Thought Leadership
White Paper

Rethinking Data Center Design

Building the room around the rack can ease cooling challenges and save money

By John Collins
Segment Manager, Data Centers
Eaton Corporation

Executive summary

For decades, most organizations have architected their data centers from the outside in. That is, they've started by designing a building and then figured out how to position and cool the server racks inside that building. Increasingly, however, that time-honored design process is proving incapable of efficiently coping with steadily climbing power densities and data center operating temperatures.

As a result, companies are slowly beginning to embrace a radical new approach to data center design in which they plan new facilities from the inside out, by selecting rack and cooling technologies first and then designing a structure around them.

This white paper discusses the issues making data center cooling more challenging than ever and explains why starting the data center design process with the racks rather than the room can help businesses reduce energy consumption, increase capacity and optimize floor space utilization.

Table of contents

Data center cooling challenges

Thanks to a variety of widely adopted new technologies, including the following, keeping data centers at safe operating temperatures has never been a more demanding task:

- **Virtualization:** Under virtualization, a single physical machine supports multiple virtual machines, each with its own operating system and applications. That, in turn, can lower hardware procurement and maintenance costs by enabling organizations to consolidate their underutilized servers onto a smaller number of more powerful host devices. However, a more thoroughly utilized server is also a hotter one. In addition, some organizations that use virtualization also use management software that dynamically relocates virtual machines onto new host servers based on load balancing requirements. In such facilities, the most heavily utilized—and hence hottest—portions of a data center may be different day to day and even hour to hour.
- **Blade servers:** Often used in conjunction with virtualization, blade servers are plug-and-play processing units with shared power feeds, power supplies, fans, cabling and storage. By compressing large amounts of computing capacity into small amounts of space, blade servers can dramatically reduce data center floor space requirements. Unfortunately, they can also raise power densities above 15 kW per rack, dramatically increasing data center heat levels.
- **Multicore processors:** Today's multicore server processors are significantly more powerful than their single-core predecessors—and far hungrier for electricity too. In fact, servers equipped with dual- or quad-core processors consume 12 to 16 kW per rack on average, resulting in significantly higher server rack operating temperatures.
- **Cloud computing:** Seeking to lower overhead and improve efficiency, businesses are adopting cloud computing solutions in ever growing numbers. Analyst firm IDC expects global spending on public cloud solutions, which deliver applications and infrastructure resources via the Internet, to jump from $21.5 billion in 2010 to $72.9 billion in 2015. Private cloud infrastructures, which employ the same basic technologies as public clouds but reside behind an individual organization's firewall, are rapidly gaining traction in corporate data centers too. Both types of cloud, however, typically utilize large numbers of commodity-priced servers, markedly increasing both power and cooling requirements.

To further complicate matters, today's data center designers must also account for these additional non-technology considerations:

- Environmental regulations aimed at lowering data center power usage and greenhouse gas emissions, such as the U.K.'s Carbon Reduction Commitment legislation, are growing more and more common, making energy efficiency an even higher priority than usual.
- Fueled by increased adoption of environmentally friendly but expensive power sources such as wind and solar, aging electrical transmission systems and utility deregulation initiatives, energy costs are on the rise, further contributing to demand for greater efficiency.
- Regulatory bodies such as the U.S. Environmental Protection Agency, along with numerous electrical utilities, are offering financial rebates to data centers that achieve significant energy savings.
- The high cost of building and operating data centers has companies seeking out ways to shrink the physical footprint of their computing facilities without compromising their ability to meet business requirements.

The evolution of data center cooling strategies

Driven by rising power densities and heat levels, data center cooling strategies have changed significantly in recent years.

Chaos air distribution strategies

Until relatively recently, most data center cooling schemes relied on so-called "chaos" air distribution methodologies, in which computer room air conditioning (CRAC) units around the perimeter of the server room pumped out massive volumes of chilled air that both cooled IT equipment and helped push hot server exhaust air towards the facility's return air ducts.

