VIRTUALIZATION
CLOUD
APPLICATION DEVELOPMENT
NETWORKING
STORAGE ARCHITECTURE
DATA CENTER MANAGEMENT
BUSINESS INTELLIGENCE/APPLICATIONS
DISASTER RECOVERY/COMPLIANCE
SECURITY

Tech Guide

Inside Data Center Cooling Approaches

Efficient cooling is a common challenge for data center administrators, but a well-designed cooling system can prevent serious failures while offering substantial cost savings.

TechTarget

Home

Editor's Note

Data Center Cooling Essentials

"Free" Data Center Cooling: Methods And Challenges

Computational Fluid Dynamics Helps Cut Cooling Costs

Saving Energy with Precise Cooling

Today, with higher operating temperatures in data centers and a continual push for IT efficiency, keeping hardware at the right temperature isn't just about filling rooms with refrigerated air. Rather than erring on the side of overcooling, data centers are adopting more controlled methods that are better for budgets (and, as a side effect, the environment).

As Robert MacFarlane shows, properly and cost-effectively cooling racks and servers is a complex undertaking, even if you're using familiar techniques that involve mechanical refrigeration, humidification and containment. In his second piece, MacFarlane goes on to consider economization, or "free" cooling, which uses outside air and can save data centers money.

As hardware runs at higher and higher temperatures, that climate doesn't have to be arctic or even cool all year round. Like all IT trends, this supposedly "free" approach doesn't come without certain trade-offs, but it's a big step toward using only as much (or as little) energy as needed.

And how can you accurately measure your data center's cooling needs? In his article, Clive Longbottom presents one possibility with an explanation of computational fluid dynamics, or CFD, which identifies areas that need cooling most, even in high-density facilities.

The technology helps to deliver cooling only where it's needed, to prevent overcooling. With persistent pressure on data centers to cut energy costs, it's important to consider any opportunity for measuring, targeting and ultimately reducing cooling needs. ■

—Laura Aberle
Associate Features Editor

Data Center Cooling Essentials

Good cooling is essential to the protection of expensive computing hardware and the critical business operations it supports. At elevated temperatures, most computers will shut down for self-preservation, but the highest-performance servers can crash in just seconds if cooling is lost. IT now has to consider designs that not only cool effectively and efficiently, but that also provide "cooling ride-through," which maintains enough cooling between a power failure and generator start to keep critical servers from overheating and crashing.

In recent years, cooling has undergone more changes than any other part of data center infrastructure because it must keep up with rapid advances in high-performance hardware. And since cooling systems are the most expensive parts of the facility to construct and to operate, developing more efficient cooling systems has become a major industry goal.

Gone are the days when data centers had to be refrigerated like meat lockers.

So how to cool? Gone are the days when data centers had to be refrigerated like meat lockers. Not only is overcooling costly in power usage, it is also environmentally irresponsible and unnecessary. In fact, one goal for the future is to be able to cool all data centers, in any climate, with no mechanical refrigeration at all.

While today this is achievable in some regions, the vast majority of organizations will still rely on mechanical cooling for years to come. Therefore, it's important to understand the cooling technologies available now, the best practices for each approach and how cooling will continue to evolve.

Home

Editor's Note

Data Center Cooling Essentials

"Free" Data Center Cooling: Methods And Challenges

Computational Fluid Dynamics Helps Cut Cooling Costs

RISING OPERATING TEMPERATURES

Modern computing equipment is designed to cool itself if the right quantity of air is delivered to it at the right temperature, known as the "inlet temperature." The American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) has confirmed that all modern and legacy hardware is designed for reliable long-term operation with air-inlet temperatures as high as 27 degrees Celsius (80.6 degrees Fahrenheit). This temperature was chosen to provide good cooling with maximum energy efficiency.

Hardware can also handle up to 35 degrees Celsius (95 degrees Fahrenheit) for several days—in the event of a partial cooling failure or an abnormally hot day, for example—without voiding warranties or significantly increasing failure rates. Energy usage increases as fans speed up, but short-term inefficiencies do not negate energy savings achieved over the long term. To be safe, many data center designers and managers choose to operate at around 24 degrees Celsius (75 degrees Fahrenheit), but there is no reason to cool below this.

CONTAINMENT

Even at higher temperatures, an effective cooling strategy is still crucial. The best approach to improved cooling is the separation of hot and cool air. "Containing" the cool air supply or the hot air discharge prevents mixing, but you need to take two steps to successfully accomplish this task.

