

How Row-based Data Center Cooling Works

White Paper 208

Revision 0

by Paul Lin and Victor Avelar

Executive summary

Row-based data center cooling is normally regarded as a “cold air supply” architecture that uses row-based coolers. However, row-based cooling is actually a “hot air capture” architecture that neutralizes hot air from IT equipment before it has a chance to mix with the surrounding air in the room. This paper discusses common misconceptions about row-based cooling, explains how row-based cooling removes hot air, and describes key design attributes that maximize the effectiveness of this approach.

Introduction

What is a data center “pod”?

A data center pod is a cluster of IT cabinets combined with power and cooling infrastructure that is deployed as a unit. Rooms are planned in advance for a number of pods, but the pods may be separately deployed or upgraded over time. Pods typically assembled on-site in a room to a standard design, but may be partially or extensively pre-manufactured.

In its most common form, a pod is a pair of rows of cabinets sharing a hot aisle.

Pod-based design is a recommended best-practice for larger data centers.

Close-coupled cooling, row-oriented architecture, and in-row cooling are some common names used to describe row-based cooling in data centers. Throughout this paper the term row cooler¹ will be used to describe this architecture. In addition, the term “pod” will be used to describe how row coolers are typically deployed (see side bar for definition).

Although row-based cooling technology has been around for about 10 years, there is still confusion surrounding how this type of cooling works. Specifically there are three common misconceptions about how it works and how it is deployed:

- **Misconception 1: Row coolers require turning vanes:** It is commonly believed that row coolers require turning vanes to direct cold air to the front of the racks. In fact, some data center designers, operators, and manufacturers design turning vanes at the outlet of row-based cooling equipment to try to blow cold air directly into the racks.
- **Misconception 2: Row coolers are needed in every row:** It is commonly believed that every row in the data center requires row cooler(s). It is counter intuitive to think that you can cool an entire pod if all row coolers are in one row and none in the opposing row.
- **Misconception 3: Row coolers can’t cool loads outside of their pod:** It is commonly believed that row coolers can only cool the loads inside the pod they are located in. Similarly, it is often thought that extra row-based cooling capacity in one pod cannot help cool another pod.

This paper uses scientific evidence and the core principle of hot air capture to explain why these are misconceptions. The paper then describes three key row-based cooling design attributes that maximize hot air capture. For more information on fundamentals of row-based cooling, see White Paper 130, [Choosing between Room, Row, and Rack-based Cooling for Data Centers](#).

The principle of hot air capture

The goal of any data center cooling system is to remove the heat added to the air by the IT equipment. A metric for measuring the effectiveness of how hot IT exhaust air is captured by cooling equipment (or how cold air is supplied to IT inlets) is called the capture index². Capture index is based solely on the airflow patterns associated with the supply of cold air to, or the removal of hot air from a rack. The capture index is typically a rack-by-rack metric and has values between zero and 100%; higher values generally imply good cooling performance and scalability of a cooling architecture.

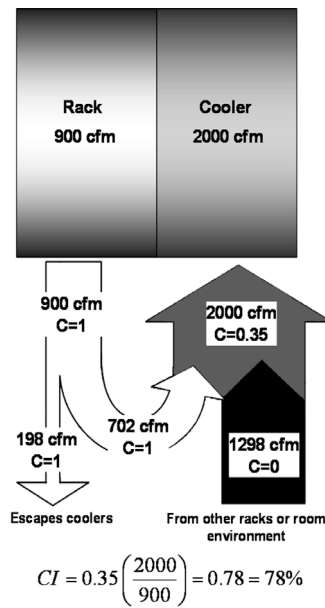
The hot air capture index is the fraction of hot air exhausted by IT equipment that is captured directly by the row coolers **within the same pod as the IT equipment**. This is the principle metric that quantifies the effectiveness of row-based cooling. This is illustrated in **Figure 1** where 78% of the hot IT exhaust air is captured by the cooler. **Figure 2** shows a CFD simulation result of temperature contour of a data center with row coolers.

¹ Note that the principles in this paper are described with the assumption that the row coolers are floor-mounted as opposed to ceiling-mounted.

² J.W. VanGilder, S.K. Shrivastava, Capture Index: An Airflow-Based Rack Cooling Performance Metric, Published in ASHRAE Transactions 2007, Volume 113, Part 1.

