



Case Study

High-performance computing

Quad-Core Intel® Xeon® processors

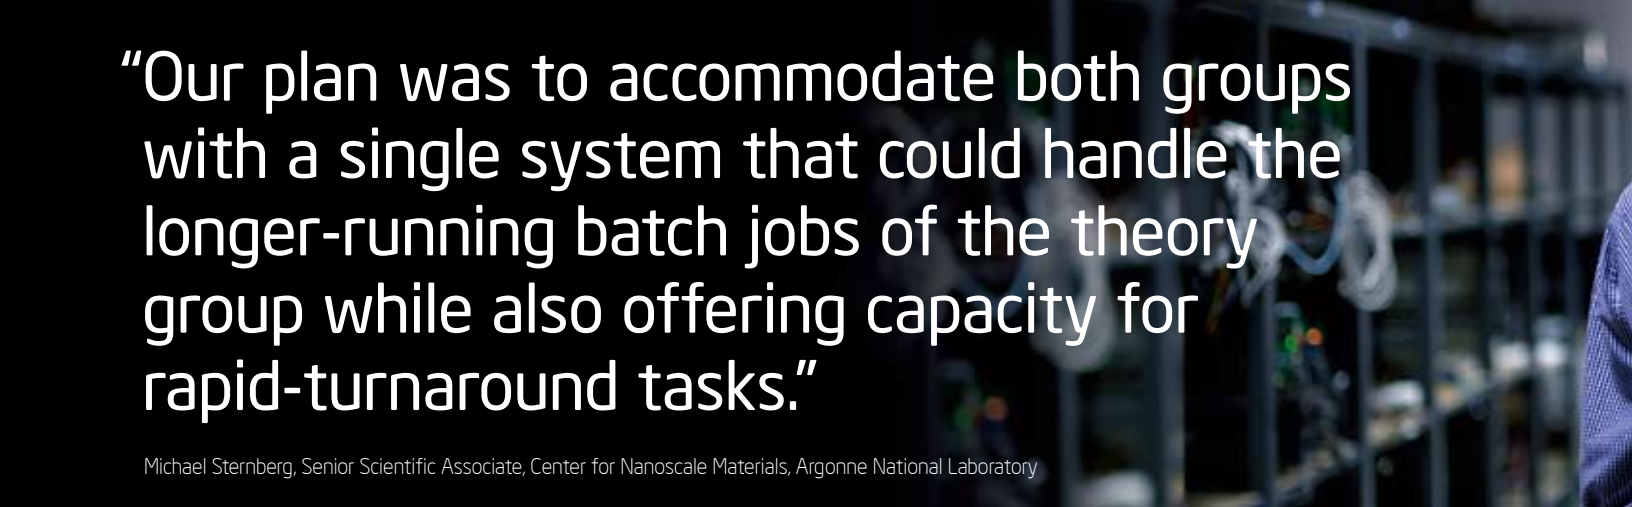
Intel® Multi-Core Performance Accelerates Nanoscience Research

The Center for Nanoscale Materials (CNM) at the U.S. Department of Energy's Argonne National Laboratory installed a dense, high-performance computing cluster using Quad-Core Intel® Xeon® processors and motherboards developed by Intel, Supermicro, and SGI to deliver maximum performance while reducing power consumption

-
- Challenges**
- Deliver the processing performance required to support world-class scientific research on nanoscale materials
 - Minimize power consumption and cooling costs by creating a dense, high-performance computing cluster
 - Adopt software tools to help accelerate development of customized scientific applications

-
- Solutions**
- The CNM team built a 1,152-core high-performance computing cluster with Quad-Core Intel® Xeon® processors plus motherboards developed by Intel, Supermicro, and SGI, delivering the extreme performance required with low power consumption and a small footprint
 - CNM application developers use the Intel® Fortran Compiler for Linux* OS and the Intel® Math Kernel Library to accelerate development of customized scientific applications

-
- Benefits**
- By implementing a high-performance computing cluster with Quad-Core Intel Xeon processors and motherboards developed by Intel, Supermicro, and SGI, CNM is maximizing the quantity and quality of research performed at the facility
 - The CNM team has created a dense cluster that delivers outstanding performance with low power, cooling, and real estate costs; with 16 CPU cores per rack unit, the Supermicro servers deliver twice the density of other high-performance computing systems considered
 - Using Intel® software tools, CNM application developers have improved the efficiency of applications by 20 to 30 percent



“Our plan was to accommodate both groups with a single system that could handle the longer-running batch jobs of the theory group while also offering capacity for rapid-turnaround tasks.”

