

Avoidable Mistakes that Compromise Cooling Performance in Data Centers and Network Rooms

By Neil Rasmussen

White Paper #49



Executive Summary


Avoidable mistakes that are routinely made when installing cooling systems and racks in data centers or network rooms compromise availability and increase costs. These unintentional flaws create hot-spots, decrease fault tolerance, decrease efficiency, and reduce cooling capacity. Although facilities operators are often held accountable for cooling problems, many problems are actually caused by improper deployment of IT equipment outside of their control. This paper examines these typical mistakes, explains their principles, quantifies their impacts, and describes simple remedies.

Introduction

Most data centers and network rooms have a variety of basic design and configuration flaws that prevent them from achieving their potential cooling capacity and prevent them from delivering cool air where it is needed. These problems are generally unrecognized because computer rooms have typically been operated at power densities well below their design values. However, recent increases in the power density of new IT equipment are pushing data centers to their design limits and revealing that many data centers are incapable of providing effective cooling as expected.

In addition to the decrease in system availability that can result from under-performing cooling systems, significant costs are incurred. This paper will describe common design flaws that can cause the efficiency of the cooling system to decrease by 20% or more. Separate studies from Lawrence Berkeley National Laboratories and from APC Corp concluded that a typical data center exhibits electrical power consumption as shown in **Figure 1**, where the electrical power consumed by the cooling system is comparable to the power consumed by the entire IT load. A 20% loss of cooling efficiency translates to an increase in total electrical power consumption of 8%, which over the 10 year life of a 500kW data center translates to a cost of wasted electricity of approximately \$700,000. *This significant waste is avoidable for little or no cost.*

Figure 1 – Breakdown of electricity consumption of a typical data center



The sub-optimization of the data center cooling system arises from a variety of sources. These problem sources include the design and specification of the cooling plant itself, and they include how the complete system delivers the cool air to the load. This paper focuses on cooling problems related to the distribution of cooling air and setup issues related to the deployment of IT equipment, for the following reasons:

- Because there are practical, feasible, and proven solutions
- Many fixes can be implemented in existing data centers

- Large improvements can result from little or no investment
- Both IT people and facilities people can contribute to fixing them
- Solutions are independent of facility or geographic location
- They lend themselves to correction through policies that are simple to implement

The paper breaks down the common flaws into five contributing categories, and addresses each in turn:

- Airflow in the rack itself
- Layout of racks
- Distribution of loads
- Cooling settings
- Layout of air delivery and return vents

For each category, a number of issues are described along with a simple description of the theory of the problem and how it impacts availability and Total Cost of Ownership. This information is summarized in the tables.

Finally, a number of policies are described which, if implemented, can significantly improve data center availability and reduce TCO.

Basic Airflow Requirements

The airflow in and around the rack cabinet is critical to cooling performance. The key to understanding rack airflow is to recognize the fundamental principle that IT equipment cares about two things:

1. That appropriate conditioned air is presented at the equipment air intake
2. That the airflow in and out of the equipment is not restricted.

The two key problems that routinely occur and prevent the ideal situation are

1. The CRAC air becomes mixed with hot exhaust air before it gets to the equipment air intake
2. The equipment airflow is blocked by obstructions.

