



Energy-efficient CoolMUC-3 Cluster comes to LRZ in Germany


insidehpc.com/2017/11/energy-efficient-coolmuc-3-cluster-comes-lrz-germany/

ralphwells

November 21, 2017

An new energy-efficient supercomputer called CoolMUC-3 has been deployed at the Leibniz Supercomputing Centre (LRZ) in Germany. The HPC cluster features Intel's many-core architecture and warm-water cooled Omni-Path switches.

Developed by MEGWARE, CoolMUC-3 is the result of an established partnership between LRZ and Megware. Similar to the delivery of the first CoolMUC in 2011 we have managed to raise the bar for energy efficiency in HPC once more by developing new, innovative technologies," said Axel Auweter, CTO at Megware.



CoolMUC-3 node developed by Megware

LRZ's procurement targets for the new cluster system were:

- To deploy state-of-the art high-temperature water-cooling technologies for a system operation that avoids heat transfer to the ambient air of the computer room.
- To supply a system to its users that is suited for processing highly vectorizable and thread-parallel applications, and
- To provide improved scalability across node boundaries for strong scaling.

The system's baseline installation consists of 148 computational many-core Intel "Knight's Landing" nodes (Intel Xeon Phi™ 7210-F hosts) connected to each other via an Intel Omni-Path high performance network in fat tree topology. The theoretical peak performance is 400 TFlop/s and the LINPACK performance of the complete system is 255 TFlop/s. A standard Intel Xeon login node is available for development work and job submission.

CoolMUC-3 comprises of three warm-water cooled racks, using an inlet temperature of at least 40 °C. A very high fraction of waste heat deposition into water is achieved by deployment of liquid-cooled power supplies, Omni-Path switches and thermal isolation of the racks to suppress radiative losses. In addition, the liquid cooled racks operate entirely without fans. A complementary rack for air-cooled components (e.g. management servers) uses less than 3% of the systems total power budget. With 4.96 GFlops/Watt (according to the strict Green500 level-3 measurement methodology) CoolMUC-3 is one of the most efficient x86 systems worldwide.

LRZ has a long-standing history in energy-efficient, warm-water cooled HPC systems. In the long run, we want to get rid of inefficient air-cooling completely. With CoolMUC-3 we take an important step into this direction.” states Prof. Dr. Dieter Kranzlmüller, Chairman of the Board of Directors at LRZ.

CoolMUC-3 is particularly suitable for parallel applications that minimize memory consumption by hybrid programming, e.g. with MPI and OpenMP and then take advantage of the vector units after optimizing the data location. For codes that require use of the distributed memory paradigm with small message sizes, the integration of the Omni-Path network interface on the chip set of the computational node can bring a significant performance advantage over a PCI-attached network card. The actually observable memory bandwidth of the high bandwidth memory yields approximately 450 GB/s per node.

Sign up for our insideHPC Newsletter