Figure 1

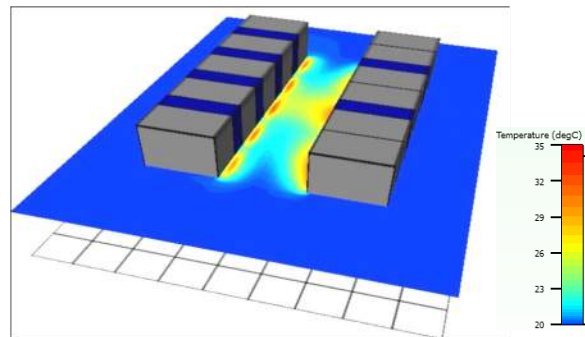
Illustration of hot air capture index



Source: J.W. VanGilder, S.K. Shrivastava, Capture Index: An Airflow-Based Rack Cooling Performance Metric

Figure 2

CFD simulation illustrating hot air capture (center aisle represents hot aisle, and blue objects represents row coolers)



A useful application of hot air capture index is optimizing the placement of racks and cooling units in a pod. The design goal is to ensure that all exhaust air is captured by the cooling units so that there is no net heating of the room. In this case, rack-by-rack capture indices explicitly show how much of each rack's airflow is captured so that the appropriate design changes may be implemented until an acceptable solution is found. The hot air capture index design target should generally be at least 90%, which means the pod layout is effective. The following design practices can help attain or exceed this value.

- **Use blanking panels and brush strips** – IT exhaust air can escape the hot aisle through open unused “U” positions in IT racks and through open areas around cable penetrations at the top and bottom of the rack. Installing blanking panels in open “U” positions and brush strips around cable penetrations is very effective at increasing the hot air capture index. For more information on airflow management with blanking panels, see White Paper 44, [Improving Rack Cooling Performance Using Airflow Management Blanking Panels](#).
- **Optimize row cooler placement** – Row coolers should be placed in between IT racks and generally not at the end of a row in order to maximize hot air capture. For uniform rack loading, the coolers should be distributed evenly along the length of the hot aisle. For non-uniform rack loading, coolers should be clustered closer to the higher loads.
- **Use side air distribution units** – Many types of switches and routers employ side-to-side airflow. This airflow configuration is generally incompatible with typical front-to-back airflow of IT enclosures. Side air distribution units allow switch and router enclosures to be placed side by side and still maintain proper airflow thereby consuming less

floor space. See White Paper 50, [Cooling Solutions for Rack Equipment with Side-to-Side Airflow](#).

- **Avoid fan trays** – Fan trays or roof fans employed in racks can actually degrade air management as the induced airflow patterns are incompatible with the front-to-back IT airflow.
- **Use deeper IT racks** – Over the last 10 years, IT equipment has trended toward deeper form factors. Using shallow-depth racks (i.e. 900mm / 35in) with this IT equipment can constrict the movement of hot exhaust air out from the back of the rack. Deeper racks allow more room for cabling which decreases resistance to airflow.
- **Use air containment** – The further row coolers are from IT racks, the more difficult it is to prevent IT exhaust air from recirculating back into IT equipment inlets. Containing the hot aisle with row coolers can improve cooling system efficiency, especially at densities below 3kW/rack. For more information on airflow management and containment please read these two white papers: White Paper 135, [Impact of Hot and Cold Aisle Containment on Data Center Temperature and Efficiency](#); White Paper 153, [Implementing Hot and Cold Air Containment in Existing Data Centers](#).

Debunking the misconceptions

The three misconceptions mentioned in the introduction are indicative of a lack of understanding in how row coolers contribute to the concept of hot air capture. This section will first describe three key design attributes of row coolers, which will then be used as the basis for debunking each misconception.

Design attributes of row coolers

Row coolers capture hot IT exhaust air and neutralize it before it has a chance to mix with the surrounding air in the room or recirculate to the front of the IT rack. Capturing 100% of hot air can improve energy efficiency and eliminate hot spots. There are three key design attributes that help row coolers capture and neutralize hot air:

- Back-to-front airflow
- Rack-based footprint
- Variable cooling capacity devices

Back-to-front airflow

Row coolers are designed with back-to-front airflow using small fans distributed across the height of the cooler. Since the rear of the row cooler is in the hot aisle, the fans are able to collect the hot IT exhaust air directly and uniformly (top to bottom) from the hot aisle as shown in **Figure 3**. It has been shown, based on computational fluid dynamic (CFD) analysis and performance testing, that the hot air capture index is generally highest for IT racks closest to the row cooler. Conversely, the hot air capture index decreases for IT racks further away from the row cooler. This relationship brings about the benefit of the second design attribute, rack-based footprint.