However, chaos air distribution commonly results in a wide range of significant inefficiencies, including these:

- **Re-circulation:** Hot exhaust air can find its way back into server air intakes, heating IT equipment to potentially dangerous temperatures
- **Air stratification:** The natural tendency of air to mass in different temperature-based layers can force set points on precision cooling equipment to be lower than recommended
- **Bypass air:** Cool supply air can join the return air stream before passing through servers, weakening cooling efficiency

For more information on this topic, please see Eaton's white paper on rack hygiene.

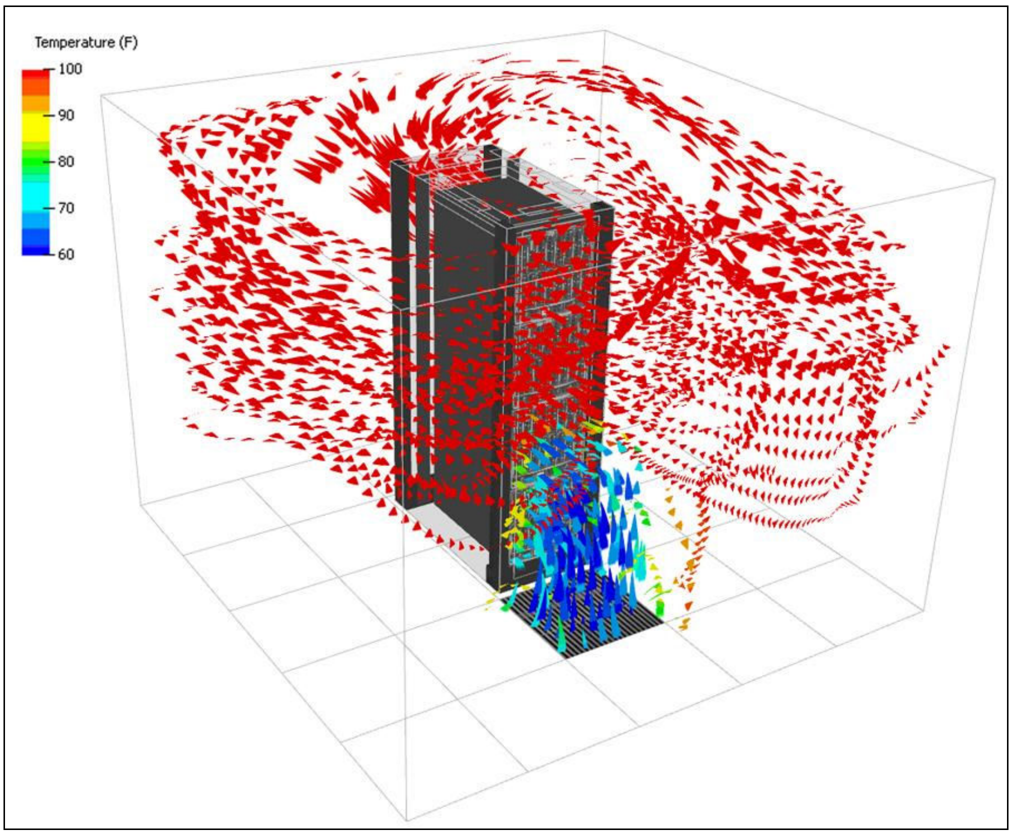

Figure 1. *Re-circulation is a phenomenon in which hot exhaust air re-enters server equipment rather than returning to the CRAC units.*

Hot aisle/cold aisle strategies

Eager to improve their cooling efficiency and keep pace with steadily climbing data center temperatures, businesses soon began adopting hot aisle/cold aisle rack orientation arrangements, in which only hot air exhausts and cool air intakes face each other in a given row of server racks. Such configurations generate convection currents that produce a cooling, continuous airflow.

Still, while superior to chaos air distribution, hot aisle/cold aisle strategies have proven only marginally more capable of cooling today's increasingly dense data centers, largely because both approaches ultimately share a common, fatal flaw: They allow air to move freely throughout the data center.