First, cabinets must be arranged front to front and back to back so that aisles of rack fronts (or "cool" aisles) alternate with aisles of exhausts (or "hot" aisles). Note that, with higher inlet temperatures, we now refer to "cool aisles" instead of "cold aisles."

Second, gaps between cabinets must be closed with fillers, and all unused panel spaces must be blocked with blanking panels (solid plates in place of missing computing hardware). This prevents hot air from the back of computing hardware from recirculating to the front and raising the equipment inlet temperature.

Even after taking these steps, air can still flow over the tops of cabinets

Home

Editor's Note

Data Center Cooling Essentials

"Free" Data Center Cooling: Methods And Challenges

Computational Fluid Dynamics Helps Cut Cooling Costs

and around the ends of rows, so the ultimate solution is to erect more barriers to contain the air. Either the hot aisle or the cool aisle can be contained in this way.

The barriers can be as simple as plastic curtains—fire-rated and anti-static, of course—hung at the ends of rows and possibly above cabinets as well. This is known as "partial containment" because some air can still leak through the curtain seams, but it's very effective nonetheless.

"Full containment" requires solid end panels with doors, plus either solid barriers between the tops of cabinets and the ceiling, or ceiling panels at cabinet height. Full containment is obviously more expensive and less flexible, so the choice between the two is often based on budget or on an existing facility's limitations.

In a new design, full containment should be considered first because it's the most effective, and the rest of the piping, electrical and cable tray can be designed to accommodate solid barriers. But if ceiling height or budget preclude this, then partial containment is better than nothing at all.

"SOURCE OF HEAT" COOLING

A lot of energy is required to push cool air across a room, under a floor or through a duct. Locating cooling units close to computing equipment significantly reduces fan energy consumption. It also makes it easier for cooling units to pull in hot return air before it can recirculate. But the heat ultimately must be moved outdoors, and this requires circulating fluid through the cooling unit. That requires piping out to the cabinet area, either under the floor or overhead.

If you're not comfortable with water being piped through your data center, the alternative is refrigerant, which if it leaks, turns into a harmless gas. But refrigerant piping is also more difficult and expensive to install, balance and alter, so it should be selected with good reason. While water in the data center brings concerns about leaks, the risk of damaging expensive hardware can be reduced with proper installation, as well as precautions like leak detection systems and drains.

Home

Editor's Note

Data Center Cooling Essentials

"Free" Data Center Cooling: Methods And Challenges

Computational Fluid Dynamics Helps Cut Cooling Costs

Another consideration is that many source-of-heat coolers remove only "sensible heat" and are designed for very high heat loads. "Sensible heat" refers to the heat you can feel or sense. It's the heat produced by electronics equipment and contains no moisture.

Therefore, many source-of-heat coolers have no humidity control. Furthermore, many source-of-heat units won't operate at low heat loads. To cool lower-heat-density racks and cabinets, as well as to provide humidity control, conventional air conditioners (like perimeter computer room air conditioners, or "CRACs") are often required as well. In some designs, separate humidification may be all that is needed. There are five fundamental types of source-of-heat cooling, plus a new tactic to consider:

1. **In-row air conditioners** are packaged like equipment cabinets and are intermixed with standard cabinets within or at the ends of cabinet rows. Some units simply discharge air out the front into the cool aisle, while others allow the air to be directed sideways toward cabinets that most need it. They then pull in hot return air from the hot aisle. That air is re-cooled and discharged again out the front.

2. **In-row coolers** (IRCs) are highly efficient for cooling, and they are even more efficient as part of a "contained" environment. Fan speed and cooling capacity are usually controlled via temperature sensors on the front of adjacent cabinets, since it's the inlet temperature that matters.

   In-row systems are available for chilled-water, condenser-water and refrigerant-based heat transfer, with and without humidity control. Chilled-water and pumped refrigerant units (systems without compressors, humidifiers or reheat) use little power, so they can also be maintained on an uninterruptible power supply (UPS) for high-density cooling ride-through after a power failure. Ride-through, however, requires enough chilled water reserve (which often exists within the header pipes), plus small pumps, also on UPS, to circulate it until generators start.

Home

Editor's Note

Data Center Cooling Essentials

"Free" Data Center Cooling: Methods And Challenges

Computational Fluid Dynamics Helps Cut Cooling Costs

Another benefit is that some IRCs can be relocated within the data center as cooling needs change.

