IDC Analyze the Future

IDC TECHNOLOGY SPOTLIGHT

The Cray CS300 Cluster's Warm Water Cooling Is at the Forefront of an HPC Industry Trend

April 2014

By Steve Conway, Earl C. Joseph, Ph.D., and Chirag Dekate, Ph.D.

Sponsored by Cray Inc.

Recent IDC surveys of the worldwide high-performance computing (HPC) market consistently show that cooling today's larger, denser HPC systems has become one of the top challenges for datacenter managers. These surveys reveal a notable trend toward liquid cooling because this method is inherently more efficient than air cooling. But not all forms of liquid cooling are equally efficient. In recent years, warm water cooling has emerged as an effective alternative to established methods that employ chilled liquids. This Technology Spotlight reviews the liquid cooling trend and the innovative use of warm water cooling in the Cray CS300 cluster supercomputer to reduce capital expense and operating costs.

Power and Cooling: A Top Issue for HPC Datacenters

In IDC's worldwide surveys since 2006, power and cooling has consistently ranked among the top 3 concerns for HPC datacenters, right behind price performance and the perennial quest for bigger budgets.

The Trend Toward Liquid Cooling

In *Power and Cooling Practices and Planning at HPC Sites,* an IDC worldwide study, the surveyed sites confirmed that cluster compute densities have increased substantially with the proliferation of multicore/manycore processors, double-dense blade form factors, wider and deeper cabinets, and related developments. Most sites employed air cooling, but all of the sites indicated they were exploring alternatives to meet their future cooling needs (see Figure 1):

- 54% of the sites were considering liquid cooling alternatives. Within this majority group, 64% of the sites were evaluating liquid cooling using chilled water and 43% were looking at liquid cooling in locus (i.e., at the rack, board, or chip package level). Sites were allowed to name more than one approach under consideration.
- Just 18.9% were evaluating air-cooled approaches. An additional 14.9% were assessing cooling based on alternative energy sources (geothermal, solar, and wind energy), while 12.2% were looking at methods fitting none of these categories.

Figure 1

Cooling Preference for Future HPC Systems

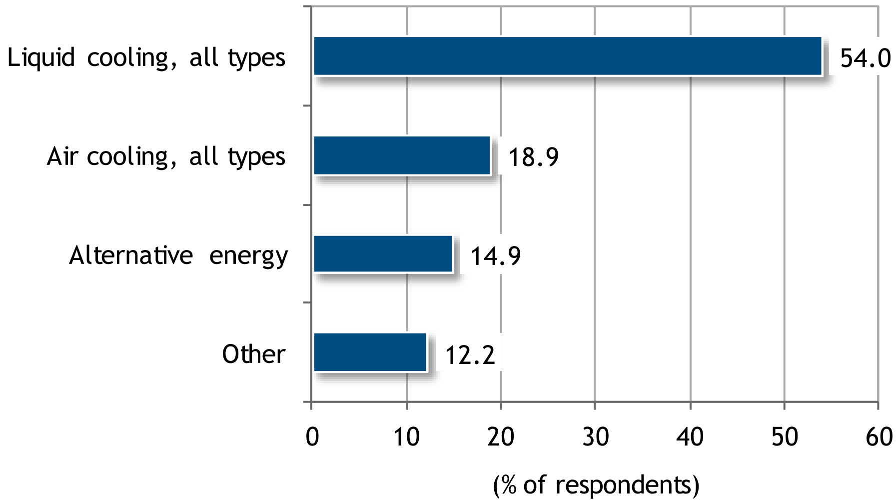

Source: IDC, 2014

The reasons for the swing toward liquid cooling are not hard to fathom:

- HPC system sizes and densities have skyrocketed. On the November 1999 list of the world's most powerful supercomputers (**www.top500.org**), the number 1 system boasted 9,632 cores and peak performance of 3.2TF. Fast forward to the November 2013 list, and the number 1 supercomputer sported 3.1 million processor cores and 54,900TF. That's a 332-fold increase in core count and a 17,156-fold peak performance gain in just over a dozen years. Average system sizes in the broader non-Top500 HPC market have undergone only slightly less spectacular growth. And when the greatly increased densities of HPC systems are taken into account, what emerges are computers that pack a whole lot more heat-generating components into much tighter confines.
- Today's largest HPC systems consume as much electricity as a small city, and their successors promise to use even more. At the same time, energy prices have risen substantially above historic levels (although prices have dropped from their 2008 highs). On average, it costs about $1 million per megawatt to operate an HPC system today. Based on that average figure, annual power and cooling costs for Top500-class supercomputers ranged from under $100,000 to $17.8 million in 2013, with costs for midrange HPC systems starting at about $10,000 per year.
- Cooling accounts, on average, for about half of the energy costs for an HPC system. It's no surprise in light of this that HPC vendors and datacenter managers alike have been seeking ways to improve cooling efficiency and lower cooling costs.
- No one cooling solution is best for every situation. Most HPC systems today are still air cooled. This means they use fans and air gaps to circulate ambient air through the computer, and the exhaust air carries heat out into the datacenter room. The thing to note here is that the exhaust air typically needs to be cooled at considerable expense by CRAC (computer room air conditioning) units in the datacenter.