Containment strategies

This flaw eventually led to the introduction of containment cooling strategies. Designed to organize and control air streams, containment solutions enclose server racks in sealed structures that capture hot exhaust air, vent it to the CRAC units and then deliver chilled air directly to the server equipment's air intakes. The end result is a series of important benefits:

- **Improved cooling efficiency:** By preventing the supply and return air streams from intermingling, well-designed containment solutions eliminate wasteful re-circulation, air stratification and bypass airflow.
- **Increased reliability:** Eliminating re-circulation spares servers from exposure to potentially dangerous warm air that can result in thermal shutdown.
- **Lower energy spending:** To counteract the effects of re-circulated exhaust air, legacy cooling schemes typically chill return air to 55ºF/12.78ºC. Containment-based cooling systems completely isolate return air, however, so they can safely deliver supply air at 65ºF/18.34ºC. As a result, containment cooling strategies typically reduce CRAC unit power consumption by an average of 16 percent.
- **Greater floor plan flexibility:** To generate the cooling convection currents that make hot aisle/cold aisle strategies work, companies must place their server racks in rigidly aligned, uniformly arranged rows. Containment strategies don't rely on convection, however, so they empower data center designers to position enclosures in any configuration that best fits their needs.

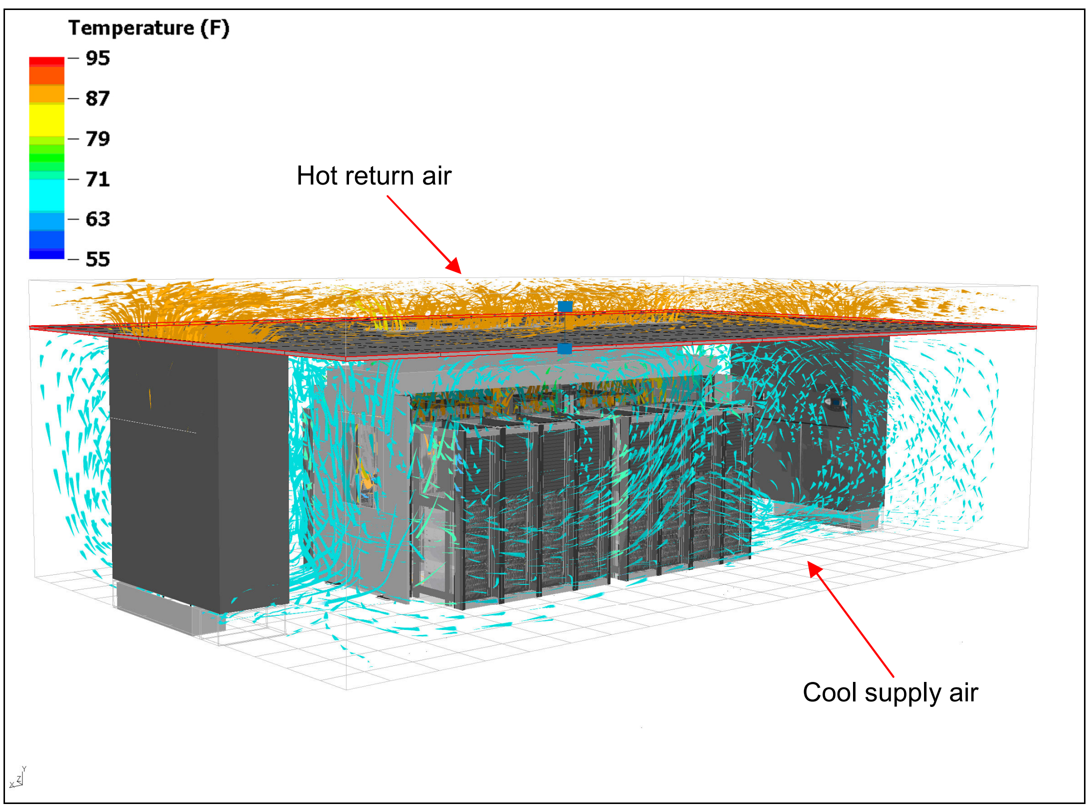

Figure 2. *Containment-based cooling strategies improve efficiency by completely isolating the supply and return air streams.*

The next step: Designing data centers around the rack

Despite the revolutionary impact of containment strategies on data center cooling, most organizations continue to plan new computing facilities the same way they always have. First they design a building and devote some of it to the data hall, or white space. Then they fill the white space with as many server racks as it will hold.

