

Overview and Design of Data Center Cabinets

In today's technology-dependent world, a reliably operating data center is a core necessity for small and large businesses alike. Without properly running network systems, most companies simply cannot run at all. In fact, almost all aspects of business revolve around the need to process, access and communicate information—the key objective of networking systems.

CONTENTS

Introduction	2
1. Power and Cooling	3
-Sustained vs. Surge Power	3
-Nameplate vs. Actual Power Usage	3
2. Understanding ΔT and CFM	4
- ΔT	4
-CFM	4
3. Raised vs. Non-Raised Floors	5
4. Energy Costs and Savings	7
5. Enclosure Overview	8
-Technology Improvements	8
-LEED	8
-Cabinet Characteristics	9
-Airflow Accessories	10
6. Cable Management	11
7. Cabinet Configurations	12
-Random Data Centers	13
-Passive vs. Active Cooling	13
-Hot Aisle/Cold Aisle Configuration	14
-HACA	15
-Floor-Ducted Solution	15
-Direct-Ducted Chimney Systems	16
-Liquid Cooling	18
-Hot Aisle/Cold Aisle Containment	20
8. Additional Cooling Solutions	22
-Heat Exchanger Doors	22
-In-Row Cooling	22
-Ceiling-Mounted Cooling Units	22
-Remote Applications	22
-Network Switches	23
9. Summary and Selection Guide	24



INTRODUCTION

DATA CENTER EQUIPMENT

To maintain maximum network availability, all equipment must be kept below a specified temperature range—a requirement that has recently become more challenging to meet. As the need for information has increased, so has the processing power of the network equipment, resulting in increasingly higher processor densities and heat levels. Elevated temperatures result in equipment failure, and ultimately, costly network downtime, making the need for efficient cooling systems more important than ever. Since nearly all power consumed by the processors is converted to heat—which must be removed via air (convection)—the data center in which this equipment resides must efficiently provide cold air to network equipment intakes and recycle the hot exhaust to remove the heat and keep vital networking equipment operating.

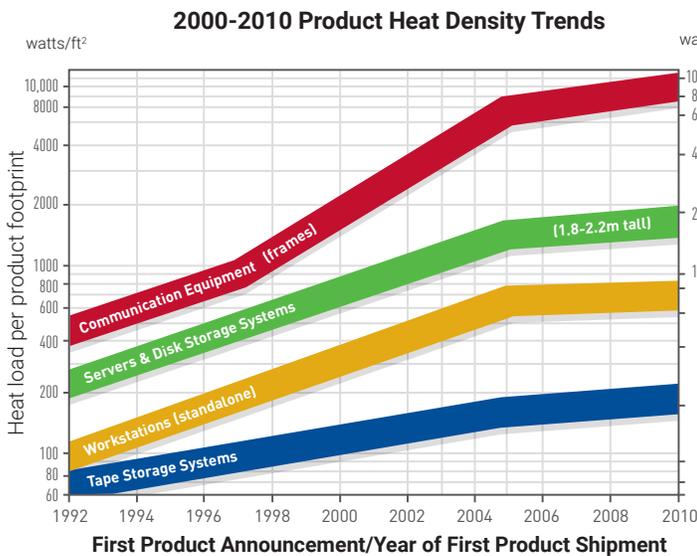
Data center equipment has followed Moore’s Law in many respects. As the transistor density increases or the amount of processing increases, the amounts of power consumed and heat produced also increase. Electrical and electronic devices have become smaller but higher performing, thus consuming more power and generating more heat in a given space. The relationship between power and heat is direct. The power entering network equipment equals the amount of heat produced: watts (power) in equals watts (heat) out. All power that comes into network equipment results in heat, with the exception of a mere fraction: 1 percent of the total power is consumed by outgoing digital signals.

To prevent equipment failure, its operating temperature must be less than 95° to 100° F, with the intake temperature below 80° to 85° F. When the intake temperatures rise above this range, network equipment

SUSTAINED TEMPERATURES ABOVE 93°F REDUCE ELECTRONIC EQUIPMENT FUNCTIONAL LIFE

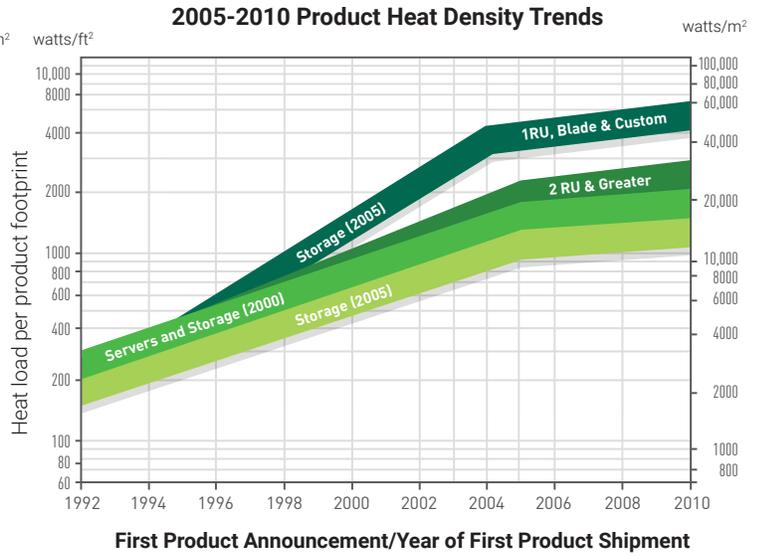
overheats. Internal thermal protection is activated, lowering the temperature by reducing processing or completely shutting down to avoid critical damage. Less intense, long-term elevated temperatures are also costly, leading to reduced reliability and product life.

To maintain reliable network systems that provide information access, processing and communication, power, cooling and airflow must be managed.



The trends on this chart show annual rates of increasing product footprint heat density ranging from 7% to 28%. Data center managers should be cautious, since the average density spread across the entire room will be reduced by the space required for aisles and site infrastructure equipment.

© 2000-2006 The Uptime Institute, Inc. Version 1.2



■ 2005 Projection for Servers — 1 RU, Blade and Custom
 ■ 2005 Projection for Servers — 2 RU and Greater
 ■ 2005 Projection for Storage
 ■ 2000 Projection for Servers and Storage

© 2006 The Uptime Institute, Inc. Version 2.0 2006-2010 Trend Prediction Data Source: ASHRAE, DataCom Equipment Power Trends and Cooling Applications, 2006. American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc., www.ashrae.org.



SUSTAINED VERSUS SURGE POWER

To maintain sufficiently cool temperatures, data center designers should plan their cooling systems according to the power used and heat generated—which are determined by the volume of network activity and processing the equipment performs, and varies throughout a normal day. The sustained power is the amount of power required to perform minimal, constant activities. During points of high network activity, more power is required, and in turn, more heat is generated.

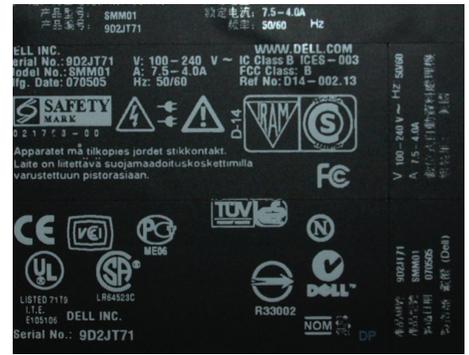
When these peak or surge power usages occur for just a few minutes in an hour, the heat load can usually be easily and quickly handled. However, the longer this activity is maintained, more power is required and more heat accumulates. Additionally, many data centers balance their power and cooling requirements by spreading out various network activities and processing runs to decrease surge times. Instead of designing multiple functions to perform simultaneously, some maintenance functions are scheduled for periods of lower power usage.

While excessive heat can be detrimental to electrical equipment, power and cooling systems based on peak load requirements will provide more capabilities than needed—and a much higher price tag (capital) than is necessary. Instead, since periods of surge power exceed the sustained power usage, data centers should be designed for an average volume of power usage.

NAMEPLATE VERSUS ACTUAL POWER USAGE

As planning power and cooling equipment in a data center simply around surge power usage results in excessive capacities and cost, so does planning a system according to nameplate requirements. Equipment contains a nameplate with a label listing the amount of power (watts) that it consumes under 100 percent utilization. However, this quantity is the maximum possible power supply that could be used—not the typical running power. In fact, this amount is rarely, if ever, achieved, as most equipment designers provide more power supply capacity than the equipment can possibly consume as a precautionary measure. In reality, network equipment running at a high percentage of its nameplate value is likely being overburdened, resulting in slow response time and is in need of replacement or reconfiguration.

Since the nameplate value is greater than the actual equipment power usage, adding the values of all equipment contained within a cabinet to determine that cabinet's power and cooling requirements results in an inaccurate and extremely high number. As an alternative, some equipment manufacturers provide guidance in typical power and cooling needs, and the American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) also provides a listing of typical running requirements. Actual power usage can also be determined with a PDU that has an Amp meter, which provides instantaneous power consumption when multiplied by the incoming voltage. Recorded over time, this provides accurate overall power consumption and is invaluable to understanding total power usage in a data center. Finally, a common industry practice is to de-rate the nameplate by 40 percent to garner an approximate actual equipment operating power.



TO PROVIDE THE MOST EFFICIENT AND COST-EFFECTIVE POWER AND COOLING SOLUTIONS, DATA CENTER DESIGNERS MUST UNDERSTAND THE IMPACT OF NUMEROUS FACTORS ON THEIR CABINET DESIGN

2. HOW TO DISSIPATE EXCESS HEAT (WATTS) – UNDERSTANDING ΔT AND CFM

WHEN IT COMES TO COOLING A DATA CENTER, ONE EQUATION IS KEY:

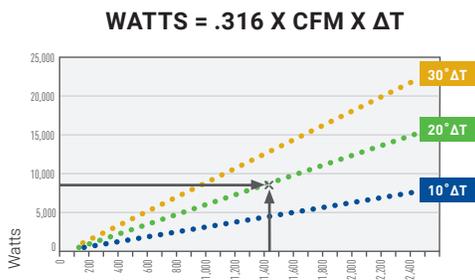
$$\text{WATTS} = .316 \times \text{CFM} \times \Delta T$$

Cold intake air and airflow (CFM) are both required to cool network equipment. These two factors work directly to dissipate, or carry away, the heat that network equipment produces. While increasing either factor will increase the amount of heat that is dissipated, there are limits. Air that is too cold results in thermal expansion and condensation issues. Airflow that is too high results in acoustic and physical limitations. A cooling system that relies on too much of only one of these factors usually results in higher capital and operating costs. Finding an ideal balance between cold air and airflow allows optimal heat dissipation, protecting network equipment from the consequences of overheating.