3. **Above-cabinet cooling** is similar to in-row cooling, but these units discharge cool air from above, directly in front of the cabinets that need it, and pull hot air back in over the cabinet tops. Some units mount on top of cabinets, and others suspend from the ceiling in the middle of the cool aisle. One system can be configured not only to provide cooling, but also to form a cool aisle "ceiling" as part of a contained solution.

   Overhead coolers are meant to cool very high-density cabinets (generally considered to be 8 kW to 25 kW or more). They remove only sensible heat, so they are used as a supplement to conventional cooling. Overhead systems are controlled by localized temperature sensors. All are refrigerant-based and are highly energy-efficient.

   As with some IRCs, the low energy use of overhead coolers can also be maintained on UPS to provide cooling ride-through until generators start. Both top-of-rack and center-of-row units can also be physically relocated with relative ease as computing equipment is reconfigured.

4. **Self-cooled cabinets** are the epitome of closed-loop systems. These cabinets recirculate air within their own enclosures for the ultimate in contained, source-of-heat cooling. They can use water or refrigerant and are highly energy-efficient. But they tend to be large, heavy and expensive. These "problem solvers" are generally used in rooms without other cooling or to provide high-density "islands" in an otherwise low-density data center without rebuilding the room.

5. **Rear-door coolers** are radiators through which chilled water circulates to remove the heat. They replace the normal rear doors on cabinets and cool the hot exhaust air as it passes through them before discharging it into the room.

   These devices were originally designed to precool very high-temperature discharge air from high-performance computers before it entered the hot aisle and returned to air conditioners for further cooling.

Newer versions can satisfy the total cooling requirement if they are installed on enough cabinets (in many instances, that means every cabinet).

This results in the air being the same temperature everywhere in the room. Since all the air can then serve as either exhaust or intake, cabinets can be arranged facing the same direction, or can still be in a hot-aisle/cool-aisle configuration. Rear-door coolers remove only sensible heat, so, as with the other supplemental systems, additional equipment is required for humidity control. Independent tests have shown fully passive rear-door coolers (those without fans) to be the most energy-efficient cooling units available.

6. **Immersion cooling** is the newest cooling option available. With only small modifications to the disk drives, equipment is fully immersed in a nonconductive oil bath that removes heat directly from circuits. In the event of a failure, the oil bath also provides an enormous thermal mass that can keep equipment cooled, with only a small circulating pump that can be run on UPS.

   This emerging technology came to market in 2011, but it is gaining interest. The main drawback is that it requires modified hardware and makes the equipment messy to work on. But as far as cooling is concerned, it is highly effective.

THE FUTURE OF DATA CENTER COOLING

All of the above approaches rely on some form of mechanical system to transfer heat to the outside air. Ideally, equipment can be cooled directly with air from outside if it is below 27 degrees Celsius. Alternatively, the circulating water could be directly cooled by outside air, requiring only pumps and no mechanical refrigeration. Both these approaches are known as free cooling.

Even if outside air conditions aren't suitable year-round, in most climates, enough days are cool enough for switching between full mechanical and free cooling to be cost-effective. The transition from mechanical to nonmechanical cooling, however, is usually the biggest challenge in a free cooling design.

But in the future, computing equipment may be designed to operate at even higher temperatures. That would enable free cooling to be used in virtually any climate all year long. And as operating temperatures continue to rise, direct liquid cooling, which was used in mainframes until about 1990, is also becoming more common.

Water is approximately 3,500 times more efficient than air in removing heat, so it is only logical to consider this option. Liquid cooling is used for some high-performance hardware and is likely to gain popularity as it becomes less realistic to move the enormous quantities of air needed to cool high-performance hardware. —*Robert McFarlane*

"Free" Data Center Cooling: Methods and Challenges

FREE COOLING ISN'T really "free" unless you can just open some windows on a nice spring day and let the air blow through. However, when you dissipate the heat from a building or data center using the lower temperature outside air instead of mechanical refrigeration, you can save so much in electrical costs that we label it "free cooling."

Another term used in the industry is economization, but regardless of what it's called, the idea is to move heat outdoors using the least energy.

USING THE AIR-SIDE METHOD FOR FREE DATA CENTER COOLING

There are two forms of free cooling, commonly referred to as air-side and water-side. Water-side (or more generically, liquid-side) is still the more common form, and uses circulating water (or a water/glycol mix in freezing climates), to carry the heat to outdoor cooling towers.

Large fans move the air through the water as it cascades down the cooling towers so that heat can dissipate into the air, even if the outside air temperature is relatively high—but fans require energy. When the air temperature is low, little to no fan power is required, but a small amount of energy is still used, mainly for the pumps that circulate water through the pipes.