- The heat-carrying capacity of air is three orders of magnitude less than that of water, requiring substantial fan power to move a given amount of heat. The thermal conductivity of air is an order of magnitude lower than that of water, requiring the ambient air to be cooled to substantially lower temperatures using mechanical chillers.

Not All Forms of Liquid Cooling Are Equally Efficient

Cooling computers with water or other liquids is not new. In the mid-1960s, early supercomputers from Control Data Corporation, Cray Research, and IBM were already using liquid cooling. In the interim, liquid cooling has morphed into various forms, each of which can be well suited to different datacenter requirements:

- Even air-cooled datacenters use liquid cooling within their CRAC and CRAH (computer room air handler) units. These units are expensive to operate and relatively inefficient because they are typically located at the outer edges of the datacenter, far from many of the heat-generating HPC system components.
- Liquid-cooled front and rear doors of racks are more efficient than CRAC/CRAH units because they operate closer to the heat-producing components. This form of liquid cooling is widely used and is appropriate for many environments. But expensive CRAC/CRAH coolant distribution units usually still need to be run in conjunction with the liquid-cooled doors. Typically, the rack-level fan power and design in these cases are very similar to those using traditional AC-cooled systems.
- Cold plates running through the servers' outer surfaces are even more efficient, but the water or other liquid they use still has to be chilled at considerable expense. In this case, the rack-level fans are significantly smaller and draw significantly less power than those in systems that are fully air cooled, or systems that are air cooled with liquid-cooled doors.
- Immersive cooling can address extreme requirements but tends to be especially expensive.

Warm Water Cooling Arrives

In recent years, a few computer makers have begun offering HPC systems whose components are cooled with water that is warm, and in some cases downright hot, yet still cooler than the HPC system's processors, memory devices, and other hot-running components. In the cooling system, the water temperatures typically range from a couple degrees above the dew point to ~40C (~104F). By bringing the warm/hot water very close to the hot components through microchannels, the system draws off enough heat to keep these components well within their recommended maximum operating temperatures. This arrangement can produce substantial energy and cost savings in several ways:

- The water flowing into the systems does not need to be chilled at substantial expense.
- The cost of buying and operating CRAC/CRAH units to cool the ambient air in the datacenter can be eliminated or greatly reduced.
- In most climates, the temperature of the water flowing out of the systems can be reduced enough by ambient free cooling to be reused in cooling the systems.
- In colder climates, this outlet water is sometimes used to heat computer buildings or other facilities.
- The rack-level component fans are significantly smaller and draw significantly less power than the fans used in air-cooled systems, depending on which blade components are cooled by warm water.

The Warm Water-Cooled Cray CS300-LC Cluster Supercomputer

Ever since the groundbreaking Cray-1 supercomputer in 1976, Cray has been recognized as an HPC industry leader in packaging and cooling technologies. In fact, the only patents issued for the Cray-1 supercomputer were for its innovative liquid cooling system. In the intervening years, Cray Research and its linear descendant, today's Cray Inc., have repeatedly helped advance the state of the art in both air-cooled and liquid-cooled supercomputer technologies.

The Cray CS300 series of cluster supercomputers is a case in point. This product series is largely derived from the company's 2012 acquisition of Appro, a successful vendor of large scale-out HPC and specialty clusters. Since the acquisition, Cray has rapidly broadened and deepened the cluster product line to address medium- to large-scale HPC capacity and capability computing and data-intensive workloads. To cover the waterfront of HPC datacenter environments and situations, the Cray CS300 cluster line is available in air-cooled, air-cooled with chilled water doors, and direct warm water–cooled versions. In addition:

- In the warm water–cooled CS300-LC cluster supercomputer, Cray employs an innovative twist. The warm water reaches the system's server nodes from a rack coolant distribution unit (CDU) that contains a liquid to liquid heat exchanger. The facilities water enters the primary water loop in the rack CDU. A double-sealed low-pressure secondary loop, with dripless quick-connects, cools the critical server components while minimizing the risk of leakage over time.
- The system's open standard nodes feature Intel Xeon processors, with options to augment these with Intel Xeon Phi coprocessors or GPU accelerators. The processors and the coprocessors, as well as the onboard memory, can be cooled with warm water.
- The warm water cooling system requires less powerful fans than the air-cooled system. Cray says that the total power for the cooling components in the warm water–cooled system (for chassis and power supply fans and liquid pumps on nodes) is 43% lower than for like-sized chassis in the air-cooled systems (for chassis and power supply fans).
- Cray reports that in practice, the warm water CS300-LC system reduces the overall datacenter power consumption by more than 50% and enables a PUE as low as 1.1.