Figure 3

Row cooler showing vertically-aligned fans and back-to-front airflow

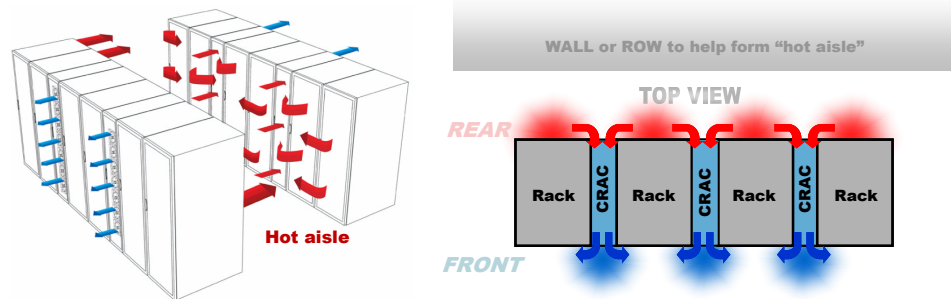
Rack-based footprint

Row coolers are designed with a footprint similar to an IT rack (full or half rack width). This design attribute allows you to distribute row coolers throughout a pod of IT racks. The maximum distance between any row cooler and IT equipment exhaust is normally within 3 meters (10 feet). The implication of this distributed cooling layout is that the hot air capture index is improved for the entire pod. When the hot exhaust air of one IT rack is beyond the “reach” of a distant row cooler, a closer row cooler will capture the majority of the hot exhaust air in that part of the pod.

If all of the hot exhaust air in a pod is captured by distributed coolers in that pod (i.e. 100% hot air capture), then by definition there can be no hot spots at the intakes of IT equipment. **Figure 4** shows an example layout of distributed row coolers. The quantity and locations of row coolers are determined according to rack densities and the extent of containment.

Figure 4

Illustration of row coolers distributed across a row of racks

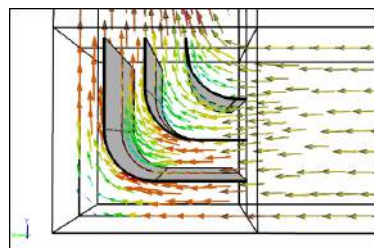


Variable cooling capacity devices

Row coolers are designed with EC fans which can adjust fan speed to change airflow and output cooling capacity. Traditional CRAHs typically only control airflow to maintain raised floor static pressure, and is not related to the IT load. This design attribute allows row coolers to balance the cooling capacity with the heat load by sensing the inlet temperature of nearby racks or the IT room.

Correcting misconception 1: Row coolers require turning vanes

If one doesn't understand the premise of hot air capture, it's easy to conclude that turning vanes are required to blow cold air in the direction of the IT equipment. However, both **design attributes described above show that turning vanes are not required**. If 100% of the hot IT exhaust air is captured and cooled, (before it has a chance to mix with surrounding air in the room), the remaining space in the room becomes a cold air supply plenum³. Therefore, it doesn't matter where the cold air (from the row cooler) goes; it just matters that the hot air (from the IT equipment) is being captured and neutralized.



CFD simulation of airflow leaving the turning vanes (Air velocity vectors show)

Turning vanes are used by some manufacturers to direct cold air into the adjacent racks, which indicates they don't understand hot air capture. The use of turning vanes not only increases the capital cost and increases the width of the cold aisle, but also results in airflow issues. The airflow leaving the turning vanes is at a higher velocity, and perpendicular to, the airflow of the adjacent IT rack. The high velocity jet flow creates regions of low pressure in front of the IT rack, interfering with rack airflow. (This is known as Bernoulli's principle,

³ Even if the initial room air temperature is higher than the row cooler supply air temperature, over time, the room reaches steady state and will be approximately the same temperature as the row cooler supply air. This assumes relatively tight room construction with minimal infiltration.

the same principle that allows wings to lift an airplane.) Unlike the gentle and fairly uniform “sucking” of hot IT exhaust in the hot aisle, the strong jets blanketing the rack inlets creates highly variable conditions from rack to rack. (Jet flow is characterized by localized pressure and velocity changes.) Note that turning vanes also have a significant pressure drop requiring more fan power.