ΔT

In the above equations, ΔT, measured in degrees Fahrenheit, refers to the difference between the intake air and exhaust air, or the amount of heat that is carried away by the stream of air exiting the heat load. A greater temperature difference signifies more heat being removed. While it is difficult to constantly deliver the coldest air to all equipment—in particular equipment installed in upper rack units—data centers that maintain a range between 10° and 30° ΔT carry a high level of confidence that equipment will dissipate the amount of heat that it produces, staying within its thermal limits.

- The area between 10° and 30° ΔT represents capable cooling in a typical well-designed data center.
- Knowing two of the three quantities allows operators to solve for the third—in this equation, knowing either watts or CFM solves for ΔT.
- For example: 9,000 watts in a data center cabinet with a ΔT of 20° requires 1425 CFM



Note: More than 30° ΔT can cause humidity and condensation issues.

CFM

Another way to increase the amount of heat dissipated or removed from network equipment is by increasing the airflow, expressed in cubic feet per minute (CFM). An increase of CFM—the amount of airflow across a given area in a given time—results in increased heat removal.

CFM can be achieved through the use of fans. The larger the fan, the more CFM it provides. An increase in RPM (the speed at which the fan circulates) as well as the size or quantity of fan blades results in a higher CFM, but as these factors increase, so do several others that can be detrimental to the data center equipment. For example, acoustic noise is one consequence of high levels of CFM. Additionally, at very high CFM, physical forces also come into play that can cause damage to electrical equipment. Plus, the higher the CFM, the higher the upfront capital costs, as well as ongoing operational expenses.

- Most data centers operate between 10° and 30° ΔT. Using this and the CFM of the fan, the necessary quantity of cooling (watts) can be determined.
- The bigger the fan, the more airflow that moves across equipment. The amount of airflow and the temperature of the air determine the amount of heat dissipation that can occur.

Note: An increase in fan size will increase acoustic noise at the same motor RPM.

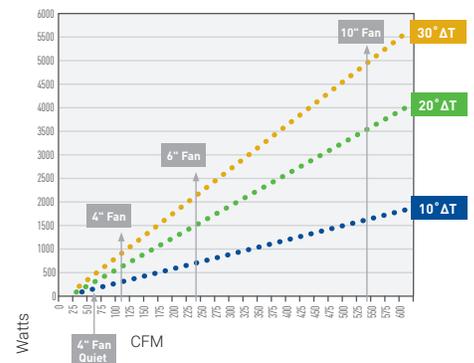
Watts = .316 x CFM x ΔT

or: **CFM** = $\frac{\text{Watts}}{(.316 \times \Delta T)}$

or: **ΔT** = $\frac{\text{Watts}}{(.316 \times \text{CFM})}$

NO MATTER HOW HIGH THE CFM, A FAN CAN NEVER COOL LOWER THAN THE SURROUNDING AMBIENT TEMPERATURE

TYPICAL 4", 6" & 10" COOLING FAN PERFORMANCE



The data center must provide enough CFM to the equipment to prevent overheating. = CFM^{req}

3. RAISED VERSUS NON-RAISED FLOOR DATA CENTERS

In raised floor data centers, perforated floor tiles are used to distribute cold air—allowing cold air to enter a cold aisle and be used for equipment intake. Typically, each floor tile is 24" x 24" and allows 200 to 500 CFM, depending on the percentage of open space and under-floor air pressure. Open space is the amount of cutout area that allows airflow through the tile: 25 percent open space signifies that of the 24" x 24" tile, 25 percent is open to airflow. Older floor tiles usually provide 25 to 35 percent open space, while some newer models feature open spaces approaching 60 percent of the total tile area. Typical floor tiles allow for 200 to 500 CFM, providing 1,500 to 4,000 watts of cooling. The exact amount of cooling (watts) per tile is determined by the actual airflow and temperature: $\text{Watts} = .316 \times \text{CFM} \times \Delta T$.

- Floor tiles are available in two basic open space styles
- Traditional 4-square tiles featuring 25 percent open space are the most common
- Newer, high-performance tiles with 56 percent or more open space are composed of grated steel, which allows them to be more than half open
- Although newer tiles feature more open space, a whole system installed with these may be difficult to support—computer room air conditioner/handler units often do not have enough air pressure beneath the floor to distribute the required CFM throughout the data center

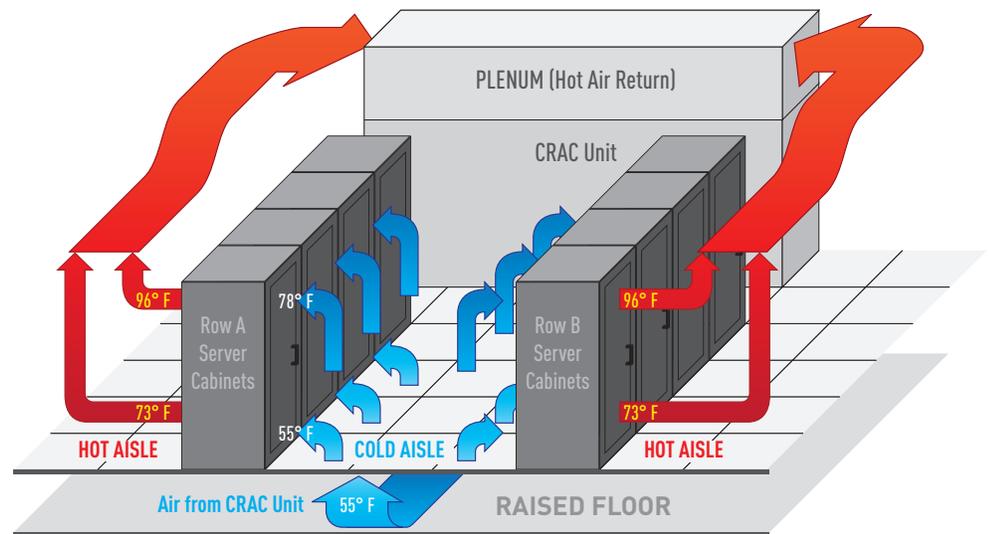
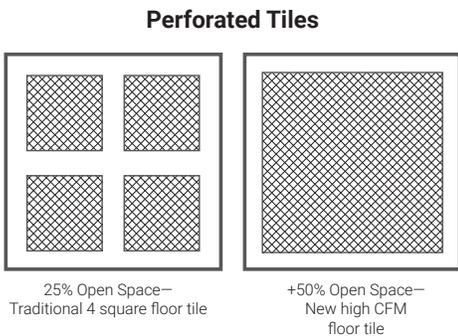
Air recycling occurs through the use of Computer Room Air Conditioners, CRAC, or computer room air handlers, CRAH, units. These air conditioners remove the heat from the hot exhaust utilizing either refrigeration or chilled water. The heat is directed outside of the building, and the reconditioned cold air is pumped into the raised floor plenum—then up through the perforated tiles and into the cold aisle.

Since the raised floor is the plenum for distributing cold air to the data center, the placement of perforated floor tiles is critical. These tiles are the point of entry for all cold air in the data center, and where they are located will determine the cooling capability and overall efficiency of the facility. They must be strategically located to ensure proper cold air distribution. The exact placement of floor tiles is often set once the cabinet and equipment are in place, through a process of measurement as well as trial and error. A more sophisticated, alternative method uses Computation Fluid Dynamics, or CFD, models to determine tile placement. In this approach, a model of the data center is created with as accurate information as possible. A floor plan is developed, and through modeling simulation, the placement of the floor tiles that will allow ideal airflow and cooling is attained.

ASSUMPTIONS:

- Average $\Delta T = 20^\circ$
- CFM = 200 to 500
- Watts cooling = 1264 to 3160 watts/floor tile

Complications can result from perforated floor tiles being positioned either too far from or too close to the CRAC/CRAH unit(s). Floor tiles placed too far from the CRAC/CRAH unit(s) produce little to no airflow. Conversely, due to the high velocity of air exiting the bottom of a unit, perforated floor tiles located in close proximity to the points at which the bottom of a cabinet meets the raised floor plenum may be exposed to the Venturi Effect. The Venturi Effect occurs when a rush of air flowing below the perforated tile causes suction into the flow. Instead of having an upward flow of cold air, a downward draw of the air from above the floor may occur—the opposite of what is needed.



3. RAISED VERSUS NON-RAISED FLOOR DATA CENTERS (CONTINUED)

TYPICAL DATA CENTER RAISED FLOOR CONDITIONS

- Plenum air temperature = 55° to 66° F
- Air temperature increases to approximately 78° F as it moves upward to higher mounted equipment
- Typical ΔT = approximately 20° F
- Hot aisle temperatures = approximately 100° F

Once the perforated floor tiles' positions are determined, additional steps can be taken to ensure proper airflow in the data center.

CHALLENGES ASSOCIATED WITH RAISED FLOOR DATA CENTERS

- Difficulty with seismic ratings. In areas with earthquakes it is almost impossible to expect five nines (99.999 percent up time) availability of the network using a raised floor.
- Raised floors were originally designed before the need for frequent moves, adds and changes due to the short operating life of IT equipment as well as frequent upgrades and other changes. Because of this, it is very difficult to access under-floor cabling to accommodate changes.
- Older locations might not have sufficient headroom for raised floors.
- The area under a raised floor is considered to be an air plenum that must meet special fire codes and include conduit requirements or special polymers and plenum rated cables.

NON-RAISED FLOOR DATA CENTERS

The primary elements of a raised floor design are intended to provide a cool air distribution system, conduits for power and data cabling, and a copper grid for grounding. However, many of these objectives can be achieved via new technologies. In fact, many large data centers now use non-raised floor designs. Along with the removal of raised floors, drop ceilings are also eliminated. Today, non-raised floor data centers can not only provide sufficient airflow and support structures, but they may also eliminate some of the issues associated with raised floor designs.

ADVANTAGES OF NON-RAISED FLOOR DATA CENTERS

- Cabling is located overhead and is easy to access.
- Cleaning the data center floors is much easier, since there is no place for dust to hide
- Weight capacity concerns are eliminated with non-raised floor data centers. Cabinets full of equipment weigh as much as 3,000 lbs, which can be a major issue with raised floor designs.
- Costs are lower since there is no raised floor with which to contend. Outside studies suggest raised floors cost approximately \$20.00/sq. ft., plus the added costs associated with power and data cabling.
- Blockages, barriers and obstructions to cold airflows are eliminated. In many instances, perforated tiles deliver substantially less airflow than expected due to blockages and obstructions under the raised floor.

