Liebert®

IT White Paper

BLADE SERVERS AND BEYOND: ADAPTIVE COOLING FOR THE NEXT GENERATION OF IT SYSTEMS

EMERSON
Network Power

Summary

Adaptive cooling principles provide a roadmap for dealing with heat densities that are increasing unpredictably.

High-density servers and communications switches, increased emphasis on business continuity, the convergence of voice and data, and new support system technologies are all driving change in the traditional data center.

In enterprise data centers, designers face the challenge of having to create facilities with a 20-year lifespan when they are unsure how technology will change in the next three to five years. This is impacting all areas of data center design, but none more than how rooms and racks are powered and cooled. Traditional approaches to cooling are only effective to a point. They lack the scalability and adaptability required to effectively cool the high-density blade server racks being deployed today, let alone those coming in the future.

As a result, data center managers are commissioning facilities that optimize current approaches to cooling while supplementing these systems with new zone- and spot-based cooling systems. This adaptive, hybrid approach provides a cost-effective solution to the requirements of today's systems while enabling the flexibility to adapt to whatever the future brings.

Adaptive cooling principles provide new and existing facilities a roadmap for dealing with heat densities that are increasing unpredictably. Adaptive cooling provides maximum flexibility and scalability with the lowest cost of ownership while maintaining or improving availability.

Equipment Trends

As computer chips pack more power into smaller spaces, the heat generated climbs dramatically. New high-performance equipment, such as dual-processor servers and high-speed communication switches, are raising rack densities well above 10 kW.

Figure 1 illustrates how the heat density of server technology has outpaced predictions made by The Uptime Institute just four years ago. With server power requirements exceeding projections, cooling strategies must adapt at a faster pace than anticipated to avoid downtime, equipment failure and reduced lifespan of electronics.

And with continued pressure to drive down data center operating costs, many organizations are attempting to pack as much equipment into as small a space as possible. As a result, rooms are heating up and organizations are feeling the affects.

Already, many of today's data centers require well over 100 Watts of power per square foot, or more than 10 times that of the average household. The latest generation of blade servers pushes power and heat levels even higher. A single rack loaded with six fully configured IBM x10 BladeCenter™ servers, each drawing 4 kW, creates a load of 24 kW in an enclosure that takes just 7 square feet of data center floor space. This shows a sharp contrast with the state of the industry just five years ago, when the average rack consumed just 1 kW of power.

Communications equipment is progressing in the same direction. Depending on its power supply configuration, the Cisco CRS-1 router creates a heat load of 15 to 16.6 kW per rack.

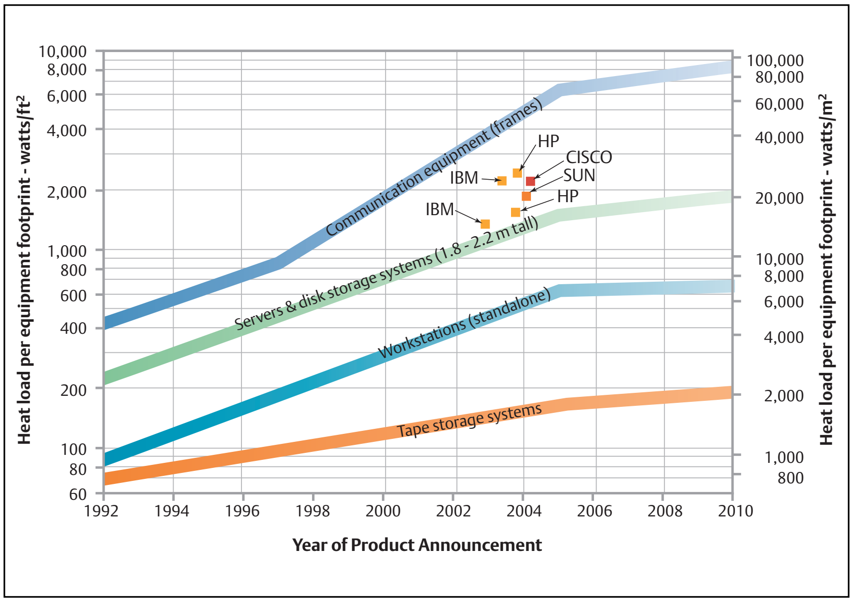

Figure 1. Equipment densities are rising even faster than predicted by The Uptime Institute.