However, designing data centers in that traditional manner can create a wide range of problems. For example, an undersized or oversized power and cooling infrastructure can limit operating capacity or increase capital expenses unnecessarily. Inconveniently located structural elements can force containment ducts to bend and detour in ways that reduce their efficiency. Excessively narrow or insufficiently long room dimensions can complicate server rack placement and produce wasted floor space.

As a result, companies are increasingly recognizing the wisdom of designing data centers not from the walls in but from the server rack out. Instead of building a room and then filling it with racks, they're selecting the ideal racks for their needs and designing the room around them. Instead of under- or overprovisioning their new facility's power and cooling resources, they're installing the optimal infrastructure for the precise array of hardware and enclosures that they'll be using. Instead of improvising solutions to efficiency-sapping structural defects, they're preventing those defects from occurring in the first place. The end result is a data center that's not only less costly to cool and maintain, but more reliable and better suited to business requirements as well.

Key steps in planning a data center around the rack

Planning data centers from the server rack out runs counter to traditional design approaches, but consistently delivers better results. Here are the essential measures to perform:

1. Collect requirements

The core advantage of designing a data center around the server rack is that it allows you to tailor the facility to your exact technical and business needs. Identifying those needs, then, should always be the first step in your planning process. In particular, be sure to evaluate your requirements in these vital areas:

- **Power density:** This is the single most important requirement to estimate, as it drives several important decisions later in the design process. Note that your power density assessment should consider not only near-term needs but long-term ones as well.
- **Budget:** Some of the cooling technologies discussed later in this white paper are more effective than others but also more expensive. A realistic understanding of both the capital and operating funds available to your new data center will help you appropriately balance efficiency with cost.
- **Location:** Where you construct your data center is also likely to influence critical design decisions. For example, organizations building in all but the hottest climates may wish to include an air-side economizer in their new facility. Air-side economizers enable companies to reduce their dependence on power-hungry CRAC units by replacing hot exhaust air with naturally cool air pumped in from outdoors. Studies have shown air-side economizers to be a practical option for at least part of the day even in mild or warm climates.

2. Decide between row- and rack-level cooling

At the highest level, most containment-based cooling schemes fall into one of two categories: Row-level and rack-level. In a row-level system, an entire row or aisle of server racks sits within a shared enclosure that's supported by shared ducts and either CRAC units or a building-wide cooling system. In a rack-level system, each individual rack gets its own enclosure, plenums and cooling devices. Though either option is capable of handling high power densities, rack-level designs tend to be both a little more powerful and a little more costly. That generally makes them a better choice for companies anticipating extreme power densities, provided they can afford the greater upfront investment.

3. Decide between passive and active containment

Most containment strategies rely mainly on passive exhaust systems in which the server hardware's built-in exhaust fans do most of the work of pulling supply air in and driving return air out. Sometimes, however, a phenomenon called backflow or back pressure inhibits airflow in ways that render server fans alone incapable of keeping the return and supply air streams moving properly. In such cases, an active exhaust system equipped with stronger, supplemental fans must be used.

When considering whether to use passive or active containment, keep these points in mind:

- As a rule, most row-based cooling solutions employ passive exhaust designs exclusively.
- It is rare for every enclosure in a rack-level cooling scheme to require active containment. Companies using rack-level cooling are generally better off making passive containment their default choice, and then installing active exhaust products selectively in specific spots prone to backflow.

4. Decide between cold aisle and hot aisle containment

Most row-level cooling schemes feature either cold aisle or hot aisle containment, both of which build off of the traditional hot aisle/cold aisle air distribution methodology. Cold aisle containment systems enclose the cold aisle portions of the server room and allow hot exhaust air to flow freely everywhere else. Hot aisle containment systems isolate the hot aisle sections of the server room and permit cold supply air to move freely elsewhere. Both approaches enhance the efficiency and capacity of your CRAC units, either by directing supply air exclusively at server equipment in the case of cold aisle solutions, or by guiding exhaust air directly back to your chillers in the case of hot aisle solutions.