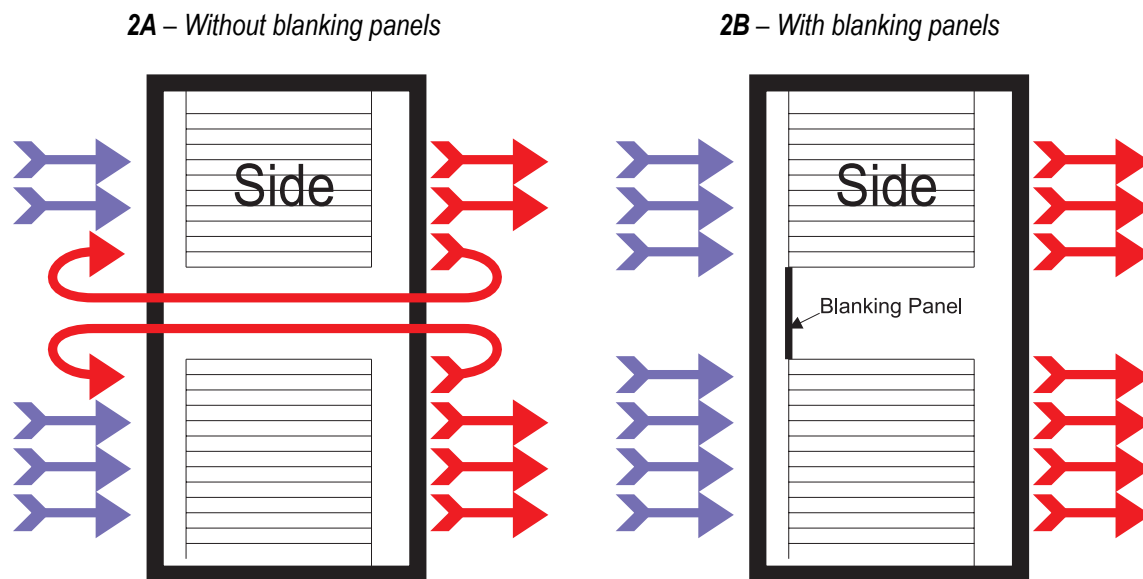
The common theme throughout the next sections is that well-intentioned implementation decisions that appear to be inconsequential actually create the two problems above, and that the common solutions routinely used to address the symptoms of these problems significantly compromise availability and increase costs.

Airflow in the Rack Cabinet

Although the rack is frequently thought of as a mechanical support, it provides a very critical function in preventing hot exhaust air from equipment from circulating back into the equipment air intake. The exhaust air is slightly pressurized, and this combined with the suction at the equipment intake leads to a situation where the exhaust air is induced to flow back into the equipment intake. **The magnitude of this effect is much greater than the magnitude of the effect of buoyancy of hot exhaust air, which many people believe should naturally cause the hot exhaust air to rise away from the equipment.** The rack and its blanking panels provide a natural barrier, which greatly increases the length of the air recirculation path and consequently reduces the equipment intake of hot exhaust air.

The very common practice of omitting blanking panels can be found to greater or lesser degrees in 90% of data centers, despite the fact that all major manufacturers of IT equipment specifically advise that blanking panels be used. The resulting recirculation problem can lead to a 15°F or 8°C rise in temperature of IT equipment. A detailed description of this effect along with experimental data is found in APC White Paper #44, "Improving Rack Cooling Performance Using Blanking Panels". Blanking panels modify rack airflow as shown in **Figure 2**. Installing blanking panels is a simple process that can be implemented in almost any data center at very low cost.

Figure 2 – Air recirculation through a missing blanking panel



Many configured racks exhibit other flaws that have the same effect as omitting blanking panels. Using wide racks with inset rails allows recirculation around the sides of the rack rails. Using shelves to mount IT equipment prevents the use of blanking panels and provides wide-open paths for exhaust air recirculation. Some standard 19" racks have inherent air recirculation paths built-in around the rails and at the top and bottom of the enclosure. In these cases the installation of blanking panels cannot completely control

recirculation. Many racks have simply not been designed to work effectively in a high density IT environment. Standardizing on the right rack and using blanking panels can greatly reduce recirculation and reduce hot-spots.

The benefits of reducing hot-spot temperature by using blanking panels and selecting racks that control recirculation are clear and offer obvious benefits in terms of system availability. However, there are other less obvious but very significant benefits that require explanation.