Figure 2 displays the results of tests conducted by Mississippi State University (MSU) on its new Cray CS300-LC warm water–cooled system while running LINPACK on the 128-node system. The system has 2 CPUs, 2 accelerators, and 64GB of memory on each node. The CPUs, accelerators, and memory are direct cooled. The results show that the overall cooling efficiency, defined as heat captured by the water divided by total power consumed, remains very consistent, ~70–80%, across the entire range of supply water temperatures.

Figure 2

Facility Fluid Temperatures Versus Cooling Efficiency

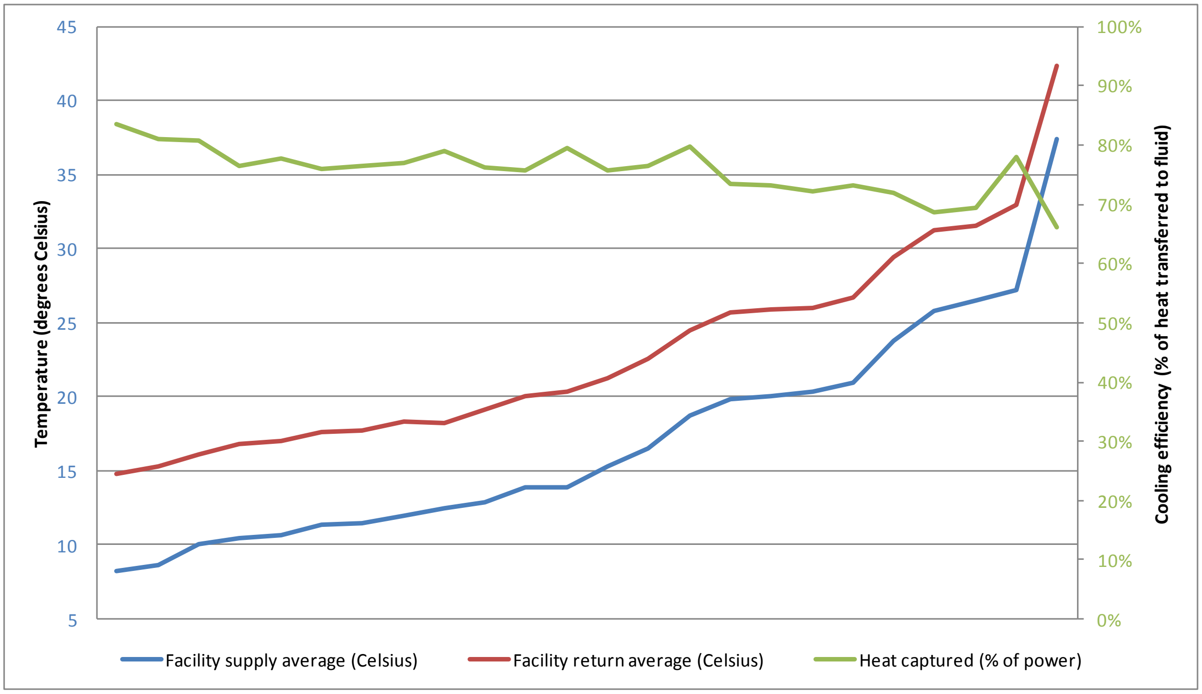

Source: Mississippi State University, 2014

A Customer Perspective: Mississippi State University

Mississippi State University installed its first water-cooled system about four years ago. It was tied into the building's chilled water system and cooled through the use of chilled rear doors. When the time came to upgrade the computing capabilities, it found that its existing chiller plant did not have enough excess capacity to handle the additional load.

Faced with the prospect of purchasing a new larger chiller for its new system, MSU quickly began studying the feasibility of using a warm water cooling technology. Because it is located in Mississippi, which has a very warm subtropical climate, it was at first unsure if it could sufficiently cool a system without using large high-capacity chillers. However, after learning that the Cray CS300-LC could be cooled with water as warm as 40C (104F), it realized that warm water cooling may be a viable solution.

MSU decided early in its procurement process that it would require some type of hardware acceleration, either Intel Xeon Phi coprocessors or NVIDIA Kepler GPUs. Ultimately, it chose to install Intel Xeon Phi coprocessors, two per node, which meant that almost half of the new system's heat load would be produced by the coprocessors. If an efficient, direct warm water–cooling technology w to be used, it would have to include cooling the coprocessors, not just the main processors and the memory. MSU concluded that the Cray CS300-LC was the only solution currently on the market that could cool all three critical components (the processors, the coprocessors, and the memory) using warm water.

