

Week of June 29, 2007 Vol. 16, No. 26

FEATURES

Getting Data Centers to Chill

by Tim McCann

SGI Chief Engineer

As increasingly dense systems drive up computer room temperatures, new cooling solutions are letting water take the heat.

The latest dense rack systems are hot -- and in more ways than one.

High-density racks are growing popular because organizations need more compute performance to run today's HPC applications. But for many institutions, labs and corporations, the "server sprawl" that resulted from years of cluster scale-outs can't continue. Simply put, computer rooms are running out of room.

So IT directors are turning to faster, multi-core processors and ultra-dense board designs to consolidate sprawling server farms into tall racks stacked with dense servers and even denser blades. As compute requirements increase, more of these racks find their way into the data center. As a result, heat very quickly becomes a major worry for system administrators.

While this is less of a concern for brand-new data centers with ventilation and cooling systems specially designed to accommodate faster, hotter deployments, most facilities aren't new. They are packed with legacy systems, each with its own particular thermal management baggage.

The Trouble with Heat

Excessive heat can introduce a constellation of problems. Within the data center, IT administrators usually arrange computer systems by trying to alternate "hot" and "cold" aisles -- in other words, they try to distribute exhaust (hot) and intake (cold) aisles throughout the room to avoid hot spots that can push ambient temperatures beyond acceptable limits. Indeed, a growing problem with high, hot racks is the tendency for some systems to emit so much heated exhaust that the computer room's ventilation system simply cannot remove it all. Too often, some of that hot air bleeds into the cold aisle, recirculating back into the system and making the rack run hotter than needed.

That phenomenon brings with it significant costs:

- **Impact on hardware.** Exposure to temperatures that exceed prescribed tolerances can cause system failures and ultimately cut short the lifespan of components. It's

easy to imagine how. Consider that dense 30 Kilowatt (kW) racks are growing more popular in today's data center. A rack pulling that much energy will roughly put out the same amount of heat as 300 100-watt light bulbs. Without an effective cooling system, that rack can reach unsafe temperatures in minutes.

- **Impact on energy costs.** As computer rooms grow hotter, they cost more to power and cool. The Green Grid, a non-profit consortium dedicated to advancing energy efficiency in data centers www.thegreengrid.org, has established a metric by which organizations can estimate the total impact of a system's power demands -- including the cost of cooling the system. The Green Grid's Power Usage Effectiveness (PUE) metric is essentially a ratio that provides a helpful multiplier for understanding the "fully loaded" energy cost of a system. For instance, if a new rack of multi-core blades draws 30kW and the PUE for the data center is 3.0, then the rack's overall power consumption is 90kW. Initial studies suggest that a PUE of 2.0 appears relatively average, with some facilities reporting PUE ratios as low as 1.3 and others more than 3.0. (The closer the PUE value is to 1.0, the better.) Even as blade and server designs themselves grow more energy-efficient, commensurate increases in density still create cooling challenges that can lead to higher PUE ratios.

Because both system density and processor speeds are increasing -- suggesting that heat dissipation will remain a challenge for years to come -- vendors and end users alike have pursued a variety of approaches to cooling. Gaining popularity on all fronts is the use of water to carry heat away from systems and the data center at large.

Water: The Coolant of Choice

As many HPC veterans know, liquid cooling isn't new. It has long been used in large, custom implementations, dating back to 1964 when IBM launched its first water-cooled computing system. For most modern implementations, water remains the coolant of choice.

Water cooling typically works rather like an automobile's radiator, only in reverse. In a car, a water/glycol mixture circulating through the radiator is cooled by the incoming airflow caused by the car's forward motion. In a computing system, fans blow heat at a water coil, which takes up the heat and either carries it away or cools it before the air reaches the data center's ambient environment. Some solutions send heated water to a central chilling station which safely dissipates the heat and pumps chilled water back to the coils in the rack systems. Plumbing runs beneath the data center floor, sharing the airflow space with cabling.

Several types of water cooling implementations are available from the major system vendors, and even from companies who specialize in providing add-on cooling solutions. All of them operate on the same fundamental approach to thermal exchange, even if their particular approach differs.