Recirculation impacts fault tolerance

Rack systems with significant recirculation have reduced fault tolerance and reduced maintainability when contrasted with properly implemented systems. In most installations, cooling is provided by an array of CRAC units feeding a common air supply plenum. In such an arrangement, it is often possible to maintain facility cooling with one CRAC system off-line due to failure or maintenance; the remaining CRAC units are able to pick up the required load. Recirculation compromises this fault tolerance capability in the following ways:

- Lower CRAC return air temperature due to recirculation causes the remaining CRAC units to operate at reduced capacity and the system is unable to meet the cooling capacity requirement
- Higher air feed velocities needed to overcome recirculation effects cannot be maintained by the remaining systems causing increased recirculation and overheating at the loads.

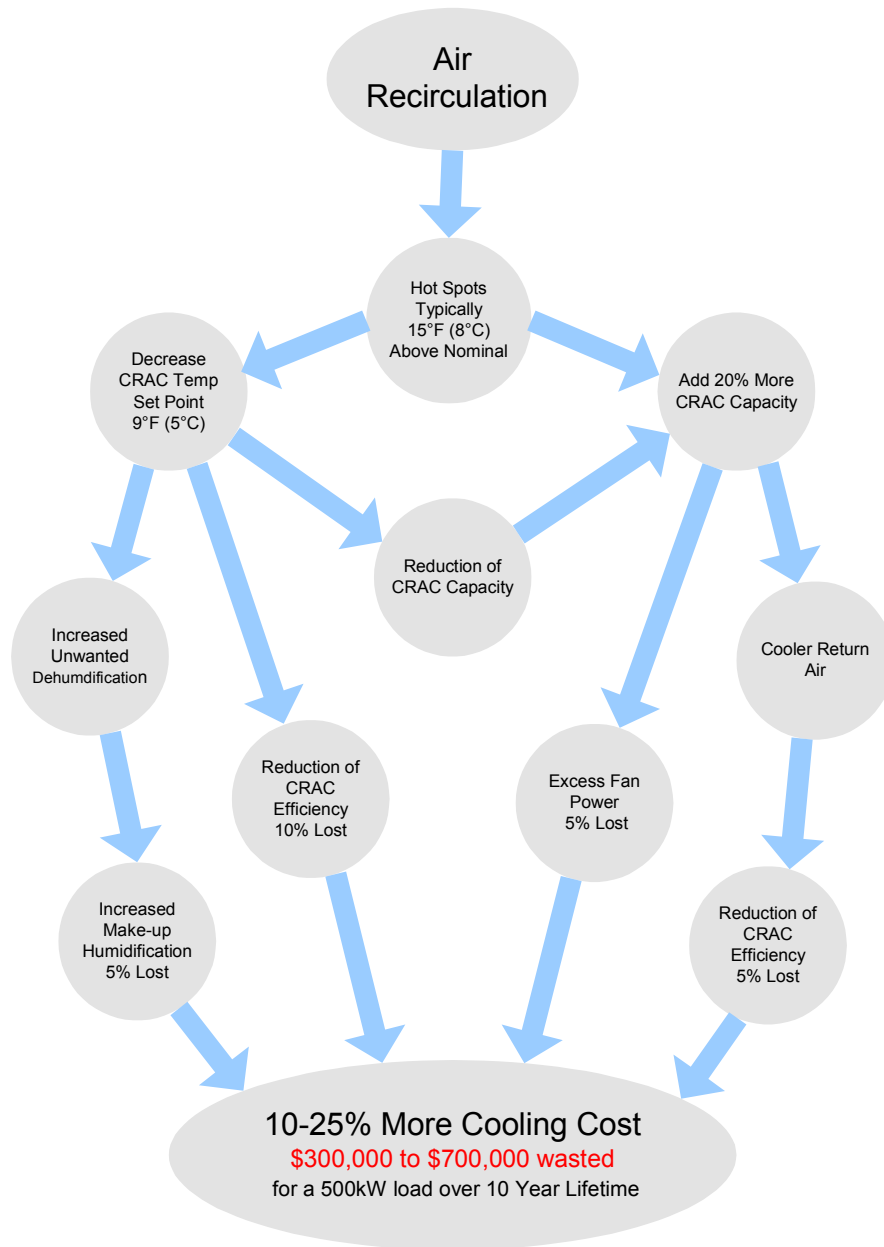
Recirculation impacts total cost of ownership

The availability issues of overheating and fault tolerance make a compelling case for the use of standardized racks and blanking panels. However, the TCO consequences of recirculation are dramatic and make the case overwhelming.

