expected to go as high as 100 degrees. These temperatures become the norm and workers within the data center need to become adjusted to the changes.

Although there are formal systems for containing cold air, many data centers use the simple (and inexpensive) expedient of installing plastic curtains to surround the individual aisles that are being cooled. In a raised-floor system, the floor of the curtained aisle gets perforated panels to deliver the cooled air to the front of the server racks, while the back of the racks are open to the room, which is used as the hot air return plenum.



**Figure 3.5: An example of a simple CACS setup.**

Keep in mind that techniques we’ve already mentioned in this chapter concerning the proper use of blanking panels and floor tiles and vents are critical in the CACS model. The hot and cold air mixing needs to be controlled as much as possible in order to deliver cooled air with the least possible energy expenditure.

There are two significant limitations to the CACS approach to cooling:

- Density limitations—Without custom-designed systems and significant modifications to the data center infrastructure, there is a practical limit to how dense the servers can be racked within a CACS. Adding fan-powered floor tiles can increase potential density, as can moving to row-based cooling systems, which removes the density limitation, from the perspective of cooling, entirely.
- Inefficiencies in air distribution—The cooled air supply (that is, the CRAC) is at a distance from the CACS, and a significant portion of the overall energy requirement is needed to simply move the air to the location necessary. It is still more efficient, however, than the traditional “whole room” data center cooling model.

The HACS avoids many of the pitfalls of the CACS approach, but it requires a significant investment in the cooling infrastructure if traditional perimeter cooling is in place. The requirements include in-row cooling with variable speed fans, temperature controlled air supply, doors at either end of the row, and a roof.





**Figure 3.6: Example of a configured HACS.**

It is possible to build a HACS environment without in-row cooling, but it would require significant custom ducting and design in order to deliver and circulate the proper amount of air. In addition, this configuration would negate many of the benefits of the HACS design.

HACS is also effectively room-neutral in terms of its impact on the data center environment. Because of its self-contained nature, complete HACS configurations can be added to existing data centers without upsetting the environmental balance in the data center. This increases the design flexibility when adding additional HACS cooled rows to an existing data center. It becomes possible to not worry about environmental issues and just concentrate on finding an appropriate amount of space.

HACS provide a higher level of operational efficiency due to a number of factors, from the ability to operate at higher temperatures (without impacting the rest of the room) to the ability to operate at higher IT load density for lower energy expenditure than a CACS solution.

In an existing data center environment, implementing CACS can be done with minimal investment and modification of the current configuration. HACS offers the ability to build higher-density data centers and can be added to an existing data center, but at a much higher price point than implementing CACS. For a point-by-point comparison of the two models, see Table 3.3.

Characteristic	Cold Aisle Containment	Hot Aisle Containment	Comment
<b>Efficiency improvements</b>	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.
<b>Ability to increase cold air supply set point without impacting entire data center</b>	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.
<b>Leverages maximum number of potential free cooling days</b>	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.
<b>Room neutral solution</b>	No	Yes	A HACS deployment is a “drop in” solution. CACS impacts the surrounding data center infrastructure.
<b>Ease of deployment with room cooling</b>	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.
<b>Ability to scale for high density</b>	No	Yes	CACS is often implemented with raised floor and inefficient fan assisted floor tiles are needed in order to achieve higher density.
<b>Room neutral design</b>	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.
<b>Adverse temperature impact on non-racked equipment</b>	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.

**Table 3.3: Comparison of HACS and CACS benefits and concerns.**

## Conclusion

Although there are few truly free things that can be done to improve the energy efficiency of your data center, we've given you the starting point to evaluate how your data center is utilizing energy as well as basic tips to improve the overall efficiency of the existing facility. Don't forget that there are always small things that can be done to improve overall efficiency that don't require major equipment configuration changes, such as the use of airflow management devices and using economy modes on cooling equipment.

Understanding the energy needs and uses within your data center is the first step towards implementing a more energy-efficient facility. Start with a basic cooling audit and evaluate the current data center cooling infrastructure, then move on with simple, inexpensive changes, such as procedural changes regarding ancillary power usage; blanking panels and efficient tile and panel configurations; and maintenance of unimpeded cooling airflow. These techniques are cost-effective measures you can take to see noticeable improvements in cooling efficiency.

## Download Additional eBooks from Realtime Nexus!

Realtime Nexus—The Digital Library provides world-class expert resources that IT professionals depend on to learn about the newest technologies. If you found this eBook to be informative, we encourage you to download more of our industry-leading technology eBooks and video guides at Realtime Nexus. Please visit <http://nexus.realtimepublishers.com>.