

A Flexible Approach to Protecting IT System Availability

Adaptive Cooling offers Scalable, Reliable and Efficient Data-Center Cooling in a Time of Density Uncertainty

by Steve Madara

(Editor's Note: This paper has been reviewed for technical merit, accuracy and usefulness by the staff of The Uptime Institute. It was prepared for inclusion in the first edition of the *Quarterly e-Journal on High-Density Computing: The Path Forward 2006* the official publication of *The Institute's High-Density Computing e-Community of Practice*.)

New high-performance computer equipment is raising rack power and heat densities well above 30kW. Further, with server-replacement cycles averaging two to three years, computer platform diversity within the data center is making efficient, effective cooling an even greater challenge.

A survey of 100 members of Emerson Network Power's Data Center Users Group reveals that data center managers' top three concerns are *power and heat* (83 percent), *availability* (52 percent), and *space constraints/growth* (45 percent).

Addressing these concerns requires approaches that deliver required reliability *plus* the flexibility to grow, while providing the lowest total *cost of ownership* (TCO) possible. The right solutions will:

- *effectively address high-density zones*
- *offer flexible, easily scalable options*
- *improve energy efficiency, and*
- *be easy to maintain and support.*

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

Dangers of Insufficient Cooling

As organizations adopt the latest server technologies into existing their data centers, they're exposed to higher failure rates, especially in the top third of the computer racks. As cooling air is supplied from the raised floor, it's often fully consumed by high-density (HD) equipment at the bottom of the rack, while the top may be deprived of required cooling air. Compounding this, HD also causes hot air to be *re-circulated back* through the top third of the rack (see the *red plume* in figure one).

The Uptime Institute reports that equipment located in the top third of a rack fails twice as often as equipment in the bottom two-thirds. Also Uptime Institute research estimates that, for every increase of 18 degrees F above 70 degrees F, long-term electronics reliability falls by 50 percent.

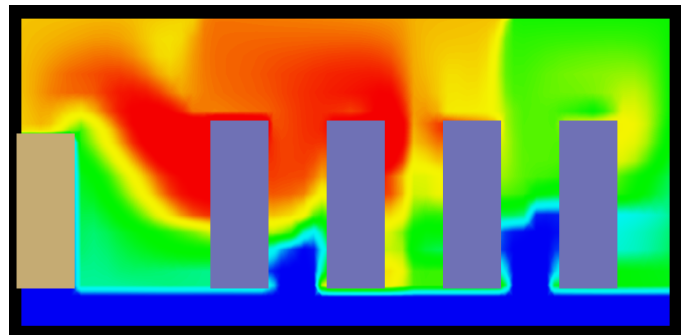


figure one—air temperatures in HDC racks

Responding to Equipment Trends

As heat densities have risen faster than many expected, and as heat loads have become increasingly diverse within the data center, managers must consider new approaches to data center cooling. Among those are:

increased spacing—Some have responded to the problem by spreading out the load. Resolving cooling capacity via rack spacing drastically reduces the number of racks the data center can accommodate. Turner Construction (New York City), builder of such state-of-the-art facilities as the Amazon.com data center in

downtown Seattle, recently analyzed construction and cooling costs to accommodate a 4,000kW load. To demonstrate the impact of different heat densities, total data center costs—construction, security, cooling, power and *uninterruptible power supply (UPS)* — were examined for three facility sizes (see figure two). The first option, designing the cooling to accommodate 50W/sf requires 80 thousand feet of space to accommodate the 4,000kW load. This facility would cost approximately \$4,000/kW of load. By increasing cooling capacity to 400W/sf, the same load can be condensed into a 10,000 square foot facility. The cost for this facility would be under \$2,400 per kW of load.