Air-side free cooling just moves air without the liquid intermediary. There's no electricity needed to pump liquid through pipes, so it should be the ultimate "green" energy saver. But like so many methods, it depends.

Many people still think the air-side method requires little more than opening a hole in the wall or extending a duct to the roof to bring in cool air at night or in the winter. It's not really that simple. When transferring heat, water is many times more efficient than air per unit volume. Therefore, it takes a much greater volume of air to cool a data center than it does water. Nearly 40,000 cubic feet per minute (cfm) of air is needed to cool a 250 kW data center. For a 1 megawatt (MW) data center, it's more than 150,000 cfm. So you'd need a really big hole or a very large duct.

The above numbers were figured at an airflow rate of 2,500 lineal feet per minute (FPM), which is 28 mph. That's quite a breeze! At that velocity, 40,000 cfm requires at least 16 square feet of duct, which might be 2 by 8 feet—doable, but still large. But 2,500 fpm is too fast and would blow air right past the cabinets, so it has to be slowed down.

A better velocity would be 1,500 FPM which, for 150,000 cfm, would require 100 square feet or something like a 4-by-25 foot hole or duct. And even that air would need to be slowed down further before it gets to the cabinets. In short, air cooling requires big ducts.

In large data centers, it's usually not practical to move and filter so much air directly through ducts from outside. That much air, through such a big hole, can bring in a lot of contaminants, so filters just as big are needed to trap particulates. It takes large fans to pull that much air through heavy filters, which add a lot of air resistance even when new, and considerably more as they fill up with dirt. That consumes energy. Then there's the filter replacement cost, which can be monthly or even more often if there's a lot of dirt or pollen in the air.

ADDRESSING AIR-SIDE COOLING CHALLENGES

It is also important to control humidity. Cooling coils can be used to condense moisture out of the air if humidity is too high, or heating coils can be used if temperature and humidity are too low and moisture needs to be added. That's energy use again.

Desiccant dehumidifiers, which use moisture-absorbing chemicals, use no energy themselves, but they still create some resistance to airflow that again requires fan energy to overcome. These factors can quickly offset energy cost savings. And we haven't even addressed gaseous contaminants that can ruin hardware over time and are difficult to extract from the air.

Home

Editor's Note

Data Center Cooling Essentials

"Free" Data Center Cooling: Methods And Challenges

Computational Fluid Dynamics Helps Cut Cooling Costs

In a sufficiently cool, clean climate, air-side free cooling may be a logical choice. In a mid-climate—including locations you might not expect, such as Hawaii—it could still be a consideration. In a hot climate, particularly one with high humidity or with blowing dust or pollutants, it's unlikely to be worthwhile.

Even with the expanded ASHRAE Thermal Envelope, which was developed partly to enable more hours of free cooling in more parts of the country, there are limits. Cost-benefit analyses must always be conducted before implementing any new technology.

The bottom line is that, no matter which method is used, dissipating heat is a process of transferring energy from one medium to another, and ultimately to outside. Mechanical refrigeration makes at least two thermal conversions. We use the unwanted heat to change a liquid refrigerant to a gas, and then use the mechanical system to change it back to a liquid so it can return through the conversion process and absorb heat all over again. Depending on the exact system used, the refrigerant may be circulated to a condenser that transfers heat to the outside air, or it may exchange the heat to water that circulates to outside cooling towers. Both methods take energy, as described above.

Water-side free cooling is simple in principle but tricky in reality.

Water-side free cooling is simple in principle but tricky in reality. It is essentially the addition of valves that bypass the mechanical refrigeration system when it's cool enough outside to carry the heat away by just circulating the liquid. Actually designing and building the system is more challenging than it sounds, and the control systems can get a little tricky.

With air-side cooling, we can avoid contamination and humidity problems by using air-to-air heat exchangers. Their purpose is to transfer heat directly from the air inside a room to cooler air outside a room, with minimal loss of inside air or infiltration of new outside air. Of course, because air is so much less efficient than water for transferring heat, the devices that accomplish this must be quite large. One such unit that has now been well proven is the Kyoto wheel.

ALTERNATIVE DATA CENTER COOLING DEVICES AND METHODS

The Kyoto wheel is a large, slowly rotating aluminum honeycomb—a modern version of the decades-old heat wheel technology—that is specifically tailored to data centers. It rotates at only 3 to 10 rpm and requires minimal energy to turn.