Reprinted with permission of The Uptime Institute from a White Paper titled Heat Density Trends in Data Processing, Computer Systems, and Telecommunications Equipment Version 1.0. Manufacturer information added by Liebert.

All that power is transformed into heat, and if the heat from a rack is not effectively removed, the performance, availability and lifespan of the equipment in the rack will be reduced significantly. Increasingly, as organizations adopt the latest server technologies into their existing data centers, they are exposed to higher failure rates, especially in the top third of the rack. As cooling air is supplied from the raised floor, it is fully consumed by high-density equipment at the bottom of the rack, while the top of the rack is deprived of the cooling air it requires. This is compounded by the fact that high densities also cause hot air to be recirculated

A single rack loaded with six fully configured IBM x10 BladeCenter™ servers, each drawing 4 kW, creates a load of 24 kW in an enclosure that takes just 7 square feet of data center floor space.

For every increase of 18 degrees F above 70 degrees F, long-term electronics reliability falls by 50 percent.

back through the top third of the rack, as illustrated by the red plume in Figure 2.

These experiences are validated by The Uptime Institute. The organization has reported that equipment located in the top third of a data center rack fails twice as often as equipment in the bottom two-thirds of the same rack. They also estimate that, for every increase of 18 degrees F above 70 degrees F, long-term electronics reliability falls by 50 percent.

The increasing failure rate at the top of racks is occurring because current air delivery through a raised floor is generally limited to cooling an average room load of about 150 Watts per square foot, or racks with 2 to 3 kW load. Beyond that point, the volume of air that can be effectively delivered to the equipment in the upper part of the rack is insufficient.

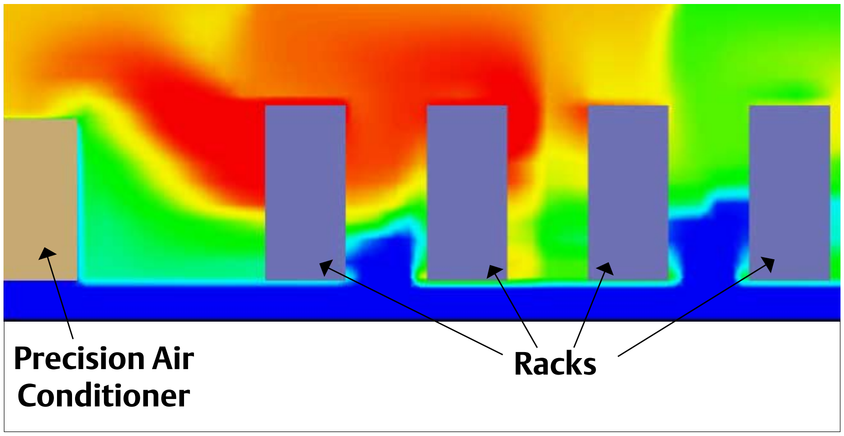

Figure 2: As heat densities rise, equipment at the bottom of the rack consumes the cold air from the floor, causing hot air to be recirculated through the top of the rack.

Responding to Equipment Trends

As equipment heat densities have risen faster than many expected, data center managers have been forced to consider new approaches to data center cooling. Among those that have been tried:

<u>Increased Spacing</u>
Some data center managers have responded to the problem by spreading out the load. They've made sure that racks are only partially populated and have increased aisle widths between high-density racks. This spreads the heat from these racks over a larger area, but consumes valuable floor space.

Hardware manufacturers are now rolling out products that will operate at more than 16 kW of power per rack. Based on field measurements at various sites, the average actual airflow in raised floor sites averages 250 cubic feet per minute (cfm) or less through each perforated floor tile, creating the ability to dissipate about 2 kW of heat. Spacing the equipment in a way that will allow existing airflow to dissipate 16 kW would require aisle widths of 16 feet.

Resolving cooling capacity via rack spacing drastically reduces the number of racks the data center can accommodate. Using traditional under-floor cooling, and spacing racks in a hot aisle/cold aisle configuration, a 10,000-square-foot data center can support only 50 racks if average rack density is 10 kW.

Turner Construction recently analyzed construction and cooling costs to accommodate a 4,000 kW load. To demonstrate the impact of different levels of heat density, three facility sizes were selected and total data center costs (construction, security, cooling, power and UPS) were examined (see Figure 3).