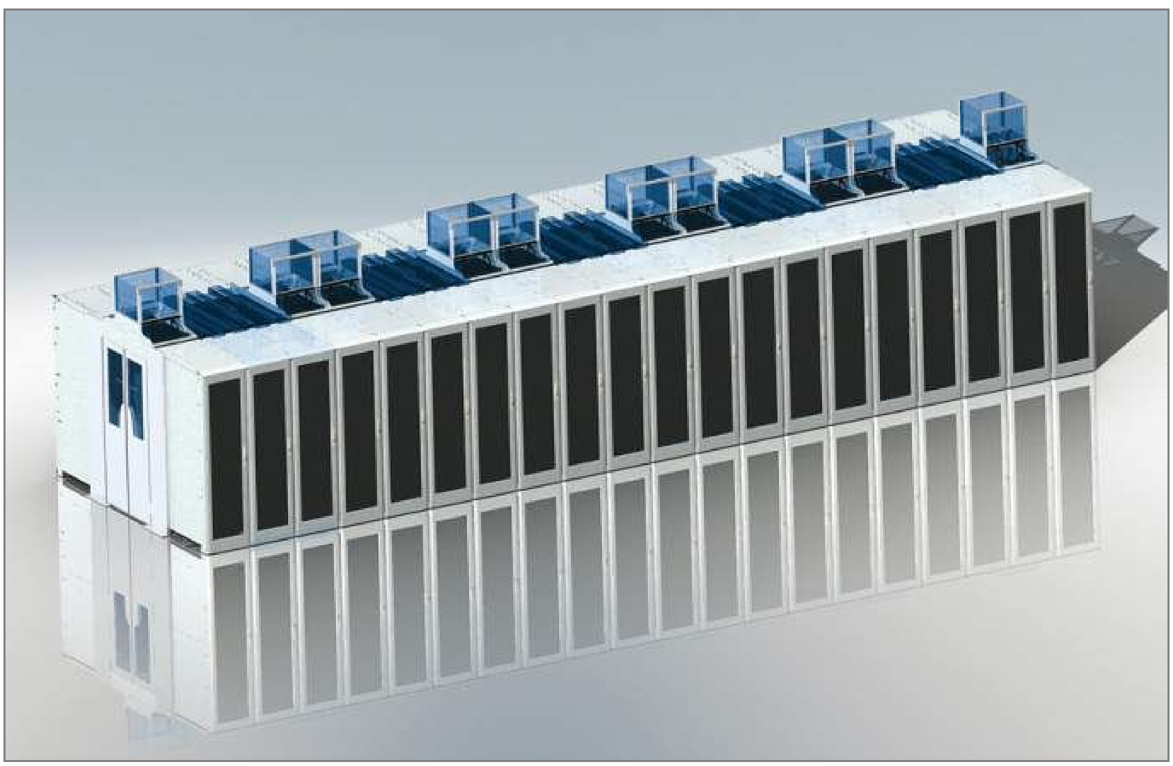

Figure 3. *Hot aisle containment systems like the one depicted here enclose the hot aisle portion of a hot aisle/cold aisle data center.*

Though companies must weigh a number of issues when choosing between cold aisle and hot aisle containment, the following two considerations can help simplify the decision in some cases:

- Since cold aisle containment allows hot return air to circulate freely, temperatures in the server room can quickly become uncomfortable for data center employees and visitors. Hot aisle containment fills the server room with cool supply air, making for better work and observation conditions. As a result, organizations that expect people to be present in the server room on a regular basis often prefer hot aisle containment systems.
- Air-side economization is generally an option only for data centers that utilize hot aisle containment, making it the row-level containment strategy of choice for organizations eager to capitalize on the savings that air-side economizers deliver.