Special fans have been added to more efficiently evacuate the air from the honeycomb, and they use energy too, but this wheel is still said to use only 8% to 25% of the power of equivalent mechanical refrigeration. There is also less than 1% air exchange between outside and inside, and the surfaces are coated to resist dirt buildup and corrosion that would reduce heat transfer efficiency.

The Kyoto wheel is a large, slowly rotating aluminum honeycomb—a modern version of the decades-old heat wheel technology—that is specifically tailored to data centers.

But the wheel is large—about 20 feet in diameter for a unit capable of 600 kW to 850 kW of heat removal—and it is contained in about 1,000 square feet of space on the roof or next to the building. But the wheel can provide 100% free cooling in Dallas for 47% of the year, and in San Jose, Calif., 81% of the year at maximum wheel capacity. Those are pretty good numbers.

Another approach is evaporative cooling. This runs the hot return air from the data center through a metal tube that is kept wet by a water spray, which keeps the tube cool by evaporation. The cool tube, in turn, cools the air that moves through it. This system works best in drier climates where evaporation is maximized, but it still takes fan power to move the air.

However, since the evaporative process can cool the air to below outside temperature, this system can be used more days of the year than the Kyoto wheel, ideally saving more refrigeration energy than it consumes in fan power. A sophisticated example of this approach is APC's relatively new Eco-Breeze indirect evaporative cooler.

The biggest drawback to the air-side method is that, in many climates, free cooling can't cool a data center 100% of the time. The newest ASHRAE Thermal Guideline may change that, but not without upgrading most of today's computing hardware.

While a lot of electricity can be saved whenever free cooling is in use, mechanical air conditioning is still necessary when free cooling is not effective.

With water-side cooling, the only real additions to the cooling plant are valves and controls. With air-side cooling, you still have the capital cost of the mechanical refrigeration system, plus the cost of the air-side cooling equipment.

With air-side cooling, you still have the capital cost of the mechanical refrigeration system, plus the cost of the air-side cooling equipment.

A thorough return on investment (ROI) analysis can determine whether a free cooling solution is really worthwhile, but upcoming energy-conservation regulations may mandate economization whether it's cost-effective or not. The real decision may be between air-side and water-side design.

Regardless of the monetary ROI, free cooling is certainly worth it environmentally, and the newer devices now available make it more realistic than ever to do some cooling with outside air. —*Robert McFarlane*

Computational Fluid Dynamics Helps Cut Cooling Costs

As server density grows, so does the need for better cooling. That's where computational fluid dynamics come in.

The first data centers housed large, incredibly energy-inefficient mainframes that emitted tremendous heat. Ways to cool, and therefore maintain, the integrity of the machine were devised. In most cases, water was used—which was an unfortunate choice if there were leaks, as computers, electricity and water do not mix.

The problem eased with the emergence of distributed computer architectures. Each server came in its own tower case, fans could be placed inside to blow air over the critical components—mainly CPUs and storage devices—and hot air could be vented to the outside.

As server density grows, so does the need for better cooling.

Once more of these servers were put in the data center, the need to provide cooler input air grew, leading to the use of computer room air conditioning (CRAC) units and raised floors to pump the air through. Different form factors, such as a move toward blades and "pizza box" servers, have eliminated the need for large axial fans that shift large volumes of air to smaller, less volume-capable radial fans.

Meanwhile, energy prices have skyrocketed and come into focus. Power usage effectiveness (PUE), which compares the amount of energy used by the total data center facility with that used by the IT equipment, shows that an "average" data center facility uses more than 1 watt in energy for cooling and other peripheral systems for every watt used for purely IT purposes.

Home

Editor's Note

Data Center Cooling Essentials

"Free" Data Center Cooling: Methods And Challenges

Computational Fluid Dynamics Helps Cut Cooling Costs

EQUIPMENT DENSITY AND ECONOMIC DRIVERS

While the energy efficiency of IT equipment is improving rapidly, the massive growth in equipment density is making it more difficult to ensure that systems are running economically.

Sure, there are many ways of approaching the overall cooling needs of a data center, such as using hot and cold aisles, running the facility at a higher overall temperature and using free air cooling. But the problem remains. Higher densities of equipment have created hot spots within the equipment that are difficult to cool effectively.

Lurking in the dark has been an approach that engineers have used for many years. Computational fluid dynamics (CFD) is a means of visualizing the heat map of an environment and then playing what-if scenarios to optimize a system. In areas such as turbine and boiler design, CFD is a proven and useful approach, but it's rarely present in data centers.