A **CRAC/CRAH CFM** must be at least equal to or more than the total CFM needed by the network equipment.

A COOLING
STRATEGY THAT
DELIVERS COLD AIR
TO THE EQUIPMENT
MUST BE
CAREFULLY
DESIGNED

4. DATA CENTER ENERGY COSTS AND SAVINGS

HYPOTHETICAL DATA CENTER

- 100 cabinets with an IT load of 5.0 kW per cabinet (average)
- 3000 sq. ft. raised floor area (5 rows of 20 cabinets with 4' aisles)
- Total IT load = 500 kW-hr (4,380,000 kW/year)
- Total power used annually 10,556,000 kW
- 0.401 kW-hr (400 watts) per sq. ft. of raised floor data center
- Energy cost \$0.10 / kW-hr
- Total annual energy costs \$1,055,600.00
- PUE = 2.4 (typical average data center) DCiE = 41.6%
- Annual cost per cabinet = \$10,556
- Annual cost per 1kW of IT load = \$2,111

SAVINGS CHART

This chart represents typical savings and is a guideline. Each data center will present many challenges, and the activities shown will vary significantly in their effectiveness and savings.

Activity	PUE	Annual Power Cost	Annual Savings
Data Center Typical (Hot Aisle/Cold Aisle)	2.40	\$1,055,600	—
Blanking Panels (In all open RU locations)	2.38	\$1,040,000	\$15,600
Floor Brushes (Tile Cutouts)	2.35	\$1,025,000	\$15,000
Perforated Tile Placement	2.30	\$1,007,000	\$18,000
CRAC/CRAH Unit – Duct Work	2.27	\$994,000	\$13,000
Drop Ceiling Return Duct Work	2.23	\$976,000	\$18,000
Cabinet Layout Optimized	2.20	\$964,000	\$12,000
Raised Floor Optimized	2.15	\$938,000	\$26,000
Containment – Cold Aisle	2.10	\$920,000	\$18,000
Hot Aisle Containment / Chimney	2.00	\$880,000	\$40,000
Liquid Cooling* (Closed Cabinet System)	1.75	\$770,000	\$286,600

*Note: Liquid cooling allows for higher density (fewer cabinets and a smaller sized data center) which is not calculated in the savings. Liquid cooling is a stand alone solution not cumulative to the other activities.

Capital construction costs are not included in calculations.

PUE
(Power Usage Effectiveness)
= Total Facility Power/
IT Equipment Power

The total amount of power used by the data center compared to the power used by the IT equipment. A typical data center with a hot aisle / cold aisle cabinet configuration has a PUE number around 2.4. A very efficient data center brings this number down to 2 or even lower. Some purpose-built data centers are claiming PUE numbers of 1.75 or even lower.

DCiE
(Data Center Infrastructure
Efficiency = IT Power
Used/Total Facility Power)

DCiE is another metric for measuring data center efficiency. It is similar to PUE but is inverted and shown as a percentage. A PUE of 2.4 = 41.6%

5. ENCLOSURE OVERVIEW

SERVER TECHNOLOGY IMPROVEMENTS AND THE NEED FOR MORE EFFICIENT COOLING.

In a modern data center, servers perform all the processing work, creating power and thermal management challenges.

Traditionally, servers were application-specific, with each server dedicated to a single client or application, which resulted in many servers sitting idle or doing very little work while still consuming power and producing heat. To improve efficiency, servers are now being virtualized. Instead of having separate units each dedicated to a particular process, servers are networked together. While fewer servers are needed because they are operating more of the time, total power usage and thermal densities increase.

Improvements in server technology have resulted in an interesting paradox. Higher amounts of processing require more power and subsequently produce greater amounts of heat that must be dissipated. Yet as servers get smaller or as fewer servers located in smaller spaces operate at higher levels (virtualization), their ability to dissipate heat becomes more challenging, since more power in less space is counterintuitive to airflow and heat removal. Additional measures must be planned to compensate for these high power and thermal loads, making the way that cabinets are arranged, configured and designed more critical than ever in ensuring a data center's success.

As noted in previous sections, the amount of heat that can be removed, or dissipated, is governed by the incoming air temperature and the volume of air.

LEED (LEADERSHIP IN ENERGY AND ENVIRONMENTAL DESIGN)

Leadership in Energy and Environmental Design (LEED) is a rating system devised by the United States Green Building Council (USGBC) to evaluate the environmental performance of a building and encourage market transformation toward sustainable design. The system is credit-based, allowing projects to earn points for environmentally friendly actions taken during construction and use of a building. LEED was launched in an effort to develop a "consensus-based, market-driven rating system to accelerate the development and implementation of green building practices." The program is not rigidly structured; not every project must meet identical requirements to qualify.

Cabinet design and configuration and data center layout can provide valuable energy savings (Energy and Atmosphere Category) and can be used with other attributes in garnering the facility LEED certification. LEED certification pertains to the entire building. Data centers are not specifically addressed but would most likely fall into the "Optimized Energy Performance" category.



PROJECT CHECKLIST: LEED – NEW CONSTRUCTION (NC) v3

Sustainable Sites		Points
Credit 3	Brownfield Redevelopment	1
Credit 5.1	Site Development, Protect or Restore Habitat	1
Credit 5.2	Site Development, Maximize Open Space	1
Credit 6.1	Stormwater Design, Quantity Control	1
Credit 6.2	Stormwater Design, Quality Control	1
Credit 7.1	Heat Island Effect, Non-roof	1
Credit 7.2	Heat Island Effect, Roof	1
Energy and Atmosphere		
Prerequisite 2	Minimum Energy Performance	Required
Credit 1	Optimize Energy Performance	1 - 19
Materials and Resources		
Credit 1.1	Building Reuse, Maintain 55%, 75%, 95% of Existing Walls, Floors and Roof	up to 3
Credit 2	Construction Waste Management, Divert 50% or 75%	up to 2
Credit 4	Recycled Content, 10% or 20% (post-consumer plus ½ pre-consumer)	up to 2
Credit 5	Regional Materials, 10% or 20%	up to 2
Innovation and Design Process		
Credit 1	Innovation in Design	1 - 5
Credit 2	LEED Accredited Professional	1
Project Totals		25 - 24

Up to 2 additional points can be earned—must be submitted and approved (not included in total)

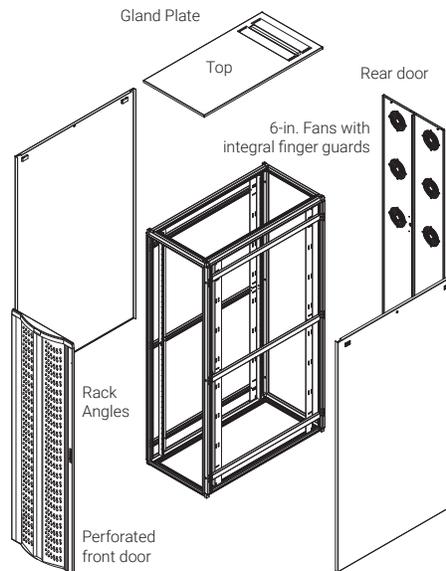
5. ENCLOSURE OVERVIEW (CONTINUED)

CABINET CHARACTERISTICS THAT AFFECT SERVER PERFORMANCE

PERFORATED DOORS

Each server manufacturer has specific guidelines concerning the amount of open space required on cabinet front doors. Open space refers to the space open for airflow through the door when compared to the equipment's intake air opening—not to be confused with the space that is open between the front of the servers and the cabinet. To provide servers with enough cold airflow to dissipate the heat generated and ensure that the cold air is not restricted from entering the installed equipment, doors must have at least 50 percent open space. Most cabinet manufacturers increase this amount to over 60 percent to avoid potential equipment manufacturers' warranty issues.

Mesh or perforated doors should be used on the front and rear of all server cabinets. While Plexiglas front doors are sometimes ordered for aesthetic reasons, they provide limited airflow and should be avoided unless non-heat producing equipment—e.g. patch panels—are utilized. Perforated or mesh doors, on the other hand, promote proper airflow to the equipment, avoiding overheating. Equally important, the rear doors should also be perforated or mesh to allow exhaust air, with heat from the servers, to flow away from the cabinet, into the hot aisle and back to the cooling unit where it can be recycled back to cold air. All servers should be placed 4 to 6 inches from the front and rear cabinet doors to provide sufficient space for accessories, such as handles and cables, as well as airflow. It is important that cabinets with logo nameplates or decorative features do not block the front perforated airflow.



REAR-MOUNT FANS

Increasing the amount of heat that is removed from a cabinet can be achieved by placing fans on the rear doors. When these fans provide higher CFM than the fans in the server, they can be used in conjunction with the server fans to increase the airflow in specific areas and assist in eliminating hot exhaust air from the cabinet. In addition to increasing airflow, rear-mount fans prevent hot spots from forming. Rear-mount fans can be monitored and controlled for maximum efficiency and can switch off or adjust their RPM to match the cabinet requirements at a specific time. Many times additional fans can also be added to the servers themselves.

CABINET TOPS AND FANS

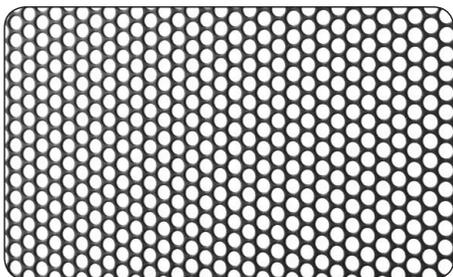
In random layout data centers, the best way to remove heat from the top of a cabinet is by installing a perforated top or a

top-mount fan. However, in more efficient hot aisle/cold aisle layouts, the use of top-mount fans can actually impede heat dissipation and overall server thermal performance by drawing cold air away from the server and mixing the cold air with hot exhaust air, wasting the energy that was originally used to create the cold air. Concerns also exist that a top-mount cooling fan can actually starve the top few servers of fresh air by overpowering the server fans.

The most efficient raised floor cooling systems include cabinets that feature solid tops, forcing hot exhaust air to exit the back of cabinets. Blocking plates utilized to keep cold air from bypassing equipment can also improve efficiency.