Correcting misconception 2: Row coolers are needed in every row

Many believe that row coolers can only cool the racks which are located in the same row as the coolers. Both design attributes described above show that, with a high hot air capture index, it doesn't matter if all the row coolers are on one side of a two-row pod (or any placement combination).

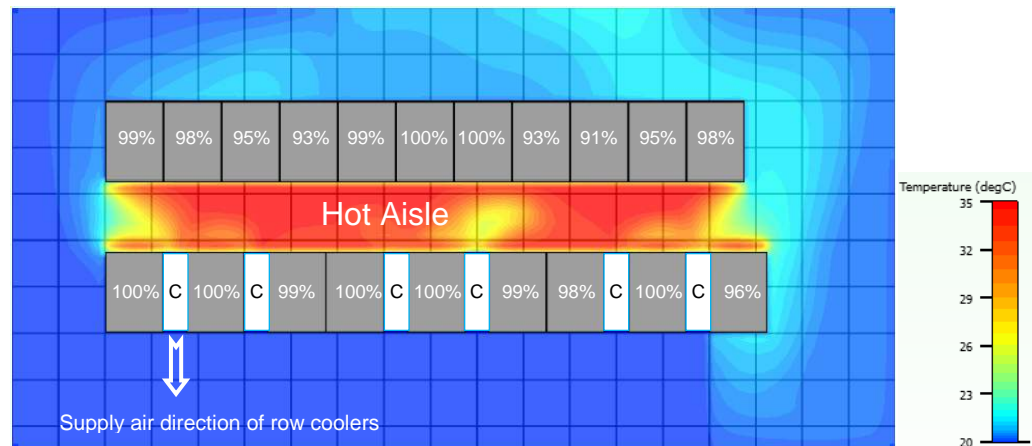
CFD modeling and actual installations show it is possible for row coolers in a single row to cool racks in both rows as illustrated in **Figure 5**. The percentage values on each rack represent the hot air capture index. Note in the figure, all the IT inlets are exposed to the blue color which is neutralized air, and the red color variation in the middle is the hot aisle.

Figure 5

Example CFD analysis with all row coolers in one row

3D view looking at the top of the racks

Average 5kW/rack with no cooling redundancy or containment



Correcting misconception 3: Row coolers can't cool loads outside their pod

This misconception has been communicated in other ways including:

- Row coolers are for spot cooling pods only
- Row coolers can't cool racks in another pod
- Row coolers can't cool a large room
- Row coolers can't cool tape storage units on the perimeter of my data center
- N+1 cooling redundancy with row cooling means that you need a redundant cooling unit in every pod

An important fact to understand about this misconception is that the hot air capture index measures the percentage of a rack's hot exhaust air that is captured directly by row cooler(s) in the same pod as the rack. By definition, if 100% of a rack's hot air were to make its way back to a row cooler far outside its pod, its hot air capture index would be 0%.

Therefore, the most predictable way to cool stand-alone ancillary IT gear, such as storage arrays, is to place a row cooler next to the gear, assuring a high hot air capture index. In essence, this creates a miniature pod with up to 100% hot air capture. This is analogous to placing a perforated tile in front of the stand-alone gear, in a raised floor cooling architecture.

The main difference is that a perforated tile operates on the premise of cold air capture which in this case is limited by how much airflow (and kW cooling capacity) is provided as discussed in WP121, [Airflow Uniformity Through Perforated Tiles in a Raised-Floor Data Center](#).

What if the stand-alone IT gear doesn't have a row cooler (i.e. in its pod)? How would this gear be cooled in a row-based cooling architecture? The answer is tied to the "rack-based footprint" and "variable cooling capacity devices" design attributes of row coolers. The rack-based footprint allows coolers to be distributed across the entire data center which creates a close-coupled cooling effect with all heat loads. The variable cooling capacity devices oversupply cold air to the cold aisle after sensing an increased room temperature. Containing the hot aisle with row coolers can also improve hot air capture, especially at densities below 3kW/rack. Will containment have a negative impact on cooling stand-alone IT gear? The easiest way to illustrate this is by comparing two identical data center layouts each with 65 racks, 235kW of IT load, two stand-alone IT devices, with and without containment.