To understand why, we have to look at where CFD and data centers have already touched. The facilities team has a strong interest in ensuring that physical disasters don't befall a data center.

Therefore, using infrared and temperature sensors in a facility could highlight issues before they become major problems. A CFD system can minimize false positives by obtaining a view of what is normal throughout a data center and producing a log of places where hot spots are acceptable.

Until now, the IT team didn't see this data. Facilities rarely saw the usefulness of sharing this information with IT because the sensors and systems that used the data were for facilities systems, such as building information modeling/management (BIM) systems.

But data center infrastructure management (DCIM) systems came along from companies such as Eaton, Nlyte Software, Romonet, Intel, Schneider Electric and others. These vendors saw how computational fluid dynamics could predict the effectiveness of a particular layout of IT equipment.

But, at that time, facilities were still predominantly purchasing DCIM. Now, IT groups know what DCIM can do for them and are seeking out products that make sense.

Home

Editor's Note

Data Center Cooling Essentials

"Free" Data Center Cooling: Methods And Challenges

Computational Fluid Dynamics Helps Cut Cooling Costs

HOW IT SHOPS USE CFD

First, IT finds the baseline in an existing data center that shows where cooling is required. Identifying the hot spots allows for targeted cooling rather than a scattershot approach to containing the average temperature of the data center within certain limits.

Many organizations are surprised by how much they can save just by targeting the cooling and removing any focus from those components that run well within a recommended thermal envelope without any cooling.

Next comes finding the know-how for optimizing the environment. For example, it may be beyond the capabilities of the existing system to effectively cool a certain rack full of spinning disks. Splitting the disks across two racks, however, plus placing low-energy network equipment in the space cleared, could enable the cooling process to deal with the heat load without incurring additional expense.

After that, IT will look into introducing new equipment. Administrators could explore in a virtual environment the question of whether existing cooling can deal with the addition of 100 new servers here or a new storage system there. They would have to choose the least expensive, most effective solution before the equipment is even on-site.

Last—and probably most important—is dealing with new architectures. How well will the facility handle a move from build-your-own racks to modular computing platforms, such as Dell vStart, IBM PureFlex or Cisco Systems' Unified Computing System? How will cooling need to change if the network is flattened to make the most of a fabric approach? And, what if the existing CRACs are replaced with free air cooling or other lower-cost systems? Where will cooling need to be focused to ensure the data center still provides continuous high availability?

CFD is a hidden technology that IT and facilities groups should consider using more often. CFD should be a hotter topic than it is today.

—Clive Longbottom

ROBERT MCFARLANE is a principal in charge of data center design at international consulting firm Shen Milsom and Wilke LLC. McFarlane has spent more than 35 years in communications consulting, has experience in every segment of the data center industry and was a pioneer in developing the field of building cable design. McFarlane also teaches the data center facilities course in the Marist College Institute for Data Center Professional program, is a data center power and cooling expert, and is widely published. He speaks at many industry seminars and is a corresponding member of ASHRAE TC9.9, which publishes a wide range of industry guidelines.

CLIVE LONGBOTTOM is the co-founder and service director at Quocirca and has been an ITC industry analyst for more than 15 years. Trained as a chemical engineer, he worked on anticancer drugs, car catalysts and fuel cells before moving into IT. He has worked on many office automation projects, as well as control of substances hazardous to health, document management and knowledge management projects.

TechTarget

Inside Data Center Cooling Approaches is a SearchDataCenter.com e-publication.

Margie Semilof
Editorial Director

Lauren Horwitz
Executive Editor

Phil Sweeney
Managing Editor

Laura Aberle
Associate Features Editor

Eugene Demaitre
Associate Managing Editor

Linda Koury
Director of Online Design

Neva Maniscalco
Graphic Designer

Rebecca Kitchens
Publisher
rkitchens@techtarget.com

TechTarget
275 Grove Street, Newton, MA 02466
www.techtarget.com

© 2013 TechTarget Inc. No part of this publication may be transmitted or reproduced in any form or by any means without written permission from the publisher. TechTarget reprints are available through The YGS Group.

About TechTarget: TechTarget publishes media for information technology professionals. More than 100 focused websites enable quick access to a deep store of news, advice and analysis about the technologies, products and processes crucial to your job. Our live and virtual events give you direct access to independent expert commentary and advice. At IT Knowledge Exchange, our social community, you can get advice and share solutions with peers and experts.