CABINET BOTTOMS

Issues can result when cooled air is delivered directly under a cabinet. First, servers located toward the bottom of a cabinet may block airflow. Additionally, the cold air that is directed toward the rear of the cabinet may spread out in all directions, which can mix hot and cold air and waste energy. These issues can be prevented by providing cold air in front of the cabinet through the perforated tile. In some cases, manufacturers use floor-ducted cabinets that ensure all equipment is provided the same cold temperature intake air by creating a cold air plenum in front of the equipment. This ensures the equipment mounted high in the cabinet receives the same cold air as the equipment located below, eliminating temperature stratification.



Note: Beyond 60% perforation appears to have diminishing returns for improving airflow.

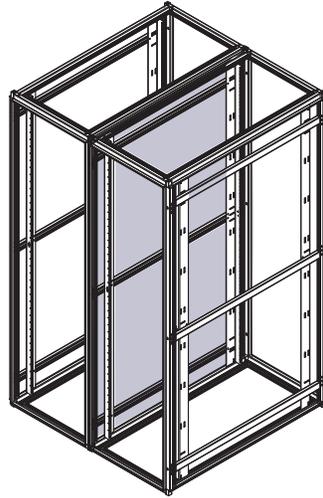
5. ENCLOSURE OVERVIEW (CONTINUED)

ACCESSORIES TO ASSIST WITH CABINET AIRFLOW

BLANKING PANELS

Blanking panels are used to fill all empty RU spaces. Blanking panels are installed to prevent cold air from rushing through the cabinet, bypassing equipment and mixing with hot exhaust air. A solid wall comprised of either equipment or blanking panels ensures that cold air can only pass through equipment to dissipate. Blanking panels provide a simple and economical way to assist in directing cold air; however, all empty rack spaces must be filled for optimal results. Based upon the data center scenario presented on page 7, the use of blanking panels in all open RU locations saves \$15,600 annually.

- Most commonly cited improvement to assist with cabinet airflow
- Prevents cold air loss in cabinet front
- Prevents recirculation
- All server manufacturers recommend the use of blanking panels
- Models include snap-in, plenum-rated plastic panels and traditional metal blanking panels



BARRIER PANELS (LOCATED BETWEEN CABINETS)

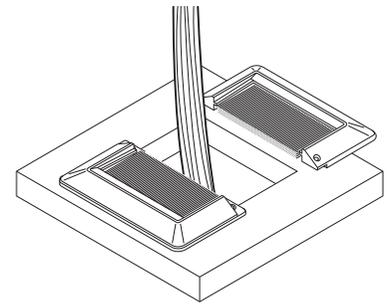
Installing barrier panels between cabinets and as end caps is extremely important in preventing air from improperly recirculating—an issue that is more prevalent when enclosures are not fully populated and when servers of multiple sizes and loads are used. Barrier panels are simply placed between servers to segregate each cabinet and better control airflow.

FLOOR BRUSHES

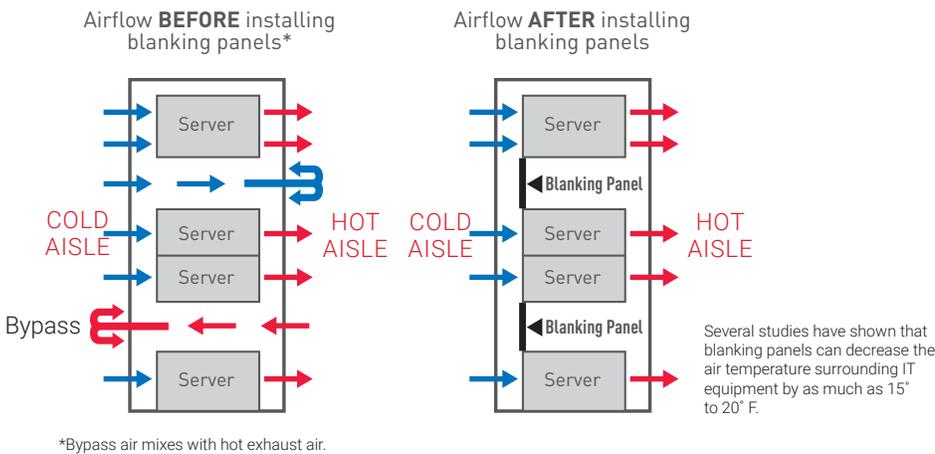
Raised floors function as more than just a plenum for distributing cold air to network equipment. They can also be the main

cable distribution means for the data center. It is common for power and data connection cables to enter a cabinet through rear cutouts. However, it is important that cold air doesn't leak through cable entryways in either the floor or the cabinet, and into the hot exhaust air—which would result in the same negative effects as when cold air bypasses network equipment altogether. To prevent this, data centers use brushes or other cable egress products to block airflow. Based upon the data center scenario on page 7, the use of floor brushes saves \$15,000 annually.

While materials such as cardboard and insulation can break up and cause particles to pervade servers, compliant materials will allow room for cables and tightly conform around them to prevent cold air leakage. Additionally, unused floor tile locations should always be sealed.



Several outside studies indicate that 50 to 80 percent of conditioned cold air can escape through open cable cutouts.



ALL COLD AIR SHOULD BE USED TO COOL EQUIPMENT. IF ALLOWED TO MIX WITH THE HOT EXHAUST AIR, THE COLD AIR IS WASTED AND EFFICIENCY IS LOWERED

6. CABLE MANAGEMENT

Cable management can have a negative impact on cooling if not properly structured and maintained. Poor cable management can lead to poor airflow, excessive heat absorption, and increased cold air bypass.

THE NEGATIVE EFFECTS OF POOR CABLE MANAGEMENT INCLUDE:

- Cable arms adversely impact airflow
- Cables placed in the hot air stream absorb and hold heat within the cabinet
- Cable pathways that extend from the cold side to the hot side can allow for cold air bypass
- Excessive and overall poor cable management leads to air blockage
- Cables that are too long for the application, excessive cable slack and unused cables can all hinder cooling capabilities
- Power cables of excessive length and obstructive PDUs can block airflow



GOOD CABLE MANAGEMENT CAN BE ATTAINED WHEN:

- Cables are neatly contained and bundled when possible
- Cables are moved out of cold and hot airstreams
- Correct cable length is used, and cable slacks or unneeded loops are avoided
- Wider cabinets with vertical cable management areas are provided when possible
- Air blocking devices are used around all cable egress points from cold to hot areas—including brushes, grommets, gland plates and other conforming devices that impede airflow
- VELCRO ties reduce cable bulk and move cables out of the airstream areas
- Vertical cable managers provide supporting fingers and space for cables outside of airstream areas
- Strategically located horizontal cable managers provide support and space, and route cables from left to right—keeping airflow areas free of obstructions
- PDUs and power cables are positioned out of airflow areas

VELCRO is a trademark of Velcro Industries B.V.

FANS AND FAN TRAYS

To eliminate hot spots within the cabinet, fans can direct the hot exhaust out of the cabinet and into a hot area. Mounted above cabinet equipment, 19" rack-mount trays are often used to eliminate hot spots from the top of the cabinet. Strategically located fans placed throughout the cabinet will assist in mitigating hot spots.



7. CABINET CONFIGURATIONS

Cabinets provide more than just network equipment support. They are key to entire data center systems, ensuring the interoperability of various rack-mount equipment. TIA/EIA standards serve as a design guide to ensure that equipment and cabinets are in sync to support each other.

Cabinet configurations include many aspects that are critical in maintaining an operative and efficient data center, including:

- Physical support for equipment
- Support of overhead, interior and sub-floor cable pathways
- Equipment access
- Access control (prevents unauthorized access to equipment)
- Equipment cable entry points
- Cable management
- Thermal management
- Aesthetics
- Flexibility and expandability to support future data center updates

There are many cabinet configuration options available to assist with providing the amount of airflow and cooling needed for servers. Options range from simple and inexpensive to costly and complex, with capabilities increasing respectively. However, the goal of each configuration is the same: help deliver the necessary amount of cool air to each server.

The data center cabinet layout, type of equipment installed and amounts of power needed and heat generated are among the factors that will influence cabinet requirements. The variety of cabinets installed is dictated by the various needs of the data center. A continuum of cabinets—from simple passive to integrated—provides a general guide to cabinet configurations.

Equipment must be no wider than 17.72"; space between rack angles for 19" rack equipment no less than 17.75". Distance between holes for 19" rack space is 18.312"

CABINET AND EQUIPMENT MANUFACTURERS ADHERE TO EIA 310 STANDARDS TO ENSURE THAT EQUIPMENT WILL FIT INSIDE CABINETS

CONTINUUM OF COOLING Future-Proof Data Center



7. CABINET CONFIGURATIONS (CONTINUED)

RANDOM DATA CENTERS

The simple side of the cooling spectrum starts with random data center layouts. In these data centers, cabinets are not organized to distribute cold intake and hot exhaust air. Heat dissipation is possible by driving hot airflows above the cabinets to avoid blowing hot exhaust air into the cabinets (hot intakes to other network equipment). Although still found in some older data centers, random cabinet orientations are being replaced by hot aisle/cold aisle as well as other more complex systems.

Data Center Trends

Data center designs continue to evolve, and the cabinets within them also change. Currently, hot aisle/cold aisle configurations are the defacto standard, but as heat dissipation needs increase, more sophisticated layouts are increasing in popularity.



PASSIVE VERSUS ACTIVE COOLING

Data centers rely on fans to direct CFM. All cabinet configurations use either passive or active cabinets, depending upon cooling needs.