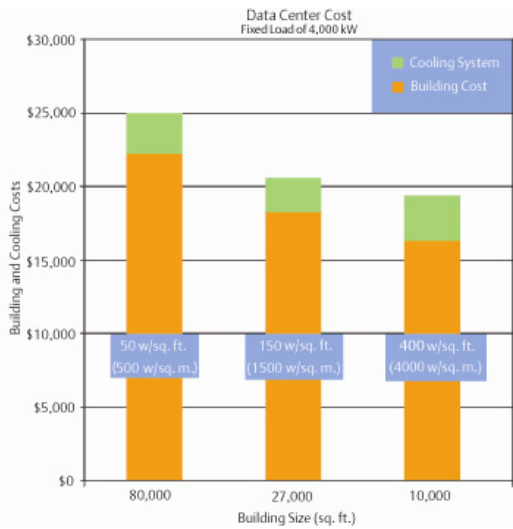


figure two—impact of different heat densities, total data center costs in three facility sizes

Clearly, real-estate costs vastly outweigh any premium required for cooling higher-density loads. By opting for a smaller facility, enabled by high-density cooling options, the facility in this example saves 45 percent of capital costs.

adding exhaust fans—Another common fix is to add exhaust fans to racks, but fans actually add to the room’s power requirements, heat load and noise; fans drawing between 200 and 500W per blower assembly are added to 500 racks in a room, then one to three additional 30-ton AC units are required just to remove the heat generated by those fans.

Rather than approaching the challenge using limited, traditional measures, a shift in approach is needed. Because the *new constant* in data center heat loads is now rapid, unpredictable change, a more adaptive approach to cooling is replacing traditional measures. Cooling requirements of high-density systems can be achieved by optimizing the cooling infrastructure and carefully selecting the two components of adaptive

cooling: traditional under-floor cooling and supplemental cooling.

Physical Infrastructure Optimization

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

hot-aisle/cold-aisle—By configuring racks front-to-front, with perforated tiles in the cold aisles only, cool intake air is isolated from hot exhaust air, optimizing cooling capacity and energy efficiency. Floor-mounted cooling is placed at the end of the hot aisles—not parallel to the rows. Parallel placement can cause hot aisle air to be drawn across the top of the racks and to mix with cold aisle air, causing insufficient cooling to equipment at the top of racks and reducing energy efficiency.

cable management—With the hot-aisle/cold-aisle approach, improved cable management 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 impede cooling air’s path. Additionally, some racks feature expansion channels that improve cable management and ease heat removal. Often, existing racks can be retrofitted with these expansion channels. Running cables above or through racks is also helping to reduce cabling under the floor. Many organizations are taking power distribution closer to the load by using advanced power strips at the rack-level, dramatically reducing the number of cables coming into the rack.

vapor seal—As equipment densities increase, a vapor barrier that isolates the controlled environment from the building environment becomes more critical. An effective vapor seal minimizes the energy required to de-humidify and re-humidify by preventing external humidity from migrating into or out of the data center.

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 to 150 Watts per square foot 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 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.

First though, traditional cooling should be maximized to ensure it provides an efficient, flexible and reliable foundation for adaptive cooling.

Optimizing Traditional Under-Floor Cooling

As demands grow, floor-mounted systems are changing to better meet new requirements. Features that will deliver the highest reliability and efficiency and the lowest total cost of ownership include:

variable capacity — Since the maximum cooling load occurs less than five percent of the time (*source: ASHRAE*), cooling systems should effectively operate at varying loads. Units' compressors should be capable of stepped unloading—or total variable capacity—to deliver the precise requirements desired without cycling the compressor off and on; *i.e.* 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.

unit-to-unit communication—Communication between units operating as a system also enhances total cooling efficiency. This is even more critical in rooms with HD loads, as zones within the room may be operating at a significantly higher temperature than other areas. This ensures units aren't countering each other by dehumidifying while others are humidifying and provides the ability to direct specific cooling to the high-heat zone, thus improving energy efficiency.

service organization availability—As heat loads increase, the margin of error in cooling system design is more critical. Available 24-hour local service and regular preventive maintenance are required to counteract mechanical wear-and-tear. Plus, new technologies improve communication functions to provide increased maintenance support. Today's options include diagnostic and support tools, maintenance notification triggers, internal maintenance-event logging, and predictive diagnostics.

adding supplemental cooling—To effectively supplement traditional cooling and address HD areas, cooling must move closer to the heat source. Two major factors guide HD cooling technology choices: *cooling fluid* and *system architecture*.

cooling fluid: water vs. refrigerant—While water is regularly used in floor-mounted cooling, when the source of cooling moves close to sensitive electrical equipment, safety becomes a key concern.

