Impact of Hybrid Warm Water Direct Cooling Solution on Energy Efficiency of a Data Center

Nitin Sukhija, Roger L. Smith, and William B. Breckenridge III
Mississippi State University (HPC^2)
200 Research Blvd
Mississippi State, MS 39762 (USA)
{nitin@cavs, roger@hpc, trey@hpc}.msstate.edu

ABSTRACT

As the mainstream computing technology is entering into a post petascale era, the number and complexity of their computational components is on a sharp increase. Moore's law scaling of doubling transistor density is often cited in this context and has held steady. However, the Dennard scaling of voltage and transistor power scaling ended several years ago. With this advancement in the computational systems, promising exascale performance in the near future, the energy consumption of these High Performance Computing (HPC) systems has rapidly increased. Moreover, with this increase, the leakage currents are expected to grow gradually resulting in increased processor chip power requirements. As an example, the server CPU thermal design power (TDP) has grown from 64W to 115W over the last few years and is expected to grow further. With the increased pressure to pack more components per rack, the power and system densities are growing. In last few years, the heat loads per standard 19" rack have increased from few KWs per rack to well over 50KW. Recent studies report that the energy efficiency improvements will play a major part in addressing the power challenges in current (petascale) and future (exascale) computing systems. Adding to this trend the additional issues of controlling power consumption, the cooling of various components in the processor sockets and resulting hotspots in the rack and datacenter is becoming increasingly difficult, thus requiring more planning and focus for improving the overall energy efficiency of HPC data centers.

In this work, we present a prototype of a hybrid direct, warm water cooling system in the Cray CS300-LC, which is instrumented to measure temperatures, flow rates of liquid and air at inlets and outlets. Depending on the system configuration (CPU, memory and accelerators in a node), 60% to 85% of the heat is removed to warm water for a wide range of datacenter inlet water temperatures, and the rest of the heat is rejected to air or indirect cooling systems, hence the term hybrid. The prototype system was implemented and tested in the High Performance Computing Collaboratory datacenter at Mississippi State University (MSU). The system contains both Intel Xeon E5 v2 CPUs and Xeon Phi coprocessors. An acceptance testing was performed by conducting multiple runs of the LINPACK benchmark to obtain detailed measurements to understand energy efficiency and computational performance under various environmental input parameters.

The results showed a clear correlation between facility input fluid temperature and processor/coprocessor performance. This appears to be mainly due to frequency throttling in the coprocessors. However, this throttling appears to be more related to power consumption than directly to cooling (although there is a correlation to power consumption based on core temperature). In fact, indications are that the warm water cooled coprocessors actually throttle less and perform better than air cooled versions during computationally intensive work loads. Tests seem to indicate that water cooled coprocessors are outperforming air cooled versions. Due to the high thermal capacity of water, case temperature differences between the first and the second processor/accelerator are much smaller than in the air cooled version. Even with very high facility input fluid temperatures (40°C) and the system at 100% computational load, all processors and coprocessors remained well within the manufacturer's recommended temperature range. As 70% to 80% of the heat generated by the system is captured through the dry cooler without taxing the datacenter CRAC/CRAH cooling system, datacenter IT capacity can be more than doubled without expanding traditional cooling systems, thereby saving datacenter capital expenditures. In addition, the operating power requirement of a dry cooler and water circulation pumps is a small fraction of equivalent power required by a traditional CRAC unit with an electrical compressor. This results in substantial operational expenditure savings, and therefore contributes to improved datacenter power usage effectiveness (PUE). There is also further energy savings in reduced chassis fan power compared to an air cooled system. An indirect benefit of a liquid cooled implementation is a substantial reduction in system noise levels (> 10DB) compared to the air cooled version due to the reduced number and speed of cooling fans

To conclude, the preliminary results with artificially elevated fluid input temperatures and 100% computational load has proven that the LC system is stable, reliable, and remains well within all operating parameters even in the subtropical climate of Mississippi. This proves that free cooling should no longer be considered an option only for those in cool climates. Thus, free cooling a CS300-LC is possible in warm, humid climates such as in Mississippi with proper design and planning. Additional details of resiliency, remote monitoring and management of the hybrid cooling system and how the CS300-LC cooling system reacts to rapidly changing application workloads and hence cooling loads will be presented in the poster session.