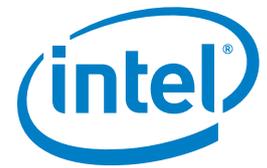


Case Study

Intel® Xeon® Processor E5 Family, Intel® Cluster Studio XE
High-Performance Computing



Delivering High-Speed Supercomputer Services

Kyoto University builds a 1,202-socket cluster supercomputer to deliver advanced information services to research institutes throughout Japan using the Intel® Xeon® processor E5 family



Academic Center for Computing and Media Studies, Kyoto University

Yoshida-Honmachi, Sakyo-ku, Kyoto
Established April 1964

Activities:

Research and development aimed at advanced applications of IT platforms and media and providing, operating, and administering IT infrastructure within the university
<http://www.media.kyoto-u.ac.jp/>

Challenges

- Delivery of high-speed supercomputer services to researchers and research institutions inside and outside academia
- Installation of cluster supercomputer with superior node and network performance
- Lower heat generation and improved power efficiency

Solutions

- Cluster supercomputer using Intel® Xeon® processor E5 family
- Intel® Cluster Studio XE suite of tools for MPI developers

Highly Versatile Supercomputer for Wide Variety of Research Applications

The Academic Center for Computing and Media Studies at Kyoto University (ACCMS) undertakes research and development aimed at advanced applications of IT platforms and media. It feeds the results of this research back into enhancements to the educational environment. As one of Japan's national centers for shared access to IT infrastructure, it also provides sophisticated computing services to researchers and other users at other universities and research institutions throughout Japan. Director of ACCMS, Professor Hiroshi Nakashima, PhD, described the center by saying, "Our two key roles are to provide IT infrastructure across Kyoto University and to conduct academic research into the building of advanced IT platforms for the future. In terms of joint research, we have an important mission as part of a network of eight computing centers at major national universities." ACCMS is actively involved in joint research and other collaborations with both private-sector companies and research institutions, with research and development organized into five departments. Besides the existing Department of Network Research, Department of Computing Research, Department of Educational Support, and Department of Digital Content Research, the center has also added a new department

called the Collaborative Research Laboratories to work on creating new IT platforms and services.

At the Department of Computing Research, to which Professor Nakashima belongs, work on supercomputers extends beyond research and development of hardware and software, and includes system operation and user support by technical staff in the IT Infrastructure Group of the IT Section. The shared access service supplies computing services, not only to researchers at Kyoto University, but also to academic researchers, private-sector businesses, and other users

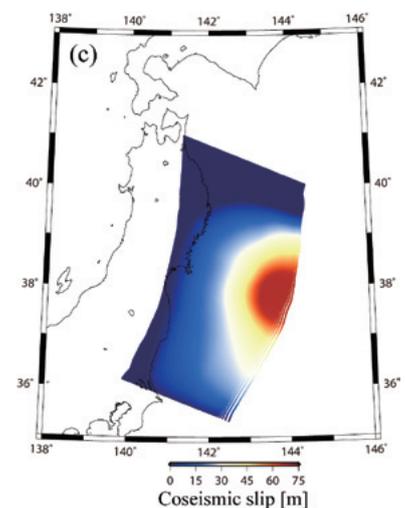
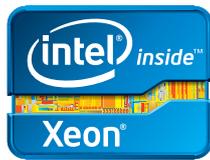


Figure 1. Computed image from simulation of major earthquake

Extracting Maximum Performance from Large Cluster Supercomputers Intel Xeon Processor Family



“Adoption of the Intel Xeon processor family provides users with a high-speed computing environment, with network performance that improves in proportion with node performance. It also reduces power consumption and dramatically enhances IT service quality.”

*Professor Hiroshi Nakashima, Ph.D
Director of Academic Center for
Computing and Media Studies,
Kyoto University*



*Professor Hiroshi Nakashima, Ph.D
Director of Academic Center for
Computing and Media Studies, Kyoto
University*

throughout Japan. Applications include large scientific and technical computations, computational chemistry, structural analysis, statistical processing, and visualization. Professor Nakashima commented that, “The supercomputers made available by ACCMS can be used by all sorts of different research institutions in a variety of fields and disciplines. This means, rather than being mission-oriented, our service must maintain a highly versatile approach. The majority of use comes from science and engineering research, with examples of large simulations on which the center has collaborated including simulations of the cycle of major earthquake events (Figure 1), and particle simulations of plasma environments in space (Figure 2).”