The largest life cycle cost related to cooling is the cost of electricity to operate cooling equipment and fans. The amount of cooling Wattage or Tonnage required by a data center is not affected by recirculation; however the efficiency of the cooling systems is significantly and adversely affected. This means that recirculation will increase the costs related to electricity. Furthermore, the costs compound as shown in **Figure 3**.

Figure 3 illustrates the sequence of consequences that typically occur from attempts to deal with the primary symptom of recirculation, namely, hot-spots. The two most common responses to hot-spots are to reduce the CRAC supply temperature, or to increase CRAC capacity, or a combination of both. These responses have significant unforeseen costs as described in the figure. Controlling recirculation by design and policy as described in this paper can be done for very little cost and avoids the consequences shown in the figure.

Figure 3 – Cascade of financial consequences of recirculation



Restriction of airflow starves equipment of fresh air, resulting in overheating. Furthermore, air restriction at the front or rear of the rack encourages recirculation through rack space without blanking panels. Therefore, it is critical to use racks that have very high door ventilation, and racks that have enough room in the back of the rack to prevent cable bundles from obstructing airflow. Users sometimes choose shallow racks believing that this will increase the floor space utilization but then are unable to utilize the density because of thermal limits due to cable-related airflow obstruction.

Table 1 – Summary of rack airflow design flaws with consequences

Design Flaw	Availability Consequences	TCO Consequences	Solution
No blanking panels Equipment on shelves Use of 23 inch (584mm) racks without rail-brushes	Hot spots, particularly at the tops of racks Loss of cooling redundancy	Electricity costs Reduced capacity of CRAC Humidifier maintenance Water consumption	Use blanking panels Do not use shelves Use racks that have no open space outside the rails Add brushes outside rails on wide racks
Under-rack wire openings without brushes	Hot spots, particularly at the tops of racks Loss of static pressure in raised floor Loss of cooling redundancy	Reduced efficiency of CRAC	Use brushes or gasketing on under-rack wire openings
Glass doors Doors with low ventilation	Overheating Amplification of problems relating to blanking panels	Decreases space and rack utilization	Use fully vented doors front and rear
Use of Fan trays and roof fans	Very little benefit Same investment could have been used for useful purpose	Wasted capital Wasted electricity	Do not use fan trays or roof fans
Shallow racks	Cable obstructions cause overheating	Decreases space and rack utilization	Use racks with enough depth to allow free air around cables

In addition to the passive means of controlling rack airflow described above, the use of rack based fan systems can be used to control rack air distribution. Some rack fan systems, like fan trays and roof fans, offer little benefit. Other fan systems, like systems that distribute under-floor air to the front of the rack, or scavenge exhaust air from the rear of the rack, can significantly improve rack airflow, reduce circulation effects, and increase rack power handling capability. Detailed discussion of these systems can be found in APC White Paper #46, “Power and Cooling for Ultra-High Density Racks and Blade Servers”. Standardizing on a rack that is designed for retrofit of effective supplemental air fan units provides for future high-density capability.

Layout of Racks

Proper rack airflow control as described in the previous section is essential to effective cooling, but is not sufficient by itself. Proper layout of racks in a room is a critical part of ensuring that air of the appropriate temperature and quantity is available at the rack. Flow of air to the rack is key.

The objective of correct rack layout is again to control recirculation, that is, to prevent CRAC air from becoming mixed with hot exhaust air before it reaches the equipment air intake. The design principle is the same: to separate the hot exhaust air from the equipment intake air to the maximum extent possible.

