



“Because this is an inter-university research institute, the HP XC4000 cluster system is often used by external staff. ...We are very happy that the HP cluster is being used so frequently in such a short period of time after its introduction. We find it quite superior to the old system.”

Objective:

The Institute of Statistical Mathematics needed rapid deployment of a computational statistics support computer system and improved arithmetic processing speed. They wanted a system with stable operation 24 hours × 365 days a year.

Approach:

- ISM adopted a packaged HPC cluster solution to achieve quicker deployment.
- ISM implemented the HP XC4000 (HP ProLiant DL145 G2) cluster system with high reliability and availability.

IT improvements:

- The solution reduced system operating loads.
- ISM has realized excellent price/performance.
- The HP cluster solution improved processing speed with 128 nodes.

Business benefits:

- Provided researchers with user-friendly computing environment.
- Improved research efficiency.



The widening world of statistical science

Ever since its creation in 1944, the Institute of Statistical Mathematics (ISM) has responded to society's need for data analysis. As Japan's only specialist statistical science institute, ISM carries out research and development into methods and theories concerning fields such as prediction and knowledge discovery, uncertainty modeling and risk analysis, data design and examination, computational inference, fundamental mathematics, and statistical resources.

HP customer case study: Japan Institute of Statistical Mathematics

Industry: Education, Public Sector

“Building clusters normally takes time, but with the HP XC4000 cluster system, the operating system, middleware and overall optimization were all included as part of the package. It’s great to be able to put something in place and deploy it immediately.”

-Professor Nakano, director, Center for Engineering and Technical Support, The Institute of Statistical Mathematics



The “Akaike Information Criterion” (AIC) is a well-known project developed at the ISM. Often called the last great discovery of the twentieth century, AIC is utilized in a variety of areas employing statistical models. Akaike Hirotugu, Professor Emeritus of The Institute of Statistical Mathematics, proposed the AIC; he was awarded the 2006 Kyoto Prize by the Inamori Foundation for this achievement.

Recently, the ISM has expanded beyond its core competency as a research organization, to develop advanced projects that are closely linked to everyday life and involve researchers and businessmen from inside and outside of the institute. Current fields of study include research in statistical methods developed for a wide variety of areas including the atmosphere, oceans, tsunamis, life sciences, the environment, medicinal products, finance, and outer space.

For example, the Prediction and Knowledge Research Centre analyzes earthquake epicenters, while the Risk Analysis Research Centre performs quantitative measurement of financial risk and similar topics. This kind of cutting-edge research is supported by two supercomputer systems: the statistical sciences supercomputer and the computational statistics support system.

Professor Tamura Yoshiyasu, deputy director of ISM, describes the computing upgrade, “Statistical scientists need to process large quantities of data. Our institute has always adopted and utilized the latest computer systems, ever since we brought in the first commercial parametron computer in 1956. We installed the HP XC4000 cluster system as our computational statistics support system to give us the most up to date computing environment available.”

Migration to new computational statistics support system

Prior to the installation of the HP cluster, the previous statistics support system used shared memory, which caused various problems. It was not possible to compare the results for each task immediately, and CPU and disk I/O performance was inadequate. The researchers requested a computing solution with higher processing speed and superior cost performance. Consequently, ISM placed these requests high on its list of priorities during its selection process.

“The factors we considered when migrating to a new system included cost, processing capability, speed of deployment, and the availability of user-friendly search functions. Statistics deal with huge volumes of data, so it is important to have a high CPU speed and fast disk input/output (I/O), while the availability of user-friendly search functions is a major attraction for the researchers. From the viewpoint of usability, it is also important to be able to use the optimum compiler for each of the various user programs,” explains Professor Tamura.

Superior price/performance

ISM compared offers from five or six companies, before choosing the HP XC4000 cluster system as its new HPC system.