Figure 6A shows the results of a CFD analysis with 14 row coolers distributed evenly throughout the data center. The density for the 20 wide racks is 5kW/rack and the 43 narrow racks are 3kW/rack. The two ancillary racks in the upper right-hand corner represent stand-alone IT gear. The hot air capture index value (with the unit of percentage) is shown on each rack. Both the 5kW stand-alone and 1kW stand-alone devices have an average inlet air temperature of 22°C (71.6°F). **Figure 6B** shows the results of a CFD analysis of the same data center except with hot aisle containment to illustrate its effect on hot air capture and inlet air temperatures of the stand-alone IT devices. Most of the racks have 100% hot air capture index. Both stand-alone IT devices have an average inlet air temperature of 21°C (69.8°F).

Figure 6A

Example CFD analysis with distributed row coolers

3D view looking at the top of the racks

235kW total load
65 total racks
Wide racks = 5kW/rack
Narrow racks = 3 kW/rack
No cooling redundancy or containment

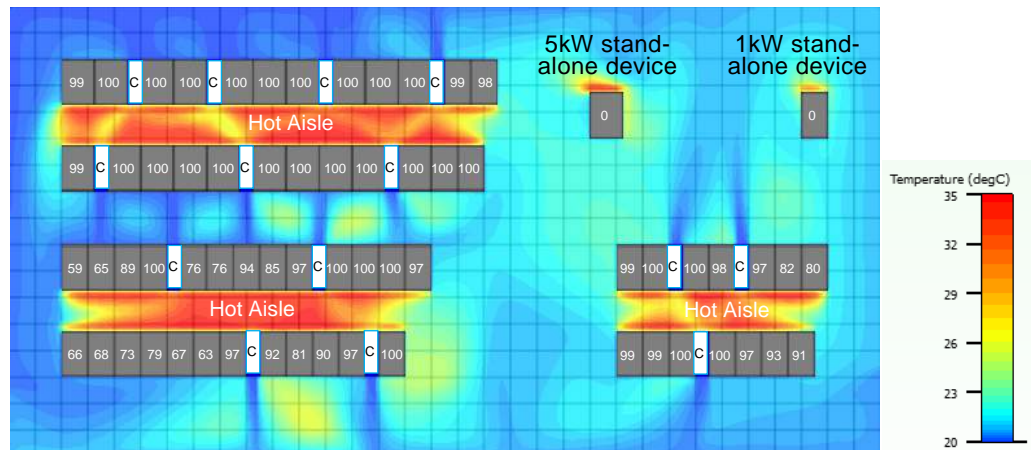
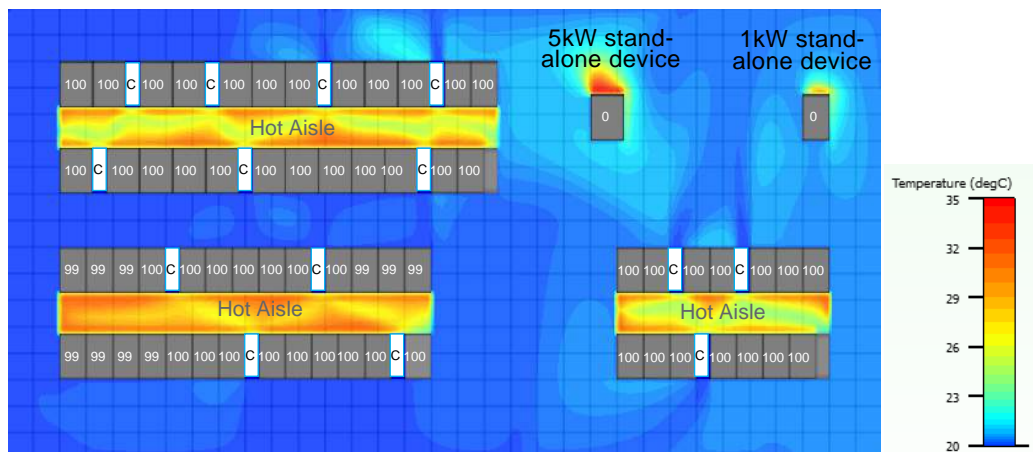


Figure 6B

Example CFD analysis with centralized row coolers

3D view looking at the top of the racks

235kW total load
65 total racks
Wide racks = 5kW/rack
Narrow racks = 3 kW/rack
No cooling redundancy but with containment