PASSIVE COOLING

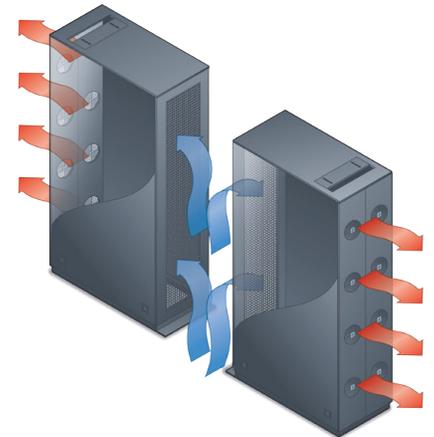
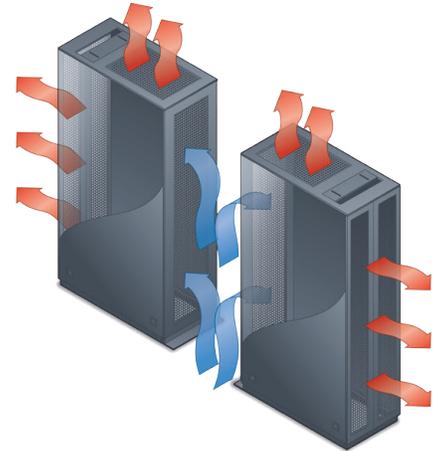
Passive cabinets are designed to maximize the ability of the internally mounted equipment to cool itself via its own fans. In this method, airflows are produced by the network equipment (e.g., fans in the servers). Factors to consider:

- Airflow must not be restricted:
 - Cables should be neatly arranged
 - Recommended open space requirements should be followed
- Heat dissipation is dependant on the equipment's CFM:
 - 1 and 2 RU servers have limited CFM
 - Blade chassis servers utilize larger fans that can produce greater CFM

ACTIVE COOLING

While passive cooling simply relies on the network equipment (servers and switches) and its fans, active cabinets utilize additional strategically placed fans to supplement airflow, thus increasing heat dissipation. These fans may be used to pump cold air to network equipment and/or assist in removing hot exhaust air from the cabinet, delivering it to the hot aisle or even back to the CRAC/CRAH unit(s). For example, adding fans to the rear door of a cabinet can essentially "super-charge" equipment and increase thermal performance by increasing airflow and thus dramatically increasing heat dissipation. Factors to consider:

- Active cooling provides increased CFM
- Fans are typically located on the exhaust side of servers
 - Rear door fans are the most common
 - Fans can also be top-mounted; however, this is typically not recommended
- Cables must be neatly arranged to not interfere with airflow



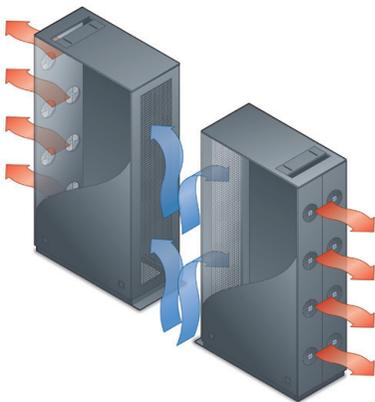
7. CABINET CONFIGURATIONS (CONTINUED)

HOT AISLE/COLD AISLE CABINET CONFIGURATION

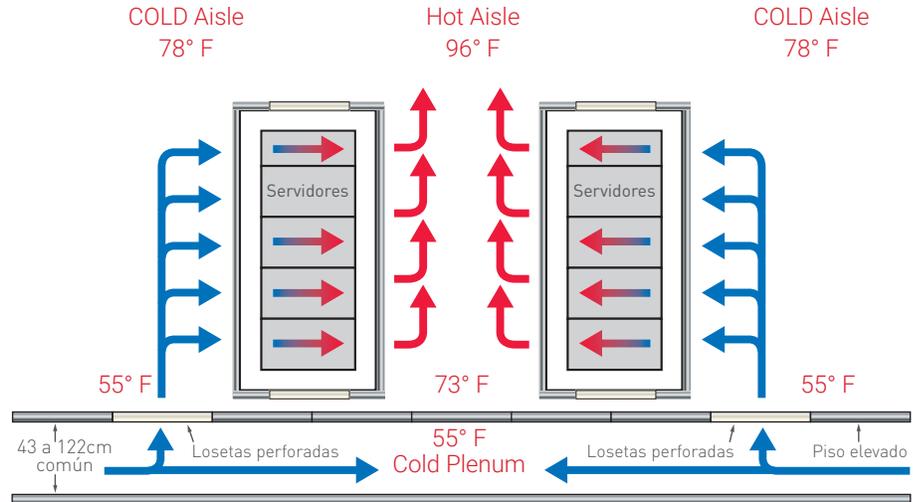
By utilizing passive or active cabinets to cool equipment, a hot aisle/cold aisle layout is designed to separate hot exhaust air from cold intake air. CRAC/CRAH units are strategically placed to draw in the hot air, recycle it to remove the heat and then disperse it into the cold aisles. Cabinets are placed on both sides of an aisle with the front sides facing each other. In turn, the backsides of the cabinets face the backsides of the next row of cabinets, and the pattern continues.

The cabinets on both sides of a cold aisle draw cold air in through their front intakes, and once it is utilized, the hot air is exhausted through the cabinet rears, creating hot aisles that alternate with cold aisles. The hot air is recirculated to the CRAC/CRAH unit(s), preventing it from being drawn into other cold air intakes. This hot aisle/cold aisle layout is universally accepted and currently the most common type of cabinet configuration.

COLD AIR THAT MIXES WITH HOT AIR WITHOUT GOING THROUGH EQUIPMENT IS WASTED



Hot Aisle/Cold Aisle Temperature Variations



KEY CONCEPTS

- Servers pull in cold air from the front to cool themselves and exhaust the hot air out the back of the cabinets
- All cold air should be utilized to carry heat away from equipment

In hot aisle/cold aisle configurations, cabinets play a critical role in airflow management to and from mounted equipment. The goal is to make certain that only cold air is used to cool network equipment, and the cabinet helps ensure this by providing pathways to equipment intakes and exhausts so cold air doesn't bypass cold aisles and mix with hot exhaust air. According to the data center scenario presented on page 8, an optimized cabinet layout that prevents mixing of hot and cold air saves \$40,000 annually.

TYPICAL HOT AISLE/COLD AISLE CONDITIONS

- Cold aisle temperature = 55° to 78° F
- Hot aisle temperature = 73° to 96° F
- Typical ΔT = 15° to 20° F
- Each cold aisle floor tile provides 200 to 500 CFM
- Coldest temperatures are always at the base of the floor

HOT AISLE/COLD AISLE DENSITIES

Although the hot aisle/cold aisle (HACA) layout is still widely used in data centers, its limitations are now more evident as more blade servers are utilized. Once densities increase to 5 kW and above, the traditional HACA layout will have difficulty handling the resulting heat loads. Layouts that use traditional 25 percent open-space floor tiles— which allow for an average of 300 to 500 CFM of cold air— can only accommodate cabinet heat dissipation of 3 to 5 kW of server or other IT equipment.

To address hot spots in this design and accommodate higher heat loads, many IT managers select grate-type floor tiles that can distribute between 700 and 900 CFM. The use of floor tiles that feature a higher percentage of open space can increase overall cabinet heat dissipation up to 8 kW per cabinet, if the CRAC/CRAH unit can provide the needed CFM. This still leaves the cabinets mostly empty due to the heat dissipation limit. Also, this increased cooling could create a new hot spot by tapping the cooling capacity from another area, if the tiles are installed improperly. Thus, installing grate-type floor tiles requires careful analysis of the IT environment.

7. CABINET CONFIGURATIONS (CONTINUED)

HACA

The HACA layout was designed to assist in airflow distribution. Although the intended arrangement is logical, modern heat loads challenge HACA's capabilities. Many outside studies of data centers state that on average, only 40 percent of available cooling is utilized, due to the bypass and over-provisioning of cool air used to eliminate hot spots. Despite this gross oversupply, approximately 10 percent of most data center racks still experience hot spots. Ensuring airflow reaches the proper areas is key to effective, efficient data center design.

HACA LEGACY DATA CENTERS

Most legacy data centers have the following characteristics:

- Air systems (usually perimeter CRAC/CRAH units) that deliver cold air at 55° F via ducting from beneath a raised floor or from overhead ducting
- Use of perforated tiles
- Standard HACA layout with a minimum of 4 ft. of separation between cold aisles (two tiles) and 3 ft. between hot aisles

Although common, this layout often proves inefficient. The most common causes of inefficiency are the following:

Bypass airflow – Valuable conditioned cold air that flows past equipment intakes and intermixes with hot exhaust air is referred to as bypass airflow. Bypass air can flow through empty RU spaces, around mounting angles, above and below equipment, and through cable cutouts and other areas.

Recirculation – When hot air from the servers or other IT equipment is exhausted, it may recirculate into its own intake at a higher temperature. It typically occurs due to hot air recirculation above a cabinet or around the side of an end-row cabinet, and it affects the loads at the highest point in a cabinet. The shortest supply path for

recirculated air is the path directly from the equipment's own exhaust. Therefore, side barrier panels and front blanking panels are very important for preventing recirculation.

Air stratification – In a data center, a layering effect of temperature gradients that occurs from the bottom to the top of the enclosure is called air stratification. In a typical raised-floor environment, air is delivered at 55° F via perforated floor tiles. Once the cold air enters the data center, it quickly heats up as the air rises. ASHRAE recommends the air temperature designated for server intakes remains between 68° and 77° F—although at the highest spots in the cabinet, the air may reach 90° F at server intakes. Data center managers may try to combat this effect to efficiently cool the highest servers in each rack, but in their attempt to accommodate for the temperature increase, many managers supply too much cold air. The best solution is not to supply more air, but rather, to avoid bypass airflow and recirculation.

THE ADVANTAGES OF SWITCHING FROM HACA TO A CLOSED-COUPLED DESIGN

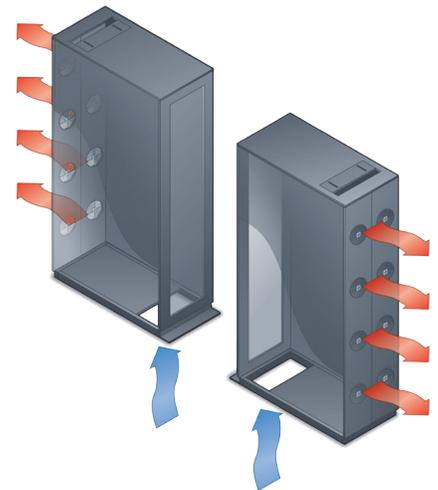
For high-density data centers, a closed-coupled design can mitigate the problems and inefficiencies of the HACA design. A closed-coupled, or closed-loop, design takes all the air supplied by CRAC/CRAH units and delivers it to rack-mounted IT equipment intakes. Then, all the hot air exhausted from the IT equipment is delivered directly back to the air intake of the CRAC/CRAH unit. Four basic cabinet layouts may be used in a closed-coupled design:

- Floor-ducted solution
- Chimney- or direct-ducted vertical exhaust
- Cold air or cold aisle containment
- Hot air or hot aisle containment



FLOOR-DUCTED SOLUTION

In this type of active configuration, the cabinet integrates a raised floor into its heat dissipation performance (see Raised versus Non-Raised Floor Data Centers, pg. 5). Ductwork and a closed area in front of the equipment ensure that each device within the cabinet is provided with uniform cold air, regardless of the cabinet position. Optimum performance, however, is dependant on properly balanced incoming and outgoing airflows. When properly designed, floor-ducted cabinets can provide extremely high heat dissipation and cooling efficiency.