The solution to this problem is well known. By placing racks in rows and to reversing the direction that alternate rows of racks face, recirculation can be dramatically reduced. The principles of this solution are described by the Uptime Institute in their white paper “Alternating Cold and Hot Aisles Provides More Reliable Cooling for Server Farms”.

Despite the clear advantages of the hot-aisle-cold-aisle system, surveys show that approximately 25% of data centers and network rooms put racks in rows that face in the same direction. **Putting racks in the same direction causes significant recirculation, virtually assures that there will be hot-spot problems, and guarantees that the cost of operating the system will be increased significantly.** The costs will vary by installation and are again illustrated by the previous **Figure 3**.

The effective application of the hot-aisle-cold-aisle technique consists of more than simply putting racks into alternating rows. Of the 75% of installations that do use the hot-aisle-cold-aisle technique, over 30% do not correctly arrange the air distribution and return systems to properly supply the rows. This is discussed later in the section titled “Layout of Air Delivery and Return Vents”

Of the sites that face racks in the same direction and do not use hot-aisle-cold-aisle techniques, surveys conducted by APC indicate that the majority are due to a management directive based on the cosmetic appearance of the data center. The surveys further suggest that these unfortunate directives would never have been made if the crippling consequences had been made clear.

For systems laid out with racks all facing the same way, many of the techniques described in this paper will be much less effective. If alternating racks is not possible, then one effective way to deal with hot-spots in this environment is to apply a supplemental air distribution unit to the affected racks.

Table 2 – Summary of rack layout design flaws with consequences

Design Flaw	Availability Consequences	TCO Consequences	Solution
Racks all facing in the same direction Hot-aisle-cold-aisle not implemented	Hot spots Loss of cooling redundancy Loss of cooling capacity Humidifier failures	Excess power consumption Water consumption Humidifier maintenance	Use hot-aisle-cold-aisle layout
Not in rows	Same problems	Same	Arrange racks in rows
In rows but not tightly butted	Same problems	Same	Bay racks together Do not space racks out

Distribution of Loads

The location of loads, particularly high power loads, can stress the capabilities of a data center. Pockets of high density loads typically occur when high-density, high-performance servers are packed into one or more racks. This situation can give rise to hot-spots in the data center and cause the operators to take corrective action such as decreasing the air temperature set point, or adding CRAC units. These actions give rise to the negative consequences summarized in **Figure 3**.

For these reasons, there is a significant advantage to spreading the load out where feasible. Fortunately, fiber and Ethernet connections are not adversely affected by spreading equipment out. Typically, the desire to co-locate such devices is driven by IT people who believe it is more convenient to co-locate devices. People attempting to co-locate high power loads should be advised of the availability advantages and cost savings which accrue from spreading out the load.

There are other options for high power racks that can avoid adverse cooling impacts. For a more complete discussion of the subject of dealing with high power racks, consult APC White Paper #46, "Power and Cooling for Ultra-High Density Racks and Blade Servers".

Table 3 – Summary of load distribution design flaws with consequences

Design Flaw	Availability Consequences	TCO Consequences	Solution
Concentrated Loads	Hot spots Loss of cooling redundancy	Excess power consumption	Spread loads out evenly as possible

Cooling Settings

The previous discussions described the adverse consequences of reducing the CRAC air temperature setting. Air conditioning performance is maximized when the CRAC output air temperature is highest. Ideally, if there were zero recirculation, the CRAC output temperature would be the same 68-77°F (20-25°C) desired for the computer equipment. This situation is not realized in practice and the CRAC output air temperature is typically somewhat lower than the computer air intake temperature. However, when the air distribution practices described in this paper are followed, it allows the CRAC temperature set-point to be maximized. To maximize capacity and optimize performance, the CRAC set point should not be set lower than that required to maintain the desired equipment intake temperatures.