The figures above help to visualize the effect of row coolers on hot air capture index values and cooling stand-alone IT device. As row coolers are variable capacity devices that are capable of over-supplying cold air to the cold aisle. When the system is in balance, the supply temperature from the cooler is equal to the IT inlet temperature. When ancillary equipment (i.e. stand-alone devices) is initially added outside of the pod, the temperature of the overall room environment increases due to air mixing. Nearby row-based cooling units sense this increase and respond by increasing cooling capacity thereby neutralizing the hot air. This increased cooling capacity is achieved by increased air flow (row coolers increase fan speed) and or lower supply air temperature (row cooler increase chilled water valve flow). These adjustments occur initially but remain unchanged once the data center reaches steady state.

The steps describing how the row-based coolers “pick up the ancillary load” are the same even when containment is used. In either case we expect the ancillary equipment to operate at higher temperatures than the row-based racks, because it is not directly coupled to row coolers in their own pod. The further away stand-alone IT gear is from row coolers, the higher their inlet air temperature will be. Should you be concerned about the elevated ancillary equipment temperature? Often times IT equipment can tolerate ASHRAE’s recommended maximum temperature of 27°C (80.6°F). If this is not acceptable, it may be appropriate to place row cooler close to the ancillary equipment rack. For more information on cooling loads outside of a pod, see White Paper 139, [Cooling Entire Data Centers Using Only Row Cooling](#) and White Paper 134, [Deploying High-Density Pods in a Low-Density Data Center](#).

Conclusion

Row-based cooling was designed with a focus on capturing hot exhaust air from IT equipment in a pod. The higher the hot air capture index values, the less hot air recirculation that occurs, which decreases temperature variations across racks (such as hot spots). It doesn't matter where the cold air goes; it matters where the hot air comes from. Row-based cooling units can also help to cool ancillary equipment or racks outside of their pod but hot air capture index is not used as a performance metric in this case. Instead row coolers use a close-coupled cooling process described in this paper. Understanding these principles helps to explain the three misconceptions described in this paper and provides a foundation for effective deployment of row-based cooling.











About the author

Paul Lin is a Senior Research Analyst at Schneider Electric's Data Center Science Center. He is responsible for data center design and operation research, and consults with clients on risk assessment and design practices to optimize the availability and efficiency of their data center environment. Before joining Schneider Electric, Paul worked as the R&D Project Leader in LG Electronics for several years. He is now designated as a "Data Center Certified Associate", an internationally recognized validation of the knowledge and skills required for a data center professional. He is also a registered HVAC professional engineer. Paul holds a master's degree in mechanical engineering from Jilin University with a background in HVAC and Thermodynamic Engineering.

Victor Avelar is the Director of Schneider Electric's Data Center Science Center. He is responsible for data center design and operations research, and consults with clients on risk assessment and design practices to optimize the availability and efficiency of their data center environments. Victor holds a bachelor's degree in mechanical engineering from Rensselaer Polytechnic Institute and an MBA from Babson College. He is a member of AFCOM and the American Society for Quality.



-  [Choosing between Room, Row, and Rack-based Cooling for Data Centers](#)
White Paper 130
-  [Improving Rack Cooling Performance Using Airflow Management Blanking Panel](#)
White Paper 44
-  [Cooling Solutions for Rack Equipment with Side-to-Side Airflow](#)
White Paper 50
-  [Impact of Hot and Cold Aisle Containment on Data Center Temperature and Efficiency](#)
White Paper 135
-  [Implementing Hot and Cold Aisle Containment in Existing Data Centers](#)
White Paper 153
-  [Airflow Uniformity Through Perforated Tiles in a Raised-Floor Data Center](#)
White Paper 121
-  [Cooling Entire Data Centers Using Only Row Cooling](#)
White Paper 139
-  [Deploying High-Density Pods in a Low-Density Data Center](#)
White Paper 134

 [Browse all white papers](#)
whitepapers.apc.com

 [Browse all TradeOff Tools™](#)
tools.apc.com

Contact us

For feedback and comments about the content of this white paper:

Data Center Science Center
dcsc@schneider-electric.com

If you are a customer and have questions specific to your data center project:

Contact your Schneider Electric representative at
www.apc.com/support/contact/index.cfm