By eliminating the need to purchase a new chiller system, MSU reports that it was able to reduce its capital infrastructure costs and could therefore use its limited resources to increase its computational capabilities. By eliminating the need to operate a new chiller system, it was also able to reduce its operational costs. Additionally, the Cray CS300-LC remote management capabilities allow MSU's staff to monitor and control the cooling parameters from anywhere in the world.

By running tests on its new system at various ambient temperatures and with various facilities input water temperature, MSU found that with a 100% system load, its new system has maintained an average of 76% heat capture into the dedicated facility water system. This facility water system uses a 25% glycol solution consisting only of a pumping station and a dry cooler for heat elimination. This cooling system draws less than 4kW of power. Traditional CRAC units would consume more than 10 times as much power to cool that same heat load. This results in an enormous savings in operational costs compared with MSU's older air-cooled systems.

Opportunities and Challenges

IDC believes that the availability of warm water cooling on the Cray CS300 system creates new market opportunities, along with some related challenges.

Opportunities

- **Make a larger dent in rising energy costs at HPC datacenters.** Warm water cooling can be even more efficient than chilled water cooling in many HPC datacenter environments, for the reasons enumerated previously. There are benefits for operating expenses, such as reducing electricity bills. There are also benefits for capital expenditures, such as reducing the need for buying additional CRAC units, as the Mississippi State University example demonstrates. IDC believes that warm water cooling will become an increasingly popular alternative to other liquid cooling methods.
- **Further tap into the market dominance of standards-based HPC clusters.** Clusters have become the dominant species of HPC servers and in 2013 accounted for about four-fifths of the $10.3 billion annual worldwide HPC server revenue, or more than $8.3 billion. Starting in 2012, the Cray CS300 product line, available in medium to large sizes, has given Cray access to a large chunk of the HPC cluster market for the first time. The availability of multiple cooling options for the product line — air cooling, air cooling with chilled water–cooled doors, and direct warm water cooling — increases the number of datacenter scenarios Cray can sell into. And since few vendors offer warm water cooling today, that option further sets the Cray CS300-LC cluster supercomputer apart from the rest of the pack.

Challenges

- **Not the only vendor with a warm water solution for HPC systems.** Not many vendors offer warm water cooling options for x86-based HPC clusters. Furthermore, in February 2014, IBM announced the planned sale of its x86-based server business, which includes the company's warm water–cooled clusters, to Lenovo. Any vendor switch like this creates an opportunity for other vendors to pursue the transferring accounts. The availability of a warm water cooling option in the CS300 cluster product line enhances Cray's ability to go after any prospect that might benefit from this option. The iconic Cray name should help in these pursuits, as should the company's ability to bundle CS300 cluster sales where appropriate with the Cray XC30 flagship product line, Cray storage and data management products, and the Urika big data appliance from the company's YarcData business unit.

- **A holistic view.** Traditionally, computational resources and datacenter operations each have their own separate budgets. Organizations that manage each area are also separate and often independent of each other. Cray needs to communicate the holistic benefits of efficient cooling technology to organizations that are traditionally only responsible for computational needs.
- **The need for Cray to spread the word.** The CS300 line has been part of Cray only since 2012 and already the company has added a multitude of new features, functions, and options to this product series. Keeping the worldwide HPC market informed about the expanding CS300 system choices will be an ongoing marketing and sales challenge.
- **Characterizing prospects.** A related sales challenge, given the company's expanding portfolio of products and options, is identifying adequate numbers of prospects and working with the prospects to determine which Cray product(s) and options will best meet their needs. Cray's long experience in the HPC market should serve the company well in this regard, and employees brought in with the 2012 Appro acquisition have a strong understanding of the cluster portion of the market.

Conclusion

Cooling today's larger, denser HPC systems is one of the biggest challenges faced by datacenter managers. This challenge is accelerating the use of liquid cooling, which is inherently more efficient than air cooling. Directly cooling HPC system components with warm water is an option that can produce exceptional energy efficiency and cost savings, as shown in tests conducted by Mississippi State University on its new Cray CS300-LC cluster supercomputer. IDC believes that the addition of this warm water–cooling option to the Cray CS300 series will add to the already impressive success of this Cray product line.

ABOUT THIS PUBLICATION

This publication was produced by IDC Custom Solutions. The opinion, analysis, and research results presented herein are drawn from more detailed research and analysis independently conducted and published by IDC, unless specific vendor sponsorship is noted. IDC Custom Solutions makes IDC content available in a wide range of formats for distribution by various companies. A license to distribute IDC content does not imply endorsement of or opinion about the licensee.