Although the CRAC temperature set point is dictated by the design of the air distribution system, the humidity may be set to any preferred value. Setting humidity higher than that required has significant disadvantages. First, the CRAC unit will exhibit significant coil condensation and dehumidify the air. The dehumidification function detracts from the air-cooling capacity of the CRAC unit significantly. To make matters worse, humidifiers must replace the water removed from the air. This can waste thousands of

gallons of water per year in a typical data center, and humidifiers are a significant source of heat, which must be cooled and consequently further detracts from the capacity of the CRAC unit. This situation is compounded when significant recirculation exists because the lower temperature CRAC air condenses more readily. Therefore it is essential not to operate a data center at higher humidity than necessary.

Some data centers, including most early data centers, had high velocity paper or forms printers. These printers can generate significant static charge. To control static discharge, the result was the development of a standard for around 50% relative humidity in data centers. However, for data centers without large high-speed forms printers a humidity of 35% relative humidity will control static charge. Operating a data center at 35% relative humidity instead of 45% or 50% can save significant amounts of water and energy, especially if there is significant recirculation.

An additional problem can occur in data centers with multiple CRAC units equipped with humidifiers. It is extremely common in such cases for two CRAC units to be wastefully fighting each other to control humidity. This can occur if the return air to the two CRAC units is at slightly different temperatures, or if the calibrations of the two humidity sensors disagree, or if the CRAC units are set to different humidity settings. One CRAC unit will be dehumidifying the air while another is humidifying the air. This mode of operation is extremely wasteful, yet is not readily apparent to the data center operators.

The problem of wasteful CRAC humidity fighting can be corrected by either A) central humidity control, B) coordinated humidity control among the CRAC units, C) turning off one or more humidifiers in the CRACS, or D) by using deadband settings. Each of these techniques has advantages, which will not be discussed in detail in this paper. When the problem occurs, the most feasible way to correct it in typical systems with independent CRACs is by verifying that systems are set to the same settings, are properly calibrated, and then expanding the deadband humidity setting, which is available on most CRAC units. When the deadband setting is set to +/-5% the problem will usually be corrected.

Table 4 – Summary of cooling settings design flaws with consequences

Design Flaw	Availability Consequences	TCO Consequences	Solution
Humidity set too high	Hot spots Loss of cooling redundancy	Excess power consumption Water consumption Humidifier maintenance	Set humidity at 35-50%
Multiple CRAC units fighting to control humidity of the same space	Loss of cooling redundancy Loss of cooling capacity	Excess power consumption Water consumption Humidifier maintenance	Set all units to the same setting Set 5% dead-band on humidity set points Use centralized humidifiers Turn off unnecessary humidifiers

Layout of Air Delivery and Return Vents

Rack airflow and rack layout are key elements to direct air to maximize cooling performance. However, one final ingredient is required to ensure peak performance, which is the layout of air delivery and return vents. Improper location of these vents can cause CRAC air to mix with hot exhaust air before reaching the load equipment, giving rise to the cascade of performance problems and costs described previously. Poorly located delivery or return vents are very common and can erase almost all of the benefit of a hot-aisle-cold-aisle design.

The key to air delivery vents is to place them as close to the equipment air intakes as possible and keep the cool air in the cold aisles. For under-floor air distribution, this means keeping the vented tiles in the cold aisles only. Overhead distribution can be just as effective as a raised floor distribution system, but again the key is that the distribution vents be located over the cold aisles, and the vents be designed to direct the air directly downward into the cold aisle (not laterally using a diffusing vent). In either overhead or under floor systems any vents located where equipment is not operational should be closed since these sources end up returning air to the CRAC unit at low temperature, increasing dehumidification and decreasing CRAC performance.