ACCMS installed its first supercomputer in 1985. While the early machines were vector supercomputers, they have switched to more versatile scalar machines since 2004. Discussing their policies for supercomputer installation and operation, Professor Nakashima said, “To keep up with other academic institutions around the world, it is important that we progress in line with the latest developments. In order to do this, we configure systems as Linux* clusters, an architecture that is in widespread use internationally. Along with its versatility, other major advantages of this approach include performance, cost, and application development efficiency.”

Evaluation of High-Speed I/O Bus Using PCI Express* 3.0 and Adoption of Intel® Xeon® Processor E5 Family, and Implementation of Compilers and Other Developer Tools

To provide its university and other users with an advanced computing environment, ACCMS has regularly updated its supercomputers every few years. In 2012, with the HPC server installed in 2008 coming up for replacement, the center undertook evaluation work to select its latest supercomputer system. After comparing and testing machines from a number of vendors, they configured a large cluster system using an HPC server fitted with the Intel Xeon processor E5 family. Explaining their reasons for choosing Intel processor-based CPUs, Professor Nakashima said, “An

essential requirement when selecting the processor was that it use the highly versatile x86 64-bit architecture. This led us to the Intel Xeon processor E5 family, which allows high-speed exchange of data between nodes using a network I/O bus that supports PCI Express 3.0.”

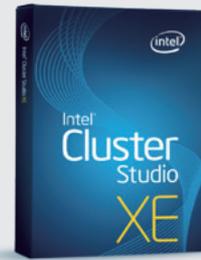
PCI Express 3.0 is a high-speed bus transmission standard able to achieve transfer speeds up to 8GT/s (giga transfers per second) per lane. At the time of writing, the Intel Xeon processor E5 family was the only processor that supported PCI Express 3.0. Since it allows faster network access than the PCI Express 2.1 standard that currently is the mainstream (maximum transfer speed: 5GT/s), the advantages of PCI Express 3.0 cannot be overestimated.

This system upgrade also included the installation of the Intel Cluster Studio XE suite of tools for message-passing interface (MPI) developers. Intel Cluster Studio XE is a package of tools for HPC clusters, including C/C++ and Fortran compilers, performance analysis tools, an MPI library, and MPI application performance analysis tools. Professor Nakashima explained the reasons for selecting Intel Cluster Studio XE by saying, “We were impressed by the software’s high degree of affinity with Intel processors and its compatibility with other x86 processors. Benchmark testing of the various tools also gave favorable results that were at a level we found satisfactory. The usability of the tools was also attractive, as was the extensive range of software included in the suite.”



Figure 2. Large-scale particle simulation of interaction between ion engine and plasma

Intel Cluster Studio XE



Intel Cluster Studio XE is a suite of development tools for MPI applications. Combining a number of highly reliable tools, including Intel's cluster software, advanced threading/memory consistency detection, and performance profiling. The software delivers significant improvements in the performance and scalability of cluster applications.

Products Included in Intel Cluster Studio XE

- **Intel® Composer XE**

Includes C/C++ and Fortran compilers, performance libraries (for numerical calculation [Intel® MKL], graphics processing [Intel® IPP], and a multithreading library [Intel® TBB]).

- **Intel® VTune™ Amplifier XE**

An analysis tool for the rapid diagnosis of performance bottlenecks. Use of templates lets you retrieve the information you need with a few mouse clicks. The intuitive user interface keeps operation simple.

- **Intel® Inspector XE**

A utility with advanced functions for detecting memory and threading errors. Supports the dynamic detection of memory problems such as memory leaks or corruption, and multithreading errors such as data conflicts or deadlocks.

- **Intel® MPI Library**

An MPI library with the scalability to handle more than 90,000 processes. Enhances the execution of applications on Intel® platform clusters.

- **Intel® Trace Analyzer/Collector**

A performance analysis tool for MPI applications. Supports event-based tracing of applications executing in parallel. Collected trace data are displayed graphically to simplify the identification of performance bottlenecks.