Michael Sternberg, Senior Scientific Associate, Center for Nanoscale Materials, Argonne National Laboratory

“It’s important to get the research done and to produce scientific results. We want to use software that enables our researchers to spend more time on their primary work and less time adjusting software.”

Michael Sternberg
Senior Scientific Associate
Center for Nanoscale Materials
Argonne National Laboratory

Created in 2006, the Center for Nanoscale Materials (CNM) at the U.S. Department of Energy’s Argonne National Laboratory offers researchers from around the world access to cutting-edge technology for studying ultra-small materials and developing innovative products—ranging from strong but lightweight fabrics to new solar cells. The CNM team responsible for creating the computational infrastructure wanted to build a dense, high-performance computing (HPC) cluster that could provide the processing power necessary for computational nanoscience and the ability to capture and analyze experimental data in real time. The team created a 1,152-core HPC cluster with Quad-Core Intel® Xeon® processors and motherboards developed by Intel, Supermicro, and SGI, while also adopting Intel® software tools to accelerate development of customized scientific applications.


Challenges

Built from the ground up, the CNM facility enables scientists and engineers to conduct a wide variety of projects in a single location. “The Center for Nanoscale Materials greatly simplifies the logistics of nanoscience research by housing so many different state-of-the-art experimental techniques under one roof,” says Larry Curtiss, acting group leader at the CNM. “Previously,

researchers had to conduct segments of research at a variety of different labs. Here, they can take a sample from one lab to another without leaving the building. At the same time, they can gain access to Argonne’s Advanced Photon Source nanoprobe beamline to help them characterize extremely small structures.”

To support the intense experimental data analysis required for that work, the CNM needed to build an HPC cluster. “We wanted to assemble a cluster that could support the work of both our theory group and our experimental scientists,” says Michael Sternberg, senior scientific associate at the CNM. “Our plan was to accommodate both groups with a single system that could handle the longer-running batch jobs of the theory group while also offering capacity for rapid-turnaround tasks.”

It was crucial for the cluster to deliver the performance required to capture and analyze a large amount of data in real time. “Some experiments produce enormous amounts of data, and researchers need to analyze that data in real time so they can reposition samples, adjust instruments, and fine-tune their experiments,” says Sternberg. “The cluster had to deliver extreme compute power and fast data transfer.”



Sternberg's team also wanted to construct a cluster that would minimize power, cooling, and real estate costs. "Research is just beginning at the CNM. We wanted to be sure that the infrastructure accommodates not only our current requirements but also our future needs," says Sternberg. "Our plan was to create a dense cluster that could reduce the physical footprint of hardware while keeping the power and cooling costs down."

The research conducted at the CNM relies on both commercial software and custom-developed applications designed for nanoscience research. The CNM application development team wanted to find ways to accelerate development while ensuring that they could fully capitalize on the power of this new cluster. "It's important to get the research done and to produce scientific results," says Sternberg. "We want to use software that enables our researchers to spend more time on their primary work and less time adjusting software."

Solutions

After thoroughly evaluating processor and motherboard technologies from a variety of vendors, the CNM team chose to work with Intel. Unlike some of the other HPC platforms, the Intel® architecture supports a wide range of scientific applications used by researchers, including The MathWorks MATLAB® software, Gaussian® chemistry software, VASP® molecular dynamics software, and Dacapo® materials science software.


Spotlight on the Center for Nanoscale Materials at Argonne National Laboratory

The U.S. Department of Energy's Argonne National Laboratory is a renowned center for finding creative solutions to pressing national problems in science and technology. The Center for Nanoscale Materials, created in 2006 through a joint partnership of the U.S. Department of Energy and the State of Illinois, provides access to some of the most advanced nanoscience technology in the world. Research performed at the facility will help lead to advances in medicine, electronics, manufacturing, and energy.

The CNM team built a 1,152-core HPC cluster using Quad-Core Intel® Xeon® 5355 processors housed in Supermicro 1U Twin* systems. With the quad-core architecture, the CNM team can maximize the cluster's processing power while conserving power consumption, cooling requirements, and real estate. The Intel quad-core architecture delivers higher compute power than single- or dual-core architectures with much lower power consumption. As a result, the CNM could create a dense cluster that leaves room for future growth. The CNM team is considering the incorporation of 45-nm Intel Xeon 5400 series processors into the cluster when they become available so the CNM can maintain high performance per watt in the future.

"Previously, researchers had to conduct segments of research at a variety of different labs. Here, they can take a sample from one lab to another without leaving the building. At the same time, they can gain access to Argonne's Advanced Photon Source nanoprobe beamline to help them characterize extremely small structures."

