CASE STUDY

High Performance Computing
Intel® Omni-Path Architecture

Oakforest-PACS Now Japan's Fastest Supercomputer

Joint Center for Advanced High Performance Computing (JCAHPC) acquires Fujitsu machine built on Intel® Xeon Phi™ Processor 7250 and Intel® Omni-Path Architecture

FUJITSU

JCAHPC's Oakforest-PACS Achieves Top 10 List:

- University of Tokyo and University of Tsukuba form JCAHPC to build Oakforest-PACS
- Oakforest-PACS performance: placed #7 in June 2017 Top500 at 13.5 petaFLOPS
- 8,208 nodes of Intel® Xeon Phi™ Processor 7250 and Intel® Omni-Path Architecture (Intel® OPA) fabric
- Second largest Intel Xeon Phi Processor deployment and largest Intel OPA deployment in the world

JCAHPC

Japan's University of Tokyo and University of Tsukuba collaboratively created the Joint Center for Advanced High Performance Computing (JCAHPC) to centralize their supercomputing resources, giving them increased computational capabilities for advanced research across a wide range of disciplines. The more than 8000-node machine, built on Intel® Xeon Phi™ Processors and Intel® Omni-Path Architecture (Intel® OPA) is the fastest supercomputer in Japan and 7th in the world, according to the Top500.

Challenge

Supercomputers are costly, but they are necessary resources to support innovative science and engineering in the world's leading institutions of academia. Without HPC, these institutions lose their ability to stay at the competitive forefront of research, and with it, the funding that supports much of that research. In public sector HPC, dollars follow FLOPS.

The University of Tokyo (U Tokyo) and University of Tsukuba (U Tsukuba) in Japan had each been operating separate supercomputing clusters for years, upgrading to new technology every few years to continue to support the ongoing research and discovery of nationwide computational scientists and engineers. Rather than continue to duplicate budgets across different institutions and limit their purchasing capability, they pooled their resources to build a record-shattering supercomputer.

Solution

"Instead of continuing to individually build HPC resources at U Tokyo and U Tsukuba," said Professor Boku of U Tsukuba, "the two universities created the Joint Center for Advanced High Performance Computing (JCAHPC) to purchase a supercomputer that would serve the needs of both our universities with more capacity than we could purchase alone."

"JCAHPC was established in 2013," added Professor Hanawa of U Tokyo. "Together, we designed Oakforest-PACS to be the next-generation and fastest supercomputer in Japan. Today, ten faculty from the two universities share operation and management of it." In November 2016, the system became Japan's fastest HPC resource, displacing Japan's K computer, and #7 in the world according to the June 2017 Top500 list.

"Oakforest-PACS was built to meet two goals:" explained Boku, "to be the fastest supercomputer in Japan and to support next-generation general-purpose computing compatible with an x86 architecture. We weren't expecting to use accelerators. For the network, we focused on performance."

Fujitsu was one of the OEMs that responded to the RFP. To meet the performance requirements for Japan's new supercomputer, Fujitsu designed Oakforest-

PACS around Intel Xeon Phi Processor 7250 with Intel OPA as the fabric. The Intel Xeon Phi Processor (codename Knights Landing), with 68 cores, 3 TFLOPS of computing performance, large onboard memory resources, and Intel® Advanced Vector Extensions 512 (Intel® AVX-512), was designed for efficient parallel processing. Intel OPA, with the PSM library support for efficient MPI routing, is designed for the kinds of codes that Oakforest-PACS would run. "This is an ideal combination for high throughput of MPI codes, which are common in the codes of science and engineering," commented Katsumi Yazawa of Intel.

JCAHPC was aware of the performance the MIC architecture and the enhancements that Intel Xeon Phi Processors offered. "For several years, developers and computational scientists at the universities have been developing and testing parallel codes using Intel AVX-512 to run on COMA, the largest Intel Xeon Phi coprocessor (Knights Corner) cluster in Japan, located at U Tsukuba," stated Boku.

"U Tokyo also introduced a small cluster with the same coprocessor" explained Hanawa. "In both universities, the advanced users optimized and evaluated their codes to make them ready for Oakforest-PACS. So, we understand the work we have to do to take advantage of the smaller cores of the Intel Xeon Phi Processor in order to achieve maximum performance from the system. We realize the memory architecture of the two processors are also different, but we have good experience with this architecture."

Result

Oakforest-PACS is the second largest deployment of Intel Xeon Phi Processors in the world and the largest deployment of Intel OPA. It placed #7 on the June 2017 Top500 with a LINPACK score of 13.5 petaFLOPS. It also placed #21 on the Green 500 ranking for June 2017, illustrating its power-efficient design overall.

Just released into production in December 2016, early researchers are seeing good performance on Oakforest-PACS. "We have been running Lattice Quantum Chromodynamics (QCD) codes and some first-order optical material science simulations on Oakforest-PACS," said Boku. "These are very quantum-level calculations and simulations. With its hundreds of mesh points, the QCD code demands a lot of memory bandwidth rather than CPU performance. We are seeing the benefits of Intel Xeon Phi Processor's MCDRAM for running Lattice QCD," he added. U Tsukuba is at the heart of Lattice QCD work in Japan, with more than 25 years of experience. The university hosts an Intel® Parallel Computing Center, which targets Lattice QCD optimization.

"Sustained performance from ARTED has reached 25 percent of theoretical peak," stated Hanawa. "That's about 700 gigaFLOPS per chip, which is much faster than what the standard Intel Xeon processor can sustain." U Tokyo early users are running atmosphere and ocean coupling and earthquake simulations using GAMERA/GOJIRA codes. Both universities expect to have published research from Oakforest-PACS later in 2017.

Solution Summary

Oakforest-PACS, a 25 petaFLOPS (peak) and 13.5 petaFLOPS (LINPACK) supercomputer, was built by Fujitsu for JCAHPC, a collaboration of U Tokyo and U Tsukuba. It is the fastest supercomputer in Japan, and number seven in the world in performance. The system is powered by 8,208 nodes of Intel Xeon Phi Processor 7250 with the Intel OPA fabric. Oakforest-PACS was released into production in December of 2016, so it is still running research that could be revealed later in 2017.

Where to Get More Information

Learn more about Oakforest-PACS at http://jcahpc.jp/eng.

Learn more about Intel Omni-Path Architecture at http://www.intel.com/hpcfabrics.

Learn more about Intel Xeon Phi processors at http://www.intel.com/content/www/us/en/products/processors/xeon-phi.html.

Solution Ingredients

- 8,208 nodes of Intel Xeon Phi 7250
- Intel Omni-Path Architecture Series 100 Host Fabric Interface Adapters
- Intel Omni-Path Architecture 100 Series Edge Switches