The key to air return vents is to place them as close to the equipment exhausts as possible and collect the hot air from the hot aisles. In some cases, an overhead dropped ceiling plenum is used and the return vents can be easily aligned with the hot aisles. When a high, open, bulk return ceiling is used, the best approach is to locate the returns of the CRAC unit as high up in the ceiling as possible and, where possible, spread out the return using ductwork in an attempt to align returns with the hot aisles. Even a crude return plenum with only a few return vents crudely aligned with hot aisles is preferred over a single bulk return at the side of the room.

For smaller rooms without raised floor or ductwork, upflow or downflow CRAC units are often located in a corner or along a wall. In these cases, it can be difficult to align cool air delivery with cold aisles and hot air return with hot aisles. Performance will be compromised in these situations. However, it is possible to improve the performance of these systems as follows:

Sealing Cable Cutouts

Cable cutouts in a raised floor environment cause significant unwanted air leakage and should be sealed. This lost air, known as bypass airflow, contributes to IT equipment hotspots, cooling inefficiencies, and increases infrastructure costs.

Many sites ignore unsealed floor openings and believe that inadequate cooling capacity is the problem. As a result, additional cooling units are purchased to address the overheating. In fact, those supplementary units may not be needed.

One alternative to minimize the cost of additional cooling capacity is to seal cable cutouts. The installation of raised floor grommets both seals the air leaks and increases static pressure under a raised floor. This improves cool air delivery through the perforated floor tiles.



- For upflow units, locate the unit near the end of a hot aisle and add ducts to bring cool air to points over cold aisles as far away from the CRAC unit as possible.
- For downflow units, locate the unit at the end of a cold aisle oriented to blow air down the cold aisle, and add either a dropped ceiling plenum return, or hanging ductwork returns with return vents located over the hot aisles.

A study of poorly placed return grilles reveals a major underlying root cause: personnel feel that some aisles are hot and some are cold and assume this is an undesirable condition and attempt to remedy it by moving cool air vents to hot aisles, and moving hot air returns to cold aisles. **The very condition that a well-designed data center attempts to achieve, the separation of hot and cool air, is assumed by personnel to be a defect and they take action to mix the air, compromising the performance and increasing the costs of the system. People do not understand that hot aisles are supposed to be hot.**

Obviously the arrangement of the distribution and return vents is easiest to establish at the time when the data center is constructed. Therefore it is essential to have a room layout with row locations and orientation before the ventilation system is designed.

Table 5 – Summary of air delivery and return design flaws with consequences

Design Flaw	Availability Consequences	TCO Consequences	Solution
Hot air return location not over hot aisle Dropped-ceiling lamp with integral air return located over cold aisle	Hot spots, particularly at the tops of racks Loss of cooling redundancy	Electricity costs Reduced capacity of CRAC Humidifier maintenance Water consumption	Locate hot air returns over hot aisle Do not use lamps with integral air returns over cold aisles, or block the return
Overhead delivery vents over hot aisles Vented floor tile in hot aisle	Hot spots Loss of cooling redundancy	Electricity costs Reduced capacity of CRAC Humidifier maintenance Water consumption	For overhead delivery always locate delivery vents over cold aisles For raised floor delivery always locate delivery vents in cold aisles
Vented floor tile near no load Overhead delivery vent open above no load Peripheral holes in raised floor for conduits, wires, pipes	Small	Electricity costs Reduced capacity of CRAC	Close vents or openings located where there is no loads
Low height of return vent in high ceiling area	Loss of CRAC capacity Loss of cooling redundancy	Electricity costs Reduced capacity of CRAC Humidifier maintenance Water consumption	Use dropped ceiling for return plenum, or extend duct to collect return air at high point

Prevention via Policies

By following the guidelines of this paper it is possible to make new data centers that are significantly more available, with fewer hot spots, and less costly to operate. Some of the techniques described can be implemented in existing data centers, but others are impractical on live systems. Naturally, it is best to avoid the problems in the first place. Surveys by APC suggest that most of the defects in cooling system design are unintentional and would not have been made if the facilities or IT staff had understood the importance of proper air distribution on the performance, availability, and cost of the data center. One way to effectively communicate the key factors to the parties involved is through the use of policies.