Closed-loop rack airflow. Available from a variety of companies, including HP, these solutions use water-chilled coils mounted inside the rack to remove the heat gained after air passes through the system's electronics. These solutions then recirculate the cooled air back into the rack. Some closed-loop systems are installed alongside a rack, and are called "sidecar" systems. One third-party option, from Knurr CoolTherm, features coils mounted below the rack's configurable space.

Open-loop rack airflow. The choice of HPC solution providers such as SGI and IBM, this approach also uses fan-blown air to cool the system's electronics. The air then is cooled through water-chilled coils, and the cooled air is exhausted at the rear of the rack. These solutions aim to keep exhausted air just slightly warmer than the ambient data center environment, thus minimizing the chance for hot spots. Open-loop solutions are available in a variety of form factors to accommodate multiple data center architectures.

Water-Chilled Doors: Convenient and Effective

Water-chilled coil solutions, whether open- or closed-loop implementations, are evolving rapidly. As more HPC users deploy these options for thermal management, vendors are finding ways to make the solutions more efficient and -- importantly -- more convenient.

One of the most significant advances is the water-chilled door. With the first server OEM implementation pioneered by SGI in 2004, this approach contains the water cooling mechanism entirely within a hinged rear door, which can be opened at any time to enable easy access to air-movers and cables, and the water-chilled door itself. Designs like these can dramatically reduce the time required to deploy the rack, not to mention the effort needed to maintain and service it.

This approach has been proven successful since its introduction. SGI alone has fielded hundreds of HPC rack installations featuring water-chilled doors. One such installation is NASA's historic Columbia supercomputer, which is powered by 10,240 processor cores packed into 20 nodes. One of the most powerful supercomputers on the planet, Columbia has delivered 142 million hours of productive use since it was installed in October 2004. Its highest density nodes are water-cooled.

Water-chilled doors are not only convenient, they're effective. For example, SGI's third-generation water-chilled doors have been shown to remove 95 percent of the heat generated by the rack system. In real terms, this means the heat expelled by a 30 kW rack would be reduced to the heat of a 1.5 kW system. Revisiting our 100-watt light bulb analogy, that's reducing the heat from 300 bulbs to just 15.

More of the Same, Only Better

Today's water-cooled thermal management solutions are designed to solve the thermal challenges posed by systems based on air-cooled, industry-standard components. As it happens, these are the very systems that are driving the majority of the growth in the HPC sector, and their popularity will in turn prompt continued refinements in cooling solutions. Because, as we've established, the future is all about density.

Meanwhile, some companies are looking at more specialized approaches to bring liquid cooling further into the rack. Current techniques range dramatically. Some systems spray coolants directly onto heat-generating components such as CPUs. Others use liquid metal to conduct heat away from heat sources. They're all after the same thing as water cooling: to minimize the problems and commensurate costs associated with excess heat.

While one or more of these sophisticated approaches may someday see broader adoption in HPC environments, the vast majority of today's systems still must meet the aggressive

price/performance targets that allow end users to acquire and deploy the resources they need, when they need them.

So it appears that, for the foreseeable future at least, we'll be letting water take the heat.

Tim McCann is chief engineer at SGI, and one of the lead architects of SGI's water-chilled doors. At the International Supercomputing Conference in Dresden, Germany, he participated in a Birds of a Feather session on energy consumption of HPC systems. For more information, visit

http://www.supercomp.de/isc2007/index.php5?s=conference&s_nav=bofs.

Full background information on all leading HPC solution providers

- [] A351) APPRO
- [] A330) Atipa
- [] A2341) BlueArc
- [] A1314) ClearSpeed
- [] A518) Cray
- [] A1198) Crosswalk
- [] A517) DataDirect Networks
- [] A373) Dell
- [] A470) HP
- [] A335) IBM
- [] A2264) Mellanox
- [] A336) Microsoft
- [] A345) Myricom
- [] A341) NEC
- [] A1942) Panasas
- [] A343) Portland Group
- [] A339) SGI
- [] A342) Sun
- [] A2396) TotalView
- [] A1380) United Devices
- [] A471) Voltaire

For sponsorship information contact: promote@hpcwire.com

www.hpcwire.com