That's why pumped *R134a refrigerant* is an ideal choice for HD applications. Because refrigerant turns into a gas when it reaches the air, a leak won't damage IT equipment or pose a safety hazard.

Pumped refrigerant solutions also provide an incremental energy savings of 25 to 35 percent, based on *kW* of cooling capacity per *kW* of heat load.

Choosing a technology platform that can scale to future needs is an important part of selecting a solution. Most major server manufacturers are working on solutions that bring refrigerant-based cooling modules into the rack to answer future heat densities of *30kW* and beyond, making refrigerant-based systems compatible with the next generation of cooling.

System Architecture: Open vs Closed

Cooling can be brought close to the load through either an open or closed architecture. In a closed architecture, the electronic and cooling equipment are located together in a sealed environment. This approach provides high-capacity cooling at the expense of flexibility and fault tolerance if failure-mode precautions are not built in.

While this approach is appropriate for small implementations, assuming failure-mode ventilation, an open-architecture approach is preferred for cooling in a data center environment. *Closed cooling* offers limited flexibility of rack combinations and often no back-up 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 15 seconds.

In *open architecture*, where modules are on or near racks, but not part of an enclosure, room air is used as a buffer in the event of a failure, making it a safer alternative in many cases. Additionally, an open architecture allows greater flexibility to re-configure as additional cooling capacity is needed.

Overhead Supplemental Cooling

An open architecture approach to high-density cooling that uses *R134a* refrigerant as a cooling medium is overhead supplemental cooling. The as-needed addition of overhead supplemental cooling units offers maximum flexibility, reliability and efficiency. Overhead units are the final building block of the adaptive cooling approach.

Overhead supplemental cooling solutions work 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. (*see figure three*).

reliability—High reliability is accomplished by placing the supplemental units close to the HD source, either in the ceiling or above the rack to supply the neces-

sary cold air to the top sections of the rack to *supplement* the air delivered from under the floor. Using *R134a* rather than water also improves reliability and eliminates the risk water poses. Additionally, since the cooling module isn't *close-coupled* with the heat load, the air in the room is used as a buffer in the event of power failure to the cooling system, providing the necessary ride-through until back-up power is restored.





figure three—overhead supplemental cooling with traditional under-floor cooling systems

energy efficiency—Overhead supplemental systems currently on the market use up to 32 percent less power than traditional floor-mounted precision air conditioners to cool *1kW* of sensible heat. One 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 three feet against zero static pressure. Since these cooling modules provide 100 percent sensible cooling, no energy is wasted controlling humidity. Additionally, refrigerant use reduces chiller capacity requirements by more than *0.2kW per 1kW* of sensible heat load. Not only does this reduce the amount of energy consumed in large implementations, it also enables additional cooling capacity without adding additional chillers or reduces the number of chillers on a new build.

maximum flexibility—Finally, overhead supplemental cooling provides the user maximum flexibility in both growth and arrangement of the data center. Since the system doesn't use any valuable raised-floor space, the user isn't constrained by 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 *16kW* modules at a time by a simple *plug-and-play* connection while the other cooling modules continue to operate.

In conclusion, adaptive cooling offers a roadmap to help managers progress toward supporting loads consistently rising in excess of predictions. This flexible approach is scalable, reliable and efficient in terms of energy, cost and space while sacrificing neither floor space or reliability. With proper planning and use of next-generation equipment, data center operating costs can be held steady or even drop, even as network availability becomes more critical. 



About the Author: Steve Madara is VP/GM, Liebert Environmental Business, (Columbus, Ohio, www.liebert.com) Emerson Network Power. He frequently provides expert comments for TechTarget's Search-DataCenter.com information portal and has presented widely on data center cooling issues. A member of ASHRAE, he holds a BSME from Bucknell Univ. and an MBA

from Wharton. **Contact:** steve.madara@emersonnetworkpower.com



© 2006, The Uptime Institute, Inc



2904 Rodeo Park Drive East, Building 100 • Santa Fe, NM 87505
 Fax (505) 982-8484 • Phone (505) 986-3900
info@uptimeinstitute.org • www.uptimeinstitute.org