The first option, designing the cooling to accommodate 50 Watts per square-foot, requires 80,000 square feet of space to accommodate the 4,000 kW load. This facility would cost approximately $17 million or $4,000 per kW of load. By increasing cooling capacity to 400 Watts per square-foot the same load can be condensed into a 10,000 square foot facility. The cost for this facility would be under $10 million, or about $2,400 per kW of load.

Clearly, real estate costs vastly outweigh any premium required for cooling higher density loads. By opting for a smaller facility, enabled by extreme density cooling options, the facility in this example saves $7 million in capital costs.

<u>Adding Exhaust Fans</u>

Another common fix is to add exhaust fans to the racks. It's important to remember, however, that fans do not remove heat — they just move it around. In fact, fans actually add to the room's power requirements, heat load and noise level. As a result, increasing the number of fans increases rather than decreases the need for additional cooling capacity. For example, if fans that draw between 200 and 500 Watts per blower assembly are added to

Clearly, real estate costs vastly outweigh any premium required for cooling higher density loads. By opting for a smaller facility, enabled by extreme density cooling options, the facility in this example saves $7 million in capital costs.

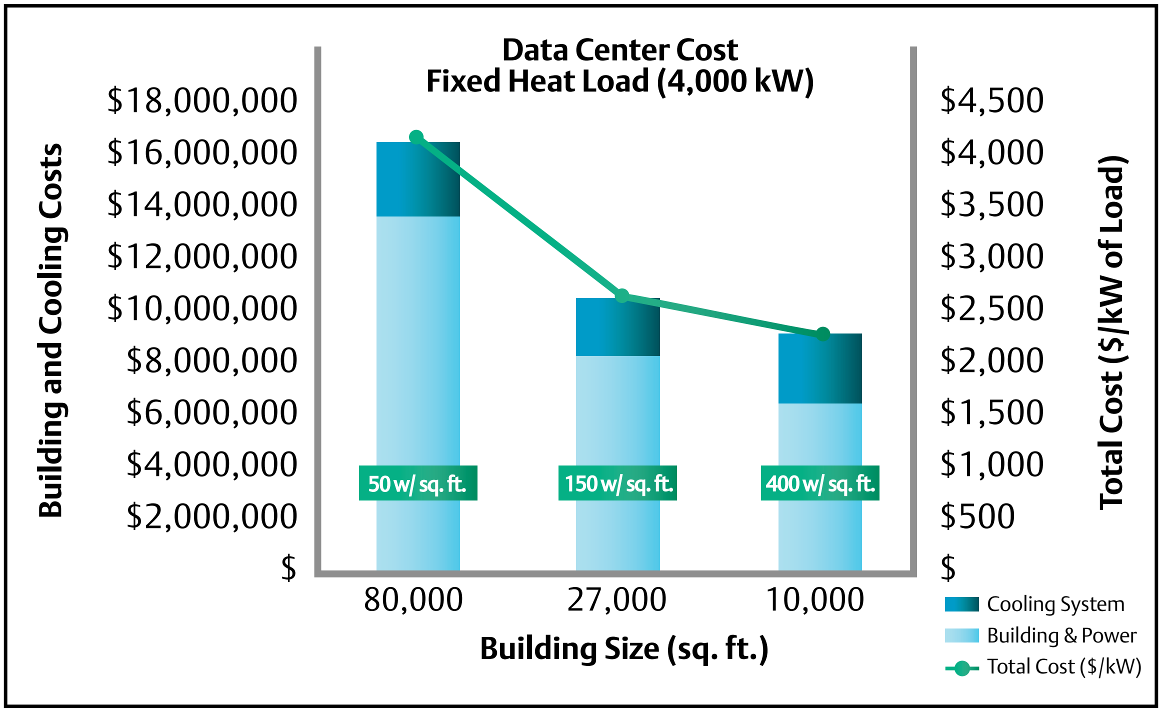

Figure 3: Cooling at higher capacities per square foot makes better use of space and reduces total data center costs significantly.

Rather than approaching the challenge of exploding heat removal requirements using limited, traditional measures, what's needed is a shift in approach.

500 racks in a room, one to three additional 30-ton air conditioning units are required just to remove the heat generated by the fans.

