

A Supercomputer with Warm Water. How Cool is That?

Peter ffoulkes, 451 Research



An energy efficient supercomputer with warm water. How cool is that?

Enlightenment has long been the ultimate pursuit of artists, philosophers, scientists, theologians and other sentient minds. Whether it is delivering the proof to support their theses, or to investigate a perplexing problem before them, they have poured a vast amount of energy into the situation. Energy has now become the problem. How do we get enough of it to feed the ever expanding goals of our ambitions and desired lifestyle, and how do we deliver and consume that energy sustainably without damaging the very environment that supports us?

Data and computationally intensive analysis has frequently been regarded as being restricted to the esoteric world of high performance computing. This is no longer the case. In the world of 'The Internet of Things', sensors, social media, and ubiquitous connectivity, the need to analyze the exponentially growing body of data is now a business imperative that far exceeds the scope and realm of the scientific research and traditional HPC community.

In the hyperscale world of global Internet companies and cloud computing providers, modern datacenters are increasingly resembling the world's largest HPC facilities from a size and energy consumption and management perspective. They may follow a few years behind the HPC trailblazers, but the core technologies required to build, operate and manage the world's leading HPC facilities often find their way into the IT mainstream and will be required for hyperscale datacenters. For energy management, that time is now.

The HPC world has hit a wall in regard to its goal of achieving Exascale systems by 2018. The current leading system on the established TOP500 methodology for evaluating the worlds most advanced computing systems rates at 34 petaFLOP/s. It consumes 17.8 MW for the machine and about 25MW with traditional air cooling. To reach Exascale would require a machine 30 times faster. If such a machine could be built with today's technology it would require an energy supply equivalent to a nuclear power station to feed it. This is clearly not practical.

Modern CPU technology is relatively energy efficient and will become even more so over the next few years as semiconductor process technology evolves, but those levels of energy efficiency will require new approaches for interconnect and memory systems to reduce consumption to more manageable levels. There are multiple initiatives in development such as silicon photonics being championed by Intel and others, and the memristor technology being championed by HP, both of which show considerable promise.

However, reducing the energy consumed by a large system is only part of the equation, it will still require cooling and conventional air cooling is very inefficient, frequently amounting to between 30 and 50 percent of total data center power consumption. As commented by Steve Hammond, Director of the National Renewable Energy Laboratory (NREL)'s Energy Systems Integration Facility (ESIF), "Air is an insulator, not a conductor. A juice glass of water has the same cooling capacity as a room full of air."

Planning for the new research facility and data center focused on a holistic "chips to bricks" approach to energy efficiency to ensure that the HPC system and data center would have a symbiotic relationship with the ESIF offices and laboratories, and to integrate the new facility into NREL's campus. HP was selected through an open RFP process to be the joint development partner for the HPC system and its cooling technology.

At the heart of NREL's new data center is the first petascale HPC system to use warm-water liquid cooling and is expected to reach an average PUE rating of 1.06 or better. Several key design specifications have led to the data center's extreme efficiency. First, high-voltage electricity (480VAC rather than the typical 208VAC) is supplied directly to the racks, which saves on power electronics equipment (cost savings), power conversions, and electrical losses (energy savings).

Secondly, the data center uses warm-water liquid cooling supplied directly to the server racks. The decision to use liquid cooling was made several years ago, before liquid-cooled systems were routinely available. This technology enables a very dense, 1.2 Pflop/s performance system with 1296 Intel Xeon processors and 576 Xeon Phi accelerators to be deployed in just 17 racks including the water cooling technology. In addition to the space efficiencies obtained the need for a large number of fans is eliminated significantly reduce the noise levels in the machine room.

There are several advantages to this approach. Water as a heat-exchange medium is three orders of magnitude more efficient than air, and getting the heat exchange close to where the heat is generated is most efficient. Also, 75°F water supplied for cooling the computers allows the data center to use highly energy-efficient evaporative cooling towers, eliminating the need for much more expensive and more energy-demanding mechanical chillers, saving both capital and operating expenses. The hydronic system features "smooth piping," where a series of 45-degree angles replace 90-degree angles wherever possible to reduce pressure drops and further save pump energy. Thus, the data center's high energy efficiency is achieved with best-in-class engineering practices and widely

A Supercomputer with Warm Water. How Cool is That?

Published on Scientific Computing (<http://www.scientificcomputing.com>)

available technologies.

By capturing the computer waste heat directly to liquid and integrating the data center into the ESIF, the data center serves as the primary heat source for office and laboratory space within the facility. Data center waste heat is also used to heat glycol loops located under an adjacent plaza and walkway, melting snow and making winter walks between buildings safer for laboratory staff.

Key elements of the HP approach to warm water cooling is the ease of installation and serviceability. Adhering to ASHRAE TC9.9 standards, the cooling system is self-contained, uses standard 'city water' and can be assembled and maintained by data center staff without the need for specialized plumbing skills. Perhaps the most significant differentiation in HP's design is in the area of serviceability. Compared with most current warm water cooling systems that take water directly to the electronic components, HP keeps all of the water in the rack itself with no direct contact with the electronics. Heat pipes transfer the energy from components to plates on the side of each tray, which have a direct metal-to-metal transfer to the circulating water cooling system in the rack. This 'dry disconnect' design delivers the equivalent serviceability of a standard blade architecture, a significant advantage over the 'drip less' connectors found in the majority of existing warm water cooling products.

As the IT industry transitions inexorably to cloud computing architectures, private datacenters, hosting, and cloud providers will frequently adopt a dense, hyperscale design, where advanced cooling and energy efficiency will be a prerequisite. For commercial deployment, ease of installation and serviceability will be a requirement, and the ability to use a technology that can allow an existing facility to be retrofitted will become a distinct advantage.

The joint development between NREL and HP of the Peregrine supercomputer and the ESIF facility is an industry leading example of the type of data center technologies that will become increasingly common in the second half of this decade.

Peter ffoulkes is Research Director, Server Infrastructure and Software, Cloud Computing at 451 Research.

Source URL (retrieved on 05/30/2016 - 10:00pm):

<http://www.scientificcomputing.com/blogs/2014/06/supercomputer-warm-water-how-cool>