5. Perform a computational fluid dynamics study

Once you've determined what kinds of racks and cooling solutions you will be using and designed a server room around them, your next step should be conducting a computational fluid dynamics (CFD) analysis. Such studies use sophisticated software to model the flow of hot and cold air in potential data center layouts. CFD assessments can help you eliminate flaws and inefficiencies in a data center design before construction begins, when you can far more easily and affordably fix them.

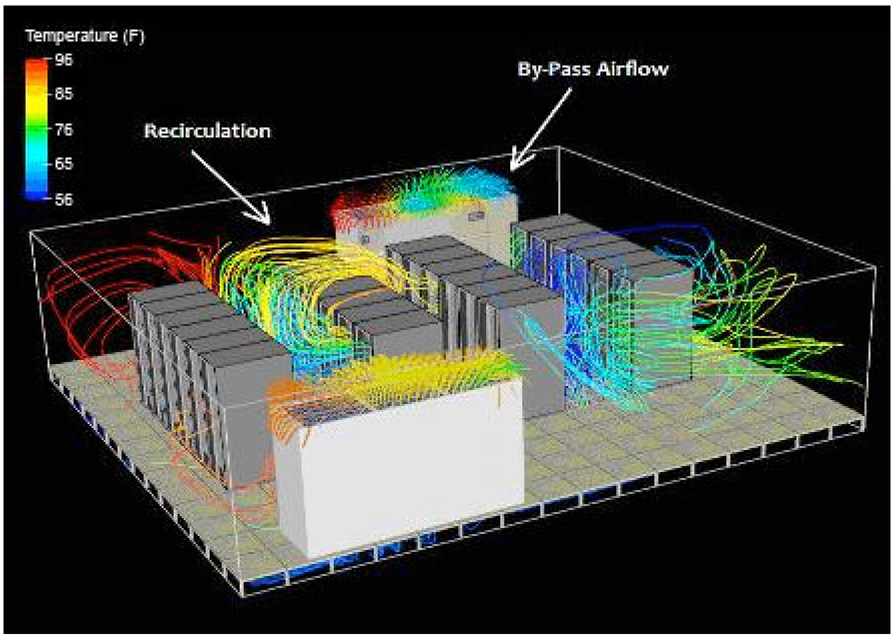

Figure 4. *A CFD study can help you spot and address cooling-related flaws in a data center design.*

Conclusion

Designing a building from the inside out violates centuries of received architectural wisdom. But by making the selection of server racks the starting point of the data center planning process instead of the end point, organizations can improve reliability and save money by ensuring that their power, cooling and IT infrastructure are all optimally suited to their requirements. Every organization, then, should make carefully analyzing their current and future needs and rigorously evaluating potential cooling and containment strategies their first step on the road to architecting new data centers. Embracing that approach to data center design will empower them to improve the efficiency and cost-effectiveness of their mission-critical computing facilities.

About Eaton

Eaton Corporation is a diversified power management company with more than 100 years of experience providing energy-efficient solutions that help our customers effectively manage electrical, hydraulic and mechanical power. With 2011 sales of $16.0 billion, Eaton is a global technology leader in electrical components, systems and services for power quality, distribution and control; hydraulics components, systems and services for industrial and mobile equipment; aerospace fuel, hydraulic and pneumatic systems for commercial and military use; and truck and automotive drivetrain and powertrain systems for performance, fuel economy and safety. Eaton has approximately 73,000 employees and sells products to customers in more than 150 countries. For more information, visit www.eaton.com. For information about server rack and airflow management solutions, visit www.eaton.com/wlrightline.

About the author

John Collins has more than 18 years of experience in the data center industry. He joined Eaton in January 2011 and is solely focused on ensuring the company's data center products and solution offerings evolve with the market. John previously held various roles in sales, sales management, and product management, where he was responsible for various global product offerings relating to power generation, power quality, and power distribution. He's currently involved in many industry groups, including The Green Grid, 7x24 Exchange and AFCOM. John received his bachelor of science in electrical engineering from the University of Rhode Island and served for 10 years in the U.S. Marine Corps.