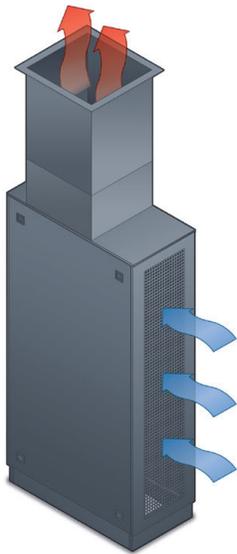


7. CABINET CONFIGURATIONS (CONTINUED)

CHIMNEY, DIRECT-DUCTED VERTICAL EXHAUST SYSTEMS

The key to an efficient data center—which results in increased capacity and lower energy costs—depends on effective cooling, which depends on segregating hot and cool air, as well as controlling and directing cold intake and hot exhaust air. The chimney approach, a direct-ducted method, consists of a passive exhaust system that prevents hot exhaust from recirculating and mixing with cold air that supplies the active equipment. The data center scenario on page 7 shows that using the chimney method can save \$40,000 annually.

Managed and contained exhaust systems are ideal for both raised floor and slab applications. Since the hot exhaust airflows are contained and directed back to the cooling units, cold air can be provided by many means: from above, perimeter, raised floor, etc.



CHIMNEY CABINET

Chimney cabinets consist of a cabinet with top ductwork that directs hot exhaust air either to an above drop ceiling, which acts like a hot plenum, or ductwork that further directs the airflow to the intakes of the CRAC/CRAH units. The ceiling plenum has a slightly lower pressure than the room, aiding in removing hot air exhaust from the data center.

AIRFLOW MANAGEMENT

Providing a clear airflow path for hot exhaust air is vital in ensuring that heat is removed from the cabinet or hot aisle. To achieve proper airflow management, the following considerations should be taken into account:

- All IT exhaust should have a direct return path to the CRAC/CRAH units, eliminating hot/cold air mixing from the top and sides, and allowing higher CRAC/CRAH unit efficiency by supplying hotter air into the return.
- Supply (cold air) temperature can be increased closer to the ASHRAE upper limit of 77° F.
- All cool air that is delivered to the active equipment is utilized, and the supply air is more flexible—not just supplied from the floor tile in front of the server. Cold air can be delivered from anywhere in the room.
- The large mass of cold air allows significant time in which a data center may continue operating during a CRAC/CRAH unit failure event, allowing for failure resolution before the data center is forced to shut down.

CHIMNEY CABINET COOLING IS ONE OF THE MOST EFFICIENT PASSIVE AIRFLOW MANAGEMENT OPTIONS FOR A DATA CENTER



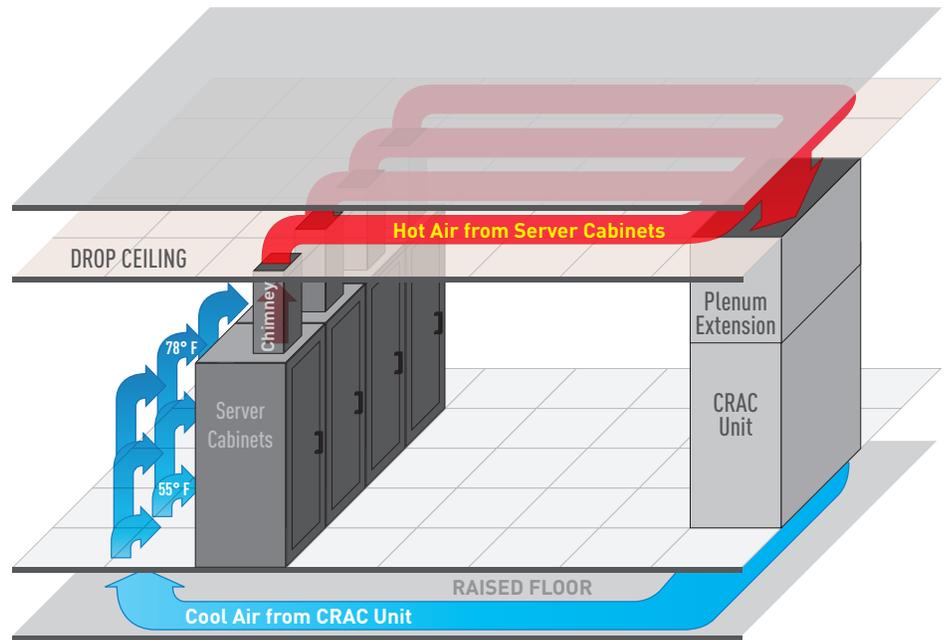
7. CABINET CONFIGURATIONS (CONTINUED)

ADVANTAGES OF THE CHIMNEY METHOD

- Many outside studies suggest that the chimney method lowers cooling system costs by at least 25 percent
- Low initial capital costs; no ongoing maintenance expenses
- The number of CRAC/CRAH units in the data center could be reduced or the future cooling capacities could be increased
- Cabinets do not have to be orientated in hot aisle/cold aisle configurations
- Outside studies have successfully tested hot air extraction for heat loads up to 20kW

CHIMNEY BEST PRACTICES

- All empty RU spaces should be filled with blanking panels and additional front openings should be sealed.
- The space behind network equipment (servers) must be free and clear of obstructions (cables, PDU power cords, etc.) to avoid “choking off” hot exhaust airflows. The rear door should be solid, not perforated; this forces the air up the chimney.
- The bottom rear of the cabinet should be analyzed. If hot eddies are evident, air deflector plates should be installed to direct airflow upward or out of the cabinet. (see picture left)
- Cable pathways should be routed overhead and on the front side of the cabinet.
- Deeper and wider cabinets will provide more space behind the servers, supplying sufficient exhaust area to remove the hot air.
- Use caution when placing the deepest equipment near the top of the cabinet, as this arrangement can create choke points for the exhaust air.



HEAT DISSIPATION CAPACITY

The amount of heat that can be successfully dissipated through the use of duct hot air systems is increased significantly: heat dissipation of over 20kW is common. Since these systems segregate and direct airflows, they improve efficiency and lower operations costs. The chimney method reduces wasted cold air by 99 percent, and less than 1 percent of cold air is allowed to mix with hot exhaust air.



Rear air deflector plate directs hot airflow efficiently up the top-mounted chimney.

7. CABINET CONFIGURATIONS (CONTINUED)

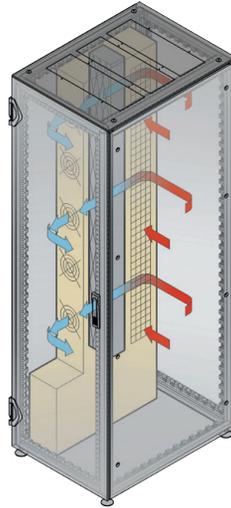
LIQUID COOLING (WATER AND REFRIGERANT)

Liquid cooling is not new to data centers. Mainframe computers from the past and most current CRAC/CRAH units require chilled water lines to carry away the heat that is generated. The use of convection or air to cool the network equipment provides a limited solution, and some data centers may require additional cooling methods. Liquid cooling provides a huge increase in heat dissipation capacity over air: water is 3,467 times more efficient than air at removing heat. Based upon the data center scenario presented on page 7, the liquid cooling closed cabinet method saves \$286,600 in annual power usage.*

The closer the cooling unit is to the heat source, the more efficient the heat transfer is—and because of this, the less costly it is to operate. The Concept is simple: a water-jacketed heat sink on top of the processor will transfer heat more efficiently than air that rushes past a heat sink and must travel all the way back to a CRAC/CRAH unit. The same Concept applies to super computers that feature the entire board and processor submerged in a liquid.

Liquid cooling systems consist of heat exchangers located inside the cabinet or next to the cabinet. Fans maintain air circulation through the active components, and the heat exchangers provide a chilled liquid that is used for heat transfer. Once the liquid is heated, it is removed from the enclosure. This set-up is referred to as a closed-coupled, or closed-loop, design.

* Note: Liquid cooling allows for higher density (fewer cabinets and a smaller sized data center) which is not calculated in the savings. Liquid cooling is a stand-alone solution not cumulative to the other activities.



LIMITATIONS OF AIR COOLING

The discussion on the limits of air cooling often generates a healthy debate. Certain variables can be manipulated to show extremely high heat dissipation from pure air solutions. In addition, it can be difficult to detect and calculate the effects of air obstructions located either beneath a raised floor or within the data center. Variables such as floor tiles, sealed openings, blanking panels, cutouts and more can also greatly influence the cooling capacity calculations. In general, standard hot aisle/cold aisle designs typically start to reach their limits at approximately 6 to 8 kW per rack. To achieve ratings close to and above 12 kW per rack with air cooling, some sort of physical segregation with a barrier is almost always required to prevent cool supply air from mixing with the warm return air. (See sections on chimney and containment solutions for additional details.) It is important to note that the theoretical limit for air cooling is based more upon the server than the room itself. The server manufacturer must supply the internal thermal management system to cool the CPU through the use of the internal fans within the server.