Close-Coupled Cooling

Another potential solution is the close-coupled cooling system, where the electronic and cooling equipment are located together in a sealed environment. This approach can provide high-capacity cooling, but at the expense of flexibility and fault tolerance. It offers almost no flexibility of rack combinations and often no backup emergency cooling. If the cooling fails, racks are isolated from room cooling and the temperature in the enclosure can reach the server over-temperature limit condition in less than 5 seconds.

Figure 4 shows how quickly various density systems can overheat without cooling. This is compounded in close-coupled systems, where the room is not used as a buffer against overheating.

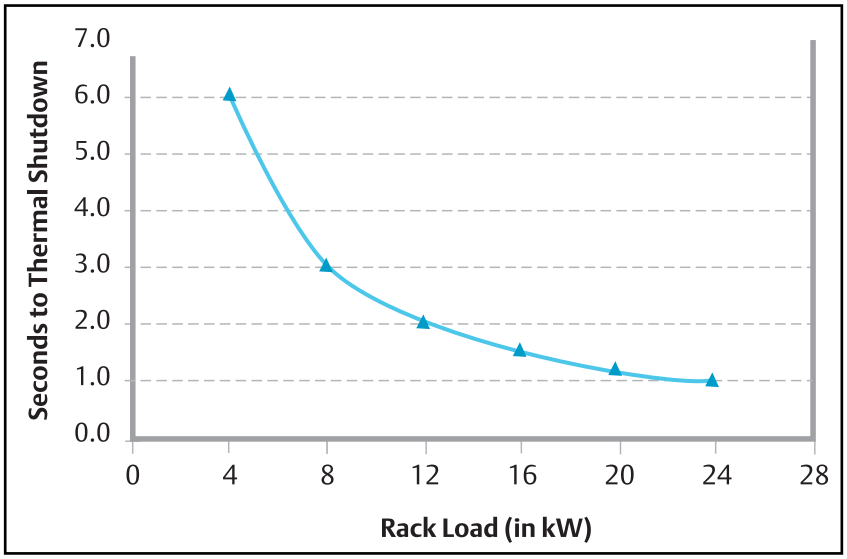

Figure 4: Time to thermal shutdown by density.

Answering Uncertainty with Adaptability

Rather than approaching the challenge of exploding heat removal requirements using limited, traditional measures, what's needed is a shift in approach. Because the constant in data center heat loads has been rapid, unpredictable change, a new, adaptive approach to cooling is replacing traditional measures as a best practice.

A survey of nearly 100 members of the Data Center User's Group* revealed that data center managers' top three concerns were density of heat and power (83 percent), availability (52 percent), and space constraints/growth (45 percent).

Answering these concerns requires an approach that delivers the required reliability and the flexibility to grow, while providing the lowest cost of ownership possible. That means:

- solutions that are able to effectively and efficiently address high-density zones
- flexible options that are easily scalable
- technologies that improve energy efficiency, and
- ease-of-use in maintenance and support to ensure maximum availability.

These requirements can be achieved by optimizing the cooling infrastructure and carefully selecting the two components of adaptive cooling: traditional under-floor cooling and supplemental overhead cooling.

* Data Center User's Group is sponsored by Liebert.

Physical Infrastructure Optimization

The following areas should be evaluated when optimizing the cooling infrastructure.

Raised Floor
Today's data centers are typically built on an 18- to 36-inch raised floor. The higher the raised floor, the greater the volume of air that can be evenly distributed under the floor and the higher the potential capacity of the cooling system. In existing data centers, however, increasing floor height presents an impractical answer to rising heat densities as it introduces a major disruption to data center operations, which not many data centers can endure. Even if a data center can manage the disruption, and if ceiling height allows for more space to be taken up under the floor, there is a limit to what can be accomplished with floor height alone. For example, a floor height of nearly 5 feet would be required to accommodate cooling for heat loads of 400 Watts per square foot.

Hot Aisle/Cold Aisle Configuration
Most equipment manufactured today is designed to draw in air through the front and exhaust it from the rear. This allows equipment racks to be arranged to create hot aisles and cold aisles. As recommended by ASHRAE TC 9.9 (American Society of Heating, Refrigerating and Air-Conditioning Engineers, Technical Committee 9.9) in their Special Publication "Thermal Guidelines for Data Processing Environments," this approach arranges racks front-to-front so the cooling air rising into the cold aisle is pulled through the front of the racks on both sides of the aisle and exhausted at the back of the rack into the hot aisle (see Figure 5). Only cold aisles have perforated floor tiles, and floor-mounted cooling is placed at the end of the hot aisles — not parallel to the row of racks.