“There were several deciding factors in our selection of the HP solution. We liked the fact that the HP cluster supported Lustre-based HP StorageWorks Scalable File Share (HP SFS), which can relieve file I/O bottlenecks. We also liked that the package included maintenance and support. The fact that the XC clusters and SFS were both optimized for High Performance Computing (HPC) was also important. The HP XC4000 had a further advantage on the cost front due to its cluster configuration.”

Professor Tamura Yoshiyasu, deputy director of ISM

Professor Nakano Junji, director of ISM’s Centre for Engineering and Technical Support, continues, “The fact that the CPUs were 64-bit AMD Opteron-based also counted in favor of the HP XC4000 cluster system. We did some benchmark tests, and it was approximately twice as fast as the other companies’ systems. The HP XC4000 cluster system came out first in the overall evaluation and achieved almost full marks in the benchmarks.”

All-in-one solution allowed rapid deployment

The XC4000 cluster system is part of the HP Unified Cluster Portfolio, (UCP) an innovative, modular package of hardware, software, and services for scalable computation, data management, and visualization. With the HP Unified Cluster Portfolio, HP provides organizations with a wide range of industry-leading cluster components and the experience needed to ensure a smooth deployment.

Because the XC4000 cluster is an all-in-one, packaged solution as part of the HP UCP, ISM was able to rapidly deploy the optimized solution. Shortly after deployment, the HP XC4000 cluster system began full-scale operations. The solution is configured with 256 AMD Opteron CPUs, 640GB of memory and a 128 node configuration, achieving an overall performance of 1331.2 GFLOPS. It is equipped with queue classes with a maximum of 64 CPUs and 256GB of memory, and installed with compilers such as the Intel Compiler (C++/FORTRAN) and PGI Compiler (C++/FORTRAN). Using an HP SFS server for storage management also made it possible to avoid the file I/O bottlenecks that were a problem with the previous Network File System (NFS).

“Building clusters normally takes time, but with the HP XC4000 cluster system, the operating system, middleware and overall optimization were all included as part of the package. It’s great to be able to put something in place and deploy it immediately. Post-installation support has been good—the system engineers in charge did their job well.”

Professor Tamura Yoshiyasu, deputy director of ISM

Customer solution at a glance

Primary applications

ISM develops applications in many areas, including:

- Modeling of causally, temporally and/or spatially inter-related complex phenomena
- Model-based statistical inference methodologies.
- Modeling and inference for statistical analysis of time series, spatial and space-time data
- Methods for extraction, processing and transformation of information in intelligent systems
- Analyses of the data generated by a system with graph structure

Primary hardware

- 128 nodes (256 AMD® Opteron™ CPUs)
- 640GB physical memory
- Peak performance: 1331.2 GFLOP/S

Primary software

- Intel Compilers (C++/FORTRAN)
- PGI Compilers (C++/FORTRAN)
- HP StorageWorks Scalable File Share (HP SFS)

Performance of HP XC4000 cluster system praised

Since it was fully commissioned as a supercomputer, use of the HP XC4000 cluster system has risen continuously, and it has gained an excellent reputation among the researchers. "Because this is an inter-university research institute, the HP XC4000 cluster system is often used by external staff. It runs 24 hours a day with an average operating ratio of 20-30 percent—sometimes even over 50 percent. There are about 10-20 consistent users, who expect that their jobs will always be run in a timely manner," explained Professor Nakano. "We are very happy the HP cluster is being used so frequently in such a short period of time after its introduction. We find it quite superior to the old system."

Professor Nakano continued, "Overseas researchers are also permitted to use the cluster over the network. Institute staff still want to use the system, even when they're away on a business trip. We do monitor the usage, but in principle, we want people to use it as freely as possible," explains Professor Nakano.

"We are hoping that HP will offer us an even faster HPC system when we come to our next system renewal," concludes Professor Tamura.

To learn more, visit www.hp.com/go/hptc

© 2007 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

AMD and AMD Opteron are trademarks of Advanced Micro Devices, Inc.

4AA1-6158ENW, November 2007