LIQUID COOLING
REDUCES ENERGY
COST BY 30 - 40%
WHEN COMPARED
TO TRADITIONAL
RAISED FLOOR
DATA CENTER
COOLING

7. CABINET CONFIGURATIONS (CONTINUED)

ADVANTAGES OF LIQUID COOLING

- Greater heat removal capacity
- Acoustic sound abatement
- Liquid cooling eliminates thermal stratification, which occurs as cold air settles at the bottom of the rack and hot air moves towards the top, allowing for increased air intake temperatures up to 80° F—further reducing energy needs
- Functions in non-raised floor environments—either by including the piping in the slab floor or installing it in a plinth base under the cabinet
- Existing data centers typically spread out the location of higher density servers, resulting in the need for more cabinets and floor space; clustering high density heat loads into liquid cooled cabinets will free up floor space within the data center. This approach works well when the existing CRAC/CRAH has been maxed out and the data center manager needs to add additional servers.
- Airflow and cooling needs are not disturbed or increased due to self-contained liquid cooled cabinets
- Liquid is 3,467 times more efficient than air in dissipating heat. For example, an air cooling ¾ HP fan providing 1,000 CFM in a 20" duct equals a liquid cooling 1/10 HP pump delivering 2 GPM in a 1" pipe.
- Higher chilled water temperatures avoid humidification issues, which decrease efficiency and increase energy costs

DISADVANTAGES AND DESIGN CONSIDERATIONS

- Fear of water in the data center
- Higher initial cabinet costs
- Higher initial cabinet costs and ongoing maintenance support

TYPES OF LIQUID COOLING SOLUTIONS

- Stand-alone cabinets
 - Totally enclosed cabinets sealed from the outside
 - Liquid-to-air heat exchangers are located inside the cabinet and produce a large quantity of cold air for cooling purposes
 - Sounds generated by the servers are contained in the sealed enclosure, resulting in an extremely quiet system
 - Can be used for remote network cabinets located outside the data center
- Above-the-cabinet row cooling
 - Heat exchangers are positioned to capture hot exhaust, remove heat via chilled water and/or refrigerant, and deposit cold air into the intake areas of the servers
 - Can be used as a supplemental or primary cooling system
 - Can be used on grade (concrete floor)
 - Ability to retrofit existing data centers
- In-row cooling with containment
 - Cold air is delivered to the cold aisle
 - Can be used as a supplemental or primary cooling system. This solution requires floor space for the in-row cooling unit.
 - When used as a primary system, a raised floor cooling system is not needed

Liquid cooling is most prevalent not as a stand-alone system that is designed upfront to handle all the heat loads in a data center, but as part of a comprehensive cooling system. Many end-users cluster extreme heat loads in a limited area of the data center so remaining areas within the data center have low heat loads, making standard cabinets acceptable for use in these areas. Liquid cooled cabinets often

Humidification,

the process of adding or removing water from the air, consumes energy. While ASHRAE recommends relative humidity of 40 to 55 percent, a higher chilled water temperature often avoids humidification needs.

provide an ideal solution in setups where an existing data center is close to maximum cooling capacity and a water supply line already exists. In these situations, a row of liquid cooled cabinets can be added to accommodate the higher heat loads. The high-density applications can be addressed, the cooling load for the rest of the data center will have decreased and future increases can be handled.



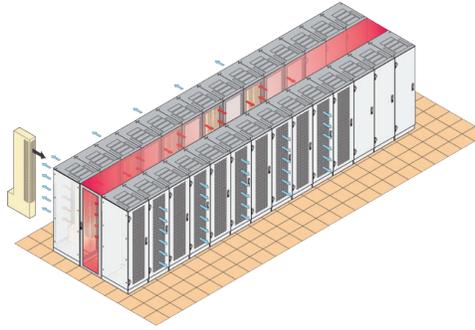
7. CABINET CONFIGURATIONS (CONTINUED)

HOT AISLE/COLD AISLE CONTAINMENT SYSTEMS

Both hot aisle containment (HAC) and cold aisle containment (CAC) systems greatly improve the data center cooling efficiency by segregating cold and hot airstreams, preventing them from intermixing. They can be used on raised floors or on grade, depending on the data center. The use of containment systems can dramatically reduce energy costs, minimize hot spots and improve the carbon footprint of the data center. The basic design principle is to enclose the hot or cold aisle through the use of aisle-way doors, roof panels and internal sealing within the cabinets.

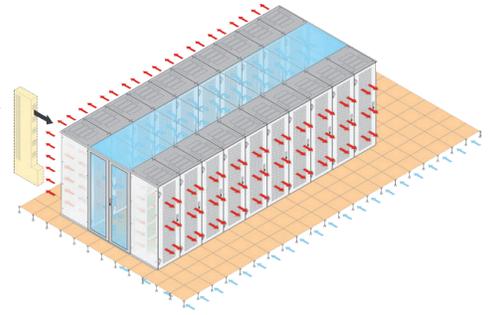
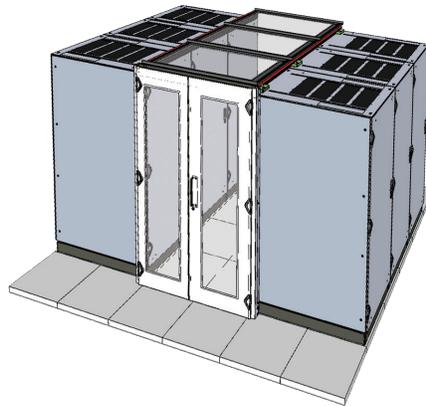
ADVANTAGES OF CONTAINMENT SYSTEMS:

- Cooling systems can be set to a higher temperature and still supply enough cold air to the servers since there is no mixing of the cold supply air and hot exhaust air. Most outside studies estimate that in traditional hot aisle/cold aisle systems, up to 40 percent of the cold supply air does not get used because hot air goes over and around cabinets, reducing the cold air supply.
- Once the containment is accomplished, increased cooling capacity—due to the separation of hot and cold air—makes it possible to cool higher heat loads per rack (e.g. more servers).
- Containment systems are able to control airflows and allow for higher intake operating temperatures, reducing and, in many cases, eliminating the need for humidification. Less humidification allows increased efficiency and lower energy usage.



HOT AISLE CONTAINMENT

In hot aisle containment, the hot airflow is contained and often directed back to the cooling units. The hot exhaust air is either directed back to the CRAC/CRAH units via overhead ductwork or air conditioned locally by in-row primary cooling units. HAC can be used with either raised or non-raised floors. It is becoming common with very high-density data centers on grade that duct the entire hot aisle air above and out the cooling units, which are placed above the roof or on an outside wall. In-row cooling moves the cooling units closer to the heat source, improving overall air handling efficiency since there is less distance to move the air. The scenario on page 7 shows an annual savings of \$175,600 for hot aisle containment.



COLD AISLE CONTAINMENT

Cold aisle containment encloses the cold air within an aisle, ensuring that all equipment uses only that air for cooling. Hot air is excluded from entering the aisle. CAC can be used with or without raised floors and can be easily retrofitted. This setup can be accomplished with external cooling, where cold air is supplied by perimeter CRAC/CRAH units—typically through use of a raised floor—or by internal cooling, a method that features a cooling source located between or above the racks.

Results from several outside studies greatly support this approach. A 2007 Pacific Gas and Electric study indicated that a 20 percent savings in chiller operating costs is possible with the use of a CAC system over traditional hot aisle/cold aisle designs, and Lawrence Berkeley National Laboratory found that electricity used to move the cold air supply could be reduced by 75 percent.¹ Additionally, several independent studies have shown cooling-related energy costs can be reduced by as much as 30 percent with a CAC system. Based upon the data center scenario presented on page 7, cold aisle containment provides an annual savings of \$135,600.

¹Gartner, 2008. McGuckin, Paul. Cool More With Less in Your Data Center.

7. CABINET CONFIGURATIONS (CONTINUED)

	Hot Aisle Containment		Cold Aisle Containment	
	External cooling – perimeter CRAC unit	Internal cooling – unit located in the containment structure	External cooling – perimeter CRAC unit	Internal cooling – unit located in containment structure
Focused cooling solution	Open cold air distribution – subject to mixing and disturbances	Open cold air distribution – subject to mixing and disturbances	Focused cooling to the server intake	Focused cooling to the server intake
Energy efficiency	Improvement over traditional HACA solutions	Improvement over traditional HACA solutions– cooling located closer to the source	Typically improves energy efficiency by 30%	Best solution for energy efficiency – can improve overall efficiency by 30 to 40%
Retrofitting requirements	Ceiling panels, doors and complete ducting to the CRAC/CRAH unit intake required	Most difficult of the four solutions to retrofit	Ceiling panels and doors required. Simplest of the four solutions to retrofit.	Ceiling panel, doors and cooling units provided.
Ability to handle high density applications	Depends on layout – typically 10-15 kW per rack is possible	More than 30 kW per rack is possible	Depends on layout – typically 10-15 kW per rack is possible	More than 30 kW per rack is possible
Use of water	No water required	Requires water piping to the heat exchanger through the data center	No water required	Requires water or refrigerant piping to the cooling unit
Floor space requirement	No additional space required	Additional floor space required for side-mounted cabinet units	No additional space required	Additional floor space required for side-mounted cabinet units
Thermal ride-through	Limited due to decreased available air space	Limited due to decreased available air space	Limited due to decreased available air space	Limited due to decreased available air space
Configuration flexibility	Limited service to rear side of panel due to higher heat environment	Limited service to rear side of panel due to higher heat environment	No limitations	No limitations

ADVANTAGES OF HAC AND CAC SYSTEMS

- HAC can increase cooling unit efficiency as hot air is ducted directly to the CRAC/CRAH unit
- HAC and CAC reduce energy costs for a cost-effective, green cooling solution
- Higher equipment densities (higher kW load per cabinet) are possible, leaving more floor space open for future expansion
- Provides a future-proof data center design; eliminates the need to purchase additional CRAC/CRAH units and avoids expensive expansion requirements
- Acoustic sound abatement

DISADVANTAGES OF HAC AND CAC SYSTEMS

- HAC requires complete hot air ducting back to the air inlet of the CRAC/CRAH
- HAC requires overhead space for the air plenum/ductwork, which may create future retrofit design challenges
- CAC systems with additional internal cooling consume valuable floor space for the additional side-mounted cooling units

- CAC requires that the enclosed area be provided enough cold air CFM to fulfill equipment requirements

DESIGN CONSIDERATIONS

- Fire suppression/sprinkler placement requirements, which vary by local jurisdiction; plastic curtains are growing in popularity, as they segregate the airflow and in the occurrence of a fire melt at 130° F
- Aesthetics and accessibility, as aisle containment designs challenge the traditional images of a data center
- Environment for personnel: within the CAC, the sound level will be higher; within the HAC, the temperatures will be significantly higher
- Lighting and access control requirements
- The data center cooling failure run time is impacted by the type of containment system that is deployed. HAC systems draw from large volumes of cold air, allowing significant run time while the cooling systems are diagnosed and repaired. CAC systems have less cold air within their containment area and, once depleted, will cause network equipment to overheat.

Containment systems

allow for higher cold temperatures per ASHRAE and ΔT , optimizing cooling systems and lowering operating costs.