Larry Curtiss
Acting Group Leader
Center for Nanoscale Materials
Argonne National Laboratory



Each of cluster's 1U systems houses a total of four quad-core processors—two on each of the two Supermicro X7DBT-INF* motherboards, which were developed by Intel, Supermicro, and SGI. With two motherboards in a single 1U enclosure, the CNM team can reduce the power requirement by half compared with other system designs—each two-board system requires only a single high-efficiency power supply. By using 16 cores in a single server, the CNM team could build a much denser cluster with the Intel® architecture than they could with other systems.

The motherboard provides the memory and networking capabilities for the rigors of nanoscience research, along with the flexibility to support future growth. "The ability to use a large amount of memory per core is essential for the memory-hungry nanoscience applications," says Sternberg. The new motherboard accommodates up to 32 GB of memory per node across 8 Dual In-line Memory Module (DIMM) slots. The CNM team currently uses 2 GB per core, but the team can add memory easily in the future if necessary. In addition, the CNM team takes advantage of the motherboard's included InfiniBand* controller, which delivers the high-performance data transfer required for research.

Intel® Software Tools Improve Application Efficiency up to 30 Percent

CNM application developers selected the Intel® Fortran Compiler for Linux* OS and the Intel® Math Kernel Library to help accelerate the development process for customized scientific software and to optimize applications running on the Intel® multi-core architecture. "There are plenty of open-source libraries and compilers available, but to use those tools you need to spend time fine-tuning them to suit your needs," says Sternberg. The Intel® software tools provided the CNM team what it needed to help accelerate development right away.

The Intel software development group also helped the CNM team optimize its code for the new cluster. By teaching the CNM developers how to rearrange loops and improve memory throughput, the Intel group helped the CNM developers achieve higher performance with their applications. With the Intel software tools and help from the Intel software team, CNM developers are taking full advantage of the compute capacity at the heart of the Intel multi-core architecture. As a result, the CNM developers have improved the efficiency of applications by 20 to 30 percent.

"The new HPC cluster has a 12-teraflop peak capacity that will undoubtedly help attract researchers to the center for conducting computational and experimental nanoscience."

Michael Sternberg
Senior Scientific Associate
Center for Nanoscale Materials
Argonne National Laboratory

Intel Multi-Core Architecture Helps Produce More Research and Higher-Quality Science

With the new HPC cluster in place, CNM is well prepared to accommodate the rigorous computational needs of nanoscience. “The new HPC cluster has a 12-teraflop peak capacity that will undoubtedly help attract researchers to the center for conducting computational and experimental nanoscience,” says Sternberg.

At the same time, the ability to pack more processing power into less space with multi-core processors and high-density server boards will help the CNM minimize ongoing power and cooling costs. The HPC cluster also has room to expand as the center grows.

Ultimately, the cluster will help to improve the quantity and quality of research that the center produces. In the past, computer resource limitations forced some nanoscience researchers to perform their computations in segments using smaller, less-powerful clusters. With high demand for those clusters, research could be a painstaking process.

Key Technologies

- Quad-Core Intel® Xeon® 5355 processors
- Supermicro X7DBT-INF* motherboard, developed by Intel, Supermicro, and SGI
- Intel® Fortran Compiler for Linux* OS
- Intel® Math Kernel Library

Using the new cluster, however, researchers can perform those computational tasks much faster and on a single system. For example, researchers no longer have to wait for data to be transferred to see an image; they can simply reposition a sample or adjust an instrument quickly and continue working. By providing a way to achieve fast results turnarounds in a single facility, the CNM can accommodate more researchers in less time—but most importantly, it can help researchers produce better science.



Find a business solution that is right for your company. Contact your Intel representative or visit the Intel® Business/Enterprise Web site at intel.com/business

Copyright © 2008 Intel Corporation. All rights reserved.

Intel, the Intel logo, Intel. Leap ahead, the Intel. Leap ahead logo, and Intel Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

This document and the information given are for the convenience of Intel's customer base and are provided

"AS IS" WITH NO WARRANTIES WHATSOEVER, EXPRESS OR IMPLIED, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT OF INTELLECTUAL PROPERTY RIGHTS. Receipt or possession of this document does not grant any license to any of the intellectual property described, displayed, or contained herein. Intel products are not intended for use in medical, life-saving, life-sustaining, critical control, or safety systems, or in nuclear facility applications.

Performance tests and ratings 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.

Intel may make changes to specifications, product descriptions and plans at any time, without notice.

0408/YMB/TDA/XX/PDF

318683-001US