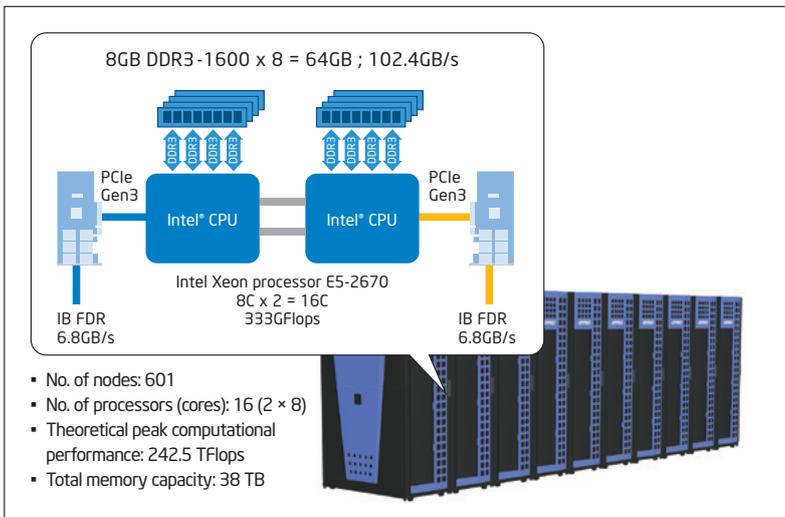


Figure 3. Configuration of Laurel Subsystem with 601 nodes

1,202-Socket Subsystem and High-Memory-Capacity Subsystem with 1.5TB per Node

The new supercomputer system entered full operation. The system was made up of one massively parallel processor (MPP) system and two cluster systems, each comprising an InfiniBand* network and HPC server fitted with the Intel Xeon processor E5 family. The peak computational performance of the overall system was 553.9 TFlops.

The Laurel subsystem has a high degree of compatibility with PC clusters, comprising 601 nodes with 64 GB of memory and 16 cores per node. It has a peak computational performance of 242.5 TFlops and total

memory capacity of 38 TB (Figure 3). The Cinnamon subsystem has 16 nodes with 1.5 TB of memory and 32 cores per node. Despite the small number of nodes, the large amount of memory per node means it will be used for applications that demand a large memory capacity. It has a theoretical peak computational performance of 10.6 TFlops and total memory capacity of 24 TB (Figure 4).

Improvements in Node Performance and Power Efficiency Deliver High-Speed Analysis at Low Cost

With 7.9 times the computational performance, 6.1 times the memory capacity, and 5.7 times the overall physical capacity of its predecessor, the new system represents

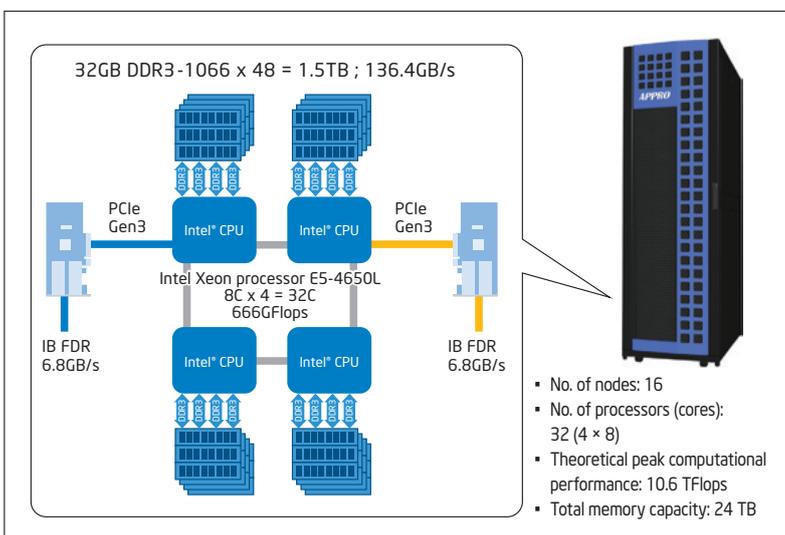


Figure 4. Configuration of Cinnamon Subsystem with 16 nodes, each with 1.5 TB of memory

a major step up in scale. The two clusters fitted with the Intel Xeon processor E5 family also deliver significant improvements in performance and power consumption. The 601-node Laurel subsystem more than doubles node performance while reducing power consumption by more than half. This corresponds to a roughly six-fold improvement in power efficiency. Referring to another major success, Professor Nakashima said, "Of particular significance is how the benefits of using PCI Express 3.0 have seen network performance improve roughly in proportion with node performance."