Parallel placement can cause air from the hot aisle to be drawn across the top of the racks and to mix with the cold air, causing insufficient cooling to equipment at the top of racks and reducing overall energy efficiency, as was seen in Figure 2.

With the hot aisle/cold aisle approach, improved cable management, both within the rack and under the floor, can also yield increased efficiency. As much as possible, cable management should be limited to the raised floor space below the hot aisle so cables do not

With the hot aisle/cold aisle approach, improved cable management, both within the rack and under the floor, can also yield increased efficiency.

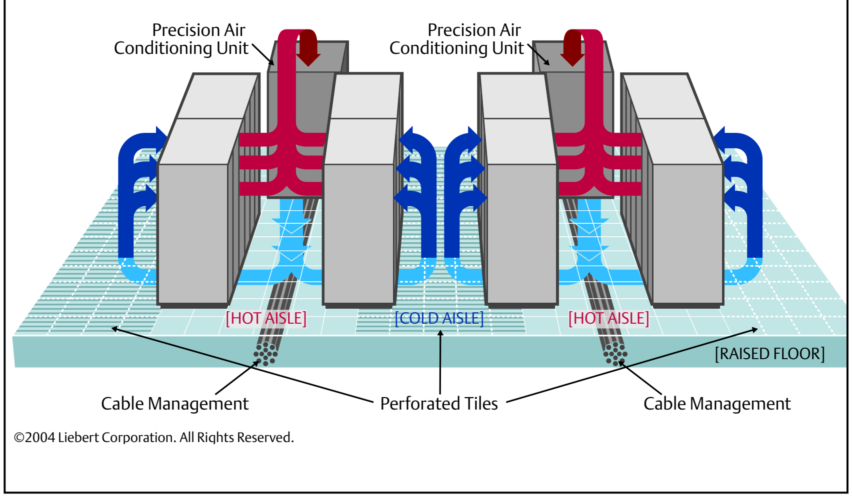

Figure 5: Racks arranged in a hot aisle/cold aisle configuration.

It's recommended that traditional systems be configured to deliver the required cooling for the first 100 Watts per square foot to 150 Watts per square of the data center heat load as well as the room's full humidification and filtration requirements.

impede cooling air's path to equipment. Additionally, new racks being introduced now feature expansion channels that improve cable management and ease heat removal for high-density racks. In some cases, existing racks can be retrofitted with these expansion channels.

While auxiliary fans are a viable option if a data center is dealing with a limited number of high-density racks, as detailed previously, fans emit heat as they hasten airflow, and are counterproductive when used beyond a few racks.

Vapor Seal
As equipment densities increase, a vapor barrier that isolates the controlled environment from the building environment becomes even more critical. Without a good vapor seal, humidity will migrate into the data center during the hot summer months and escape during the cold winter months. An effective vapor seal minimizes the energy required to either dehumidify or re-humidify.

Responding to Changing Requirements

Traditional floor-mounted cooling systems with under-floor air delivery will continue to play an essential role in data center cooling. It's recommended that traditional systems be configured to deliver the required cooling for the first 100 Watts per square foot to 150 Watts per square of the data center heat load as well as the room's full humidification and filtration requirements. With floor-mounted cooling systems optimized, the next element of adaptive cooling is overhead supplemental cooling, which can take a data center beyond the 150 Watts per square foot limit of traditional cooling solutions to well over 500 Watts per square foot.

Traditional Under-Floor Cooling

First though, traditional cooling should be maximized to make sure it is as efficient, flexible and reliable as possible. As demands grow, floor-mounted systems are changing to better meet new requirements. Features that will deliver the highest possible reliability and efficiency and the lowest total cost of ownership include:

Variable Capacity: ASHRAE has determined that the maximum cooling load occurs less than 5 percent of the time. Accordingly, cooling systems should effectively operate at varying loads. Units should have compressors that are capable of stepped unloading – or total variable capacity – to deliver the desired cooling requirements without cycling the compressor off and on. For example, a system operating with two compressors partly loaded will consume approximately 50 percent of the energy of a fully loaded system, but will deliver 76 percent capacity because the condenser and evaporator are sized for full load. New variable capacity systems provide even more precise control of capacity.