CONTAINMENT SYSTEMS CAN INCREASE ENERGY EFFICIENCY BY AS MUCH AS 70 PERCENT OVER TRADITIONAL HOT AISLE/COLD AISLE SYSTEMS

8. ADDITIONAL COOLING SOLUTIONS

HEAT EXCHANGER DOORS

Heat exchanger doors at the rear of server cabinets cool the hot exhaust as the air stream moves out of the cabinet into the hot aisle. These are considered a supplementary cooling measure and can improve the capacity of the cabinet by accommodating more equipment—and therefore, more heat. The doors have liquid cooling lines that carry away the heat absorbed from the hot exhaust airflow. However, the flexible joints at the doors' hinge points cause concern for potential leaks. This solution is typically a very large and difficult undertaking for providing additional cooling.

IN-ROW COOLING

In-row cooling utilizes cooling units—placed between cabinets—that capture hot aisle exhaust air and neutralize it before it mixes with cold air. This air then cools the network equipment in an arrangement that allows the cycle to repeat itself. The cooling units use either compressor-based cooling or liquid “chilled water” cooling to remove the heat and direct it outside of the room. This solution provides very high heat dissipation, and when used in conjunction with containment systems, is commonly capable of at least 20 kW to 30 kW of cooling per cabinet.

ADVANTAGES

- High heat dissipation capabilities
- Additional cooling is possible if CRAC/CRAH unit is fully utilized
- Supplemental cooling technique for high density applications

DISADVANTAGES

- Reduced reliability due to a number of potential failure points
- Potentially higher maintenance costs than traditional CRAC/CRAH data center designs
- Access to a cooling unit is limited due to its location between cabinets
- Units consume valuable floor space, typically consuming a third of the width of a cabinet
- To avoid false cooling readings, in-row cooling units must be placed directly across from one another; optimal results are achieved by establishing equal rows

CEILING-MOUNTED COOLING UNITS

Ceiling-mounted cooling units are placed above the cold aisle to deliver cold air. These units utilize compressor-based, liquid Freon, or liquid chilled water cooling systems. Because cold air is denser than hot air, it sinks into the cold aisle. The intake side of the cooling units faces the hot aisle and repeatedly draws in hot air to be cooled. These units can be used on raised floors or on grade. They can serve as a supplemental or primary cooling system depending on the quantity of units used. The major drawbacks of this solution are high upfront capital costs and reduced reliability. Liquid cooling lines above equipment cabinets are a concern because of potential fluid leaks. A high number of connection points must be monitored, and much preventative maintenance is needed. Thus, operational costs can be much higher than traditional CRAC/CRAH units.



REMOTE APPLICATIONS

Applications outside the data center rely on some of the same equipment and cooling methods used in data centers. The basic function of these cabinets is to provide a data center within the cabinet.

AC stand-alone cabinets feature a dedicated air conditioner mounted to the cabinet to provide cooling for the remote application. The cabinet is completely closed to all outside environments—no outside air is allowed to enter. The AC reconditions hot exhaust air by transferring the heat outside the cabinet.

Filter fan cabinets are used when the air temperature in an environment is sufficient to cool the equipment but the air needs to be filtered, such as in a dusty environment. The top-mount fans direct hot exhaust out of the cabinet, while a lower filter intakes outside air, removing contaminants and drawing cold air from the floor. As previously noted, fans can only cool to the surrounding air temperature regardless of the amount of airflow. This solution is ideal for warehouses, factory floors, schools, etc.

8. ADDITIONAL COOLING SOLUTIONS

NETWORK SWITCHES

Not all data center equipment employs front-to-back airflow. Some equipment utilizes alternate airflow designs that can be accommodated through one of several techniques.

SPECIALIZED NETWORK EQUIPMENT

Some specialized network equipment requires unique cold and hot air paths. Devices like CISCO 6509 & 6513 network core switches utilize a right-to-left airflow. To be used alongside servers in hot aisle/cold aisle data centers, this equipment requires specialized ductwork that provides cold air to the equipment, as well as a pathway for the hot exhaust to exit. This technique ensures equipment such as network core switches may utilize the same cooling strategies already in place for servers.

Without this ductwork, the hot exhaust from one switch will be directed toward the air intake of the adjacent switch or server cabinet—typically resulting in overheating. The optimal system design not only addresses the switch's unusual airflow,

but it also manages the high-density cabling used in these applications.

SWITCH CABINETS

Relatively new to data centers, switch cabinets can be used to house and support large network core switches, patch panels and cables. Their rapid rise in importance is the direct result of server virtualization, which allows servers to be pooled together and used much more efficiently. Important factors to consider include providing enough heat dissipation, a high cable capacity for high-density cables and proper cable management.

TYPICAL SWITCH CABINET SPECIFICATIONS

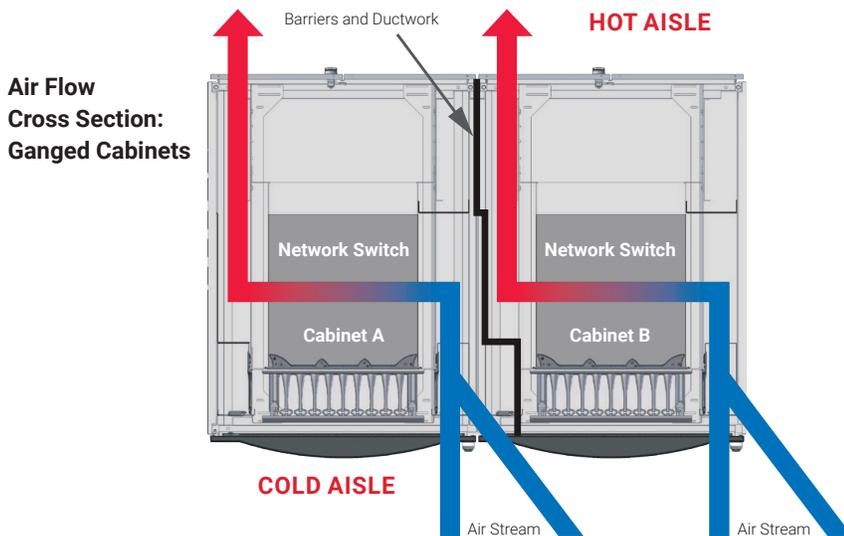
- Widths: 800 mm (32"), meets CISCO operating requirements
- Height: 2000 mm (84"), 42-45 RU
- Depth: 1000 mm (39.4"), deeper is better

Avoid cabinet widths of less than 32." CISCO requires six inches of clearance between the sides of 6500 Series Switches and the sides of a cabinet.



IMPORTANT CABINET ATTRIBUTES FOR NETWORK SWITCHES

- Individual cable management provided at the front of the cabinet plus management of cable bundles at the rear of the cabinet
- Front air dams and blanking panels used to prevent hot exhaust recirculation
- Installing and stacking two large network core switches into a single cabinet should be avoided due to the large volume of cables that must be routed to adjacent cabinets as well as the need to provide the required cold air to the switch intakes
- Cable management featured at both sides of the switch, allowing cable to be routed in two paths if desired; also facilitates one-path routing in instances where the end-user must ensure the fan tray is accessible (in case of switch failure)
- Some data centers employ open frame racks to support network switches. While they provide easy cable access and management, they inadequately address thermal management issues, which leads to poor reliability and short equipment life.



9. SUMMARY AND SELECTION GUIDE

DATA CENTER CABINET BEST PRACTICES

- All unused RU spaces must be filled. Blanking panels can be used to fill empty spaces and increase airflow efficiency.
- Use side barriers to prevent unwanted airflow, including heat and cold air leakage to adjacent cabinets.
- Use brushes or other airflow-blocking cable egress devices to prevent unwanted hot and cold airflow through cable openings.
- Doors should meet equipment manufacturers' open space requirements to avoid warranty issues. Open space above 60 percent meets most manufacturers' requirements.
- Split rear doors improve cabinet-to-cabinet equipment access and provide space for walking around equipment.
- Mount servers and other equipment not less than 4 inches from the inside of the front door to allow sufficient airflow and space for equipment handles.
- Route cables neatly away from incoming and existing airflow to mitigate interference.
- Servers usually require EIA Universal Spaced Square Holes (.375" x .375").
- Order cabinets with the maximum depth available to provide necessary space for PDUs, cables and equipment. Avoid depths less than 1000 mm (39").
- Install fans in locations where they do not pull cold air into heated airstreams.
- Mount PDUs in available vertical and horizontal spaces.
- Install equipment to promote shortest cable routes. Avoid large slack bundles which adsorb heat and block airflow.
- Avoid cable arms that block rear exhaust airflow.
- Utilize the equation $CFM = Watts / (.316 \times \Delta T F)$ for proper cooling solution sizing

DATA CENTER STRATEGIES - COOLING

Strategy	Capital Costs	Raised Floor	Operating Costs/kW	Average Max kW per Cabinet	Reliability (Complexity)
Random (chaos)	\$	Yes	Very High	2,500 kW	Very Low
Hot aisle / cold aisle	\$	Yes	High	6 kW – 8 kW	Low
Chimney	\$\$	Yes or No	Moderate	10 kW – 20 kW	Medium
Containment (raised floor)	\$\$\$	Yes	Moderate	15 kW – 20 kW	Medium / High
Closed-loop, ducting, hot and cold containment	\$\$\$\$	Yes or No	Low	10 kW – 20 kW	Medium
In-row cooling & containment (heat exchanger)	\$\$\$\$\$	No	High	10 kW – 20 kW	High
Liquid cooling (heat exchanger)	\$\$\$\$	No	Low	20 kW – 40 kW	High
Chip-level liquid cooling	\$\$\$\$\$\$	No	Very Low	+40 kW	Very High

DATA CENTER STRATEGIES - RAISED FLOOR

Raised Floor Cooling	Percent of Cold Air Wasted	Operating Costs / kW	Reliability (Complexity)	Upfront Capital Costs
Random (chaos)	40%	High	Low	Low
Hot aisle / cold aisle	25% +	High	Low / Medium	Low / Medium
Chimney	3%	Low	Medium	Medium
Containment	1%	Low	Medium / High	Medium
In-row cooling & containment	1%	Med	Medium / High	High

DATA CENTER STRATEGIES - ON GRADE (CEMENT SLAB)

On Grade (Slab) Strategies	Percent of Cold Air Wasted	Operating Costs / kW	Reliability (Complexity)	Upfront Capital Costs Comparison
Above cabinet cooling units	20%	High	Low	Low
Hot aisle ducting/outside cooling units	5%	Low	Medium	Very High
Chimney	3%	Low	Medium	Medium
In-row cooling	1%	Low	Medium / High	Medium / High



nVent.com

Our powerful portfolio of brands:

CADDY ERICO HOFFMAN RAYCHEM SCHROFF TRACER