While the Cinnamon subsystem, with its large memory capacity, has only improved node performance by about 20 to 30 percent, power consumption has been cut to one-tenth that of the system it is replacing. Also, the ability of the entire system to fit into a single rack means it takes up only one-tenth as much space as the previous system. This cut installation costs significantly.

The new supercomputer system also brings benefits for research. The improvement in underlying performance means users can obtain the results of even large and complex calculations quickly and cost-efficiently. In an academic context, the faster speed will prove valuable because it allows calculations to be executed for more parameters than would otherwise be possible in the limited time available. The new system also has benefits for ACCMS in its role as a service provider. In particular, Professor Nakashima notes that, "Under current operating practices in which usage fees are calculated based on the amount of power that users consume, a major benefit is the near six-fold improvement in power efficiency, which means that we can provide users with roughly six times as much computing capacity within the allocated budget."

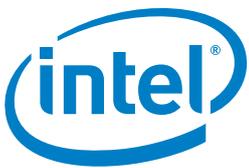
Regarding their future plans, ACCMS has already decided to install an additional new supercomputer fitted with the next generation of Intel Xeon processors in 2014. The new subsystem is expected to

have a peak computational performance of 400 TFlops. Combined with the existing system, this will result in a supercomputer with performance approaching the 1 PFlops range. Looking to the future, the center is currently at the stage of testing the next generation of technology, with a presumption that the processors used will be made by Intel. "I have been entirely satisfied with the Intel Xeon processor E5 family and Intel Cluster Studio XE products used in our new system," said Professor Nakashima. "I look forward to Intel's ongoing technological innovation and its development of fascinating products."

For its part, through ongoing technological innovations in its development tools and the Intel Xeon processor, Intel intends to contribute to further enhancements to the IT infrastructure that ACCMS is seeking to build.

Find the solution that's right for your organization. Contact your Intel representative, visit Intel's Business Success Stories for IT Managers (www.intel.com/itcasestudies), or explore the Intel.com IT Center (www.intel.com/itcenter).

For more information on the Intel Xeon processor, visit <http://www.intel.co.jp/xeonE5/>



Performance tests and ratings contained within this document are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, visit Intel Performance Benchmark Limitations. Intel does not control or audit the design or implementation of third party benchmarks or Web sites referenced in this document. Intel encourages all of its customers to visit the referenced Web sites or others where similar performance benchmarks are reported and confirm whether the referenced benchmarks are accurate and reflect performance of systems available for purchase.

When used on compatible microprocessors, Intel® compilers will not necessarily achieve the same level of optimization as achieved on Intel microprocessors. This includes optimization for the Intel® Streaming SIMD Extensions 2 (Intel® SSE2), Intel® Streaming SIMD Extensions 3 (Intel® SSE3), and Intel® Supplemental Streaming SIMD Extensions 3 (SSSE3) instruction sets, as well other optimization. Intel assumes no responsibility for the provision, functions, or effects of optimization on microprocessors not made by Intel. The microprocessor-specific optimization performed by this product is intended solely for Intel microprocessors. Certain optimization that is not specific to the Intel® microarchitecture is reserved for use with Intel microprocessors. For more information about the specific instruction sets to which this disclaimer applies, please refer to the user reference guides for the respective products.

This paper is for informational purposes only. THIS DOCUMENT IS PROVIDED "AS IS" WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, NON-INFRINGEMENT, FITNESS FOR ANY PARTICULAR PURPOSE, OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION, OR SAMPLE. Intel disclaims all liability, including liability for infringement of any proprietary rights, relating to use of information in this specification. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted herein.

Intel, the Intel logo, Xeon, and Xeon Inside are trademarks of Intel Corporation in the U.S. and other countries.

Microsoft and Windows are trademarks of Microsoft Corporation in the U.S. and other countries.

* Other names and brands may be claimed as the property of others.

Copyright © 2012 Intel Corporation. All rights reserved

327412 -001US
JPN/1207/PDF/SE/ESS/FN