By reducing compressor cycling, variable capacity systems reduce compressor starts and stops (on/off), which is one of the leading causes of wear in a compressor.

Unit-to-Unit Communication: Communication between units operating as a system (team) also enhances total cooling efficiency. This is even more critical in rooms with high-density loads, as zones within the room may be operating at a significantly higher temperature than other areas. This ensures that units are not fighting each other by dehumidifying while others are humidifying and provides the ability to direct specific cooling to the high-heat zone, thus improving the energy efficiency of the data center.

Service Organization Availability: As the heat load increases, the margin of error in a cooling system design becomes more critical. A solid preventive maintenance program will keep equipment in top operating condition. Available 24-hour local service and regular preventive maintenance by trained professionals to counteract mechanical wear-and-tear, are required.

In addition, new technologies exist that improve communication functions to provide increased support for maintenance programs. State-of-the-art options include diagnostic and support tools, maintenance notification triggers and internal logging of maintenance events, including predictive diagnostics.

Overhead Supplemental Cooling

The final building block of the adaptive cooling approach is the as-needed addition of overhead supplemental cooling units.

This supplemental cooling solution works in concert with traditional under-floor cooling systems for both existing and new data centers by providing effective cooling where the under-floor system leaves off. This method provides the necessary incremental, 100 percent-sensible cooling for the added heat load and leaves the requirements of humidity control and filtration to the base under-floor cooling system.

This approach delivers the reliability expected in a data center while providing a roadmap for growth for high- density loads through scalable modules with virtually no space requirements and a lower incremental energy cost. These elements combine to provide the most effective cost structure for growth.

High reliability is accomplished by placing the supplemental units close to the high density source, either in the ceiling or above the rack. This configuration supplies the necessary cold air to the top sections of the rack to “supplement” the air delivered from under the floor (See Figure 6).

As densities increase, cooling loops can be interlaced to provide maximum cooling redundancy serving several racks. Additionally, since the cooling module is not directly “close-coupled” with the heat load, the air in the room is used as a buffer in the event of a power failure to the cooling system, providing the

This supplemental cooling solution works in concert with traditional under-floor cooling systems for both existing and new data centers by providing effective cooling where the under-floor system leaves off.

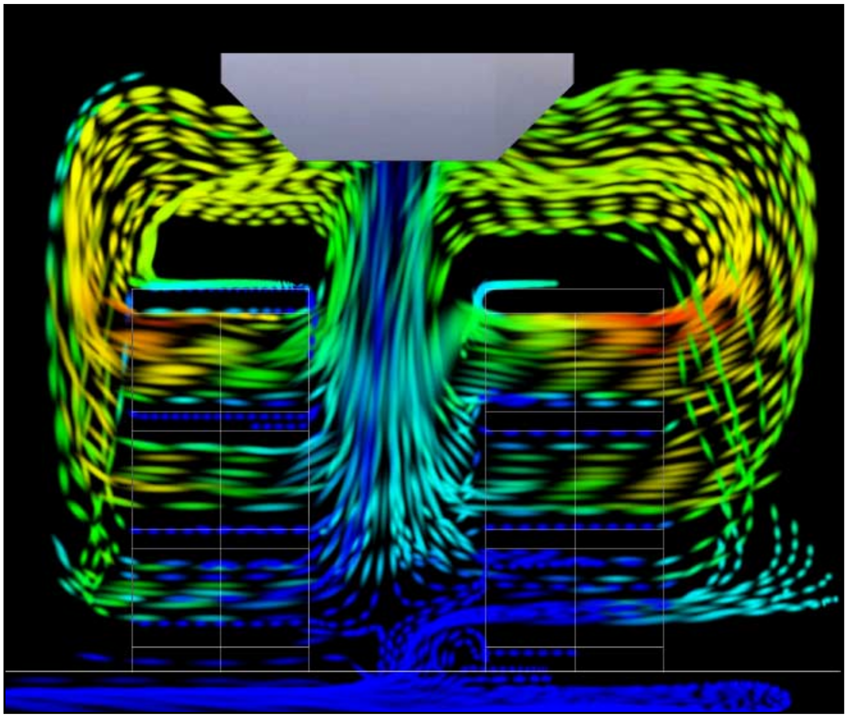

Figure 6: Cool air reaches both the top and bottom of the rack when overhead cooling is used to supplement the under-floor system.