Table 6 – Suggested data center design policies

Policy	Justification
Use hot-aisle-cold-aisle rack layout	The separation of hot and cold air reduces hot spots, increases fault tolerance, and significantly reduces the consumption of electricity. It is a well known fact that facing all rows in the same direction causes each row to be fed hot exhaust air from the row in front of it, leading to overheating and dramatically reduced air conditioner performance.
Use blanking panels in unused positions in all racks	Blanking panels prevent hot exhaust air from equipment from returning to the equipment intake, preventing hot spots and increasing equipment life. All server and storage manufacturers specify that blanking panels should be used.
Use gaskets or brushes on all under-rack wire openings in raised floors	The purpose of the raised floor air distribution system is to deliver cool air to the equipment intakes. These intakes are located on the front of the racks. Openings below the racks feed cool air to the equipment exhaust, bypassing the equipment and reducing the performance of the cooling system.
Do not attempt to correct the temperature in hot aisles. They are supposed to be hot.	The purpose of the hot aisle is to separate the hot exhaust air from the cool equipment intake air. Any attempt to defeat this function will compromise the design of the system, reduce equipment reliability, and increase operating cost. The exhaust air from the equipment is supposed to be hot, and the hot aisle is intended to take this hot air back to the air conditioning system. Having the hot aisle be hot helps ensure that the equipment intakes on the cold aisle are kept cold.
Standardize racks	Racks serve an essential function as part of the cooling system and are not simply mechanical supports. Rack features that prevent exhaust air from reaching equipment intakes, provide for proper ventilation, provide space for cabling without airflow obstruction, and allow the retrofit of high density supplemental cooling equipment should be part of a rack standard.
Spread out high density loads	Concentrating high power loads in one location will compromise the operation of those loads and typically increase data center operating costs. Fault tolerance in the air delivery system is typically compromised when high power loads are concentrated. The entire data center temperature and humidity controls may need to be altered in a way that compromises cooling capacity and increases cooling cost.

Establishing policies can force constructive discussions to occur. In addition to establishing policies, communication can be facilitated by signage or labeling. An example of a label that is located on the rear of racks in hot aisles is shown in **Figure 4**. Personnel such as IT personnel will often view the hot aisle as an

undesirable problem or defect. This label helps them understand why one area of the data center is hotter than another.

Figure 4 – Label communicating the purpose of the hot aisle



Conclusions

The air distribution system is a part of the data center that is not well understood, and facility operators and IT personnel often take actions involving airflow that have unintentional and adverse consequences to both availability and cost.

Flawed airflow implementation has not been a serious problem in the past, due to low power density in the data center. However, recent increases in power density are beginning to test the capacity of cooling systems and give rise to hot-spots and unexpected limitations of cooling capacity

Decisions such as facing all racks in the same direction are often made for cosmetic reasons to project image; but as users and customers become more educated they will conclude that people who do not implement airflow correctly are inexperienced, which is the opposite of the original intent.

Adopting a number of simple policies and providing a simple justification for them can achieve alignment between IT and Facilities staff resulting in maximum availability and optimized TCO.

About the Author:

Neil Rasmussen is the Senior VP of Innovation for APC- the Critical Power and Cooling Services Business Unit of Schneider Electric. He establishes the technology direction for the world's largest R&D budget devoted to power, cooling, and rack infrastructure for critical networks. Neil is currently working to advance the science of high-efficiency, high-density, scalable data center infrastructure solutions and is the principal architect of the APC InfraStruXure system.

Prior to founding APC in 1981, Neil received his Bachelors and Masters degrees from MIT in electrical engineering where he did his thesis on the analysis of a 200MW power supply for a tokamak fusion reactor. From 1979 to 1981, he worked at MIT Lincoln Laboratories on flywheel energy storage systems and solar electric power systems.