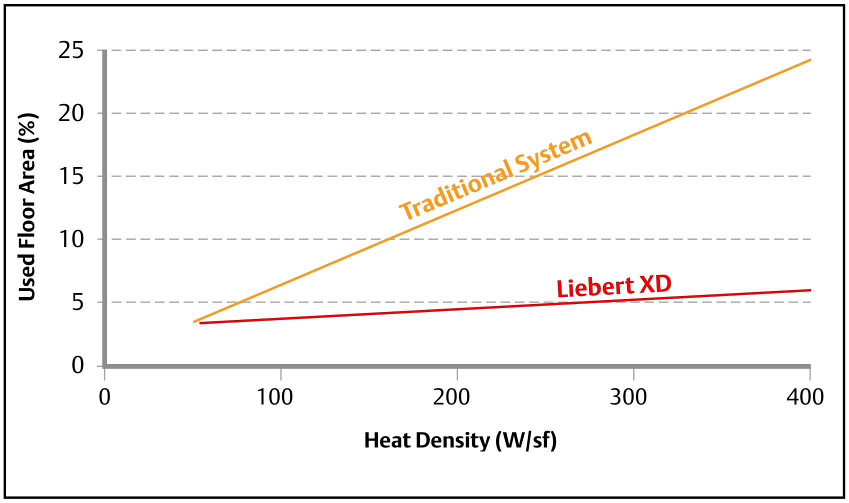

Figure 7: Supplemental cooling dramatically increases space utilization.

necessary ride-through until back-up power is restored. To verify the performance, this adaptive cooling method was modeled with Computational Fluid Dynamics (CFD) to over 500 Watts per square foot. Additionally, an actual data center was developed to prove and demonstrate the operation and performance of this approach.

But the most important feature of supplemental cooling is the energy savings. Overhead supplemental systems that are currently on the market utilize 17 percent less energy than traditional floor-mounted precision air conditioners. The reason for this savings is the fan horsepower required to move the air is 64 percent less, since it has to move the air less than 3 feet against zero static pressure. Additionally, since these cooling modules provide 100 percent sensible cooling, there is no wasted energy in dehumidifying the air and then re-humidifying the air to maintain the required room humidity level.

Finally, this element of the adaptive approach provides the end user maximum flexibility in both growth and arrangement of the data center. Since valuable raised floor space is not consumed, it does not constrain the end user with unnecessary rack orientations, condensation removal, ducting or equipment placements. Plus, today's supplemental cooling conforms to any rack manufacturer's equipment. With available pre-piping options, the facility can be equipped with the necessary piping in the ceiling that permits the end user to add or move 8 kW modules at a time by a simple "plug and play" connection while the other cooling modules continue to operate. This pre-piping option can be handled at one time for a fraction of the cost of system equipment costs.

This adaptive cooling approach permitted Virginia Tech to solve the space, heat density and energy requirement challenges for their Super Computer site, which was initially configured to accommodate over 200 Watts per square foot. But with flexibility in mind, Virginia Tech reconfigured the data center in half the space and reallocated the cooling modules in the final space to accommodate well over 350 Watts per square foot of heat. (See Figure 8.)

Figure 8: Supplemental cooling solutions at Virginia Tech's data center

Conclusion

An adaptive cooling architecture permits the flexibility to grow with constantly changing electronics systems. With proper planning and the use of next-generation equipment, the total costs of operating a data center can hold steady or even decrease, even as network availability becomes more critical and environmental regulations tighten.

Components of adaptive systems incorporate the latest generation of floor-mounted cooling in a raised-floor environment, overhead extreme-density cooling over hot zones, rack-mounted cooling to address hot spots, and an environmental monitoring system to detect and address problems before they threaten network availability.

Adaptive cooling offers a roadmap to help data center managers progress toward supporting loads that are consistently rising in excess of predictions. The adaptive approach is flexible, scalable, reliable and efficient in terms of energy, cost and space. Unlike other approaches, it does not sacrifice floor space or reliability.

Liebert Corporation
1050 Dearborn Drive
P.O. Box 29186
Columbus, Ohio 43229
800.877.9222 (U.S. & Canada Only)
614.888.0246 (Outside U.S.)
Fax: 614.841.6022
www.liebert.com

Printed in U.S.A. 0405 WP406