CASE STUDY

Intel® Data Center Manager (Intel® DCM)
Reducing the Power Consumption of HPC Environments

Monitoring the temperature and power consumption of each server with Intel® DCM and reducing power consumption by up to 5-8% through allocation to nodes with high power efficiency

Satisfying facility power regulations and social demand for energy conservation

Academic Center for Computing and Media Studies, Kyoto University

Location: Yoshida-honmachi, Sakyo-ku, Kyoto-shi, Kyoto

Established: April 1969

Activities:
Research and development related to advanced use of IT platforms and media, and development, operation, and management of the campus' IT environment

http://www.media.kyoto-u.ac.jp/

Hiroshi Nakashima (Center)
Professor, Ph. D.
Academic Center for Computing and Media Studies, Kyoto University

Keiichiro Fukazawa (Right)
Associate Professor, Ph. D.
Academic Center for Computing and Media Studies, Kyoto University

Junichi Hikita (Left)
Leader
Supercomputing Section
Planning and Information Management Department
IT Service Division, Kyoto University

Challenge

- Monitoring CPU and memory power consumption to improve HPC power efficiency
- Power capping to satisfy the facility and social demands

Solution

- Intel® Data Center Manager (Intel® DCM)
- Cluster-type HPC equipped with Intel® Xeon® and Intel® Xeon Phi™ processors

Providing advanced IT services to Kyoto University and research institutes across Japan

The Kyoto University Academic Center for Computing and Media Studies (ACCMS) promotes research and development using a five-department structure which includes the fields of networks, supercomputing, educational systems utilizing multimedia, and digital academic content, as well as collaborative research. It also functions as a base which incorporates IT platform network centers from eight universities nationwide for joint use and collaborative research, supporting the utilization of its supercomputer system (HPC) and collaborating with researchers in various computational science fields. The office of Professor Hiroshi Nakashima engages in various research related to HPC and works in cooperation with the University's Institute for Information Management and Communication IT Platforms Department to help improve services.

A vector-type HPC system was first introduced in 1985, with a transition to a scalar type occurring in 2004. Since then, the system has been refreshed roughly every four years, with overhauls in 2008, 2012, and 2016. The HPC environment as of 2018 consists of three systems: two cluster systems comprised of HPC servers featuring Intel® Xeon® processors, and one MPP system comprised of HPC servers featuring Intel® Xeon Phi™ processors. This system provides an overall computational performance of 6.5524 PFlops. "At present, 40% of the use of the ACCMS HPC system occurs within Kyoto University, with the remaining 60% occurring externally. Although the number of users is growing year upon year, our policy is to maintain an operating ratio of approximately 70% to allow us room to cope," explains Professor Nakashima.

Improving power efficiency rates essential for reducing user costs

HPC operation is significantly affected by the power consumption of servers. One reason for this is a facility-related issue, with the maximum amount of suppliable power (power cap) being determined based on the facility. "In particular, there was a movement to cap HPC power consumption across Japan in the wake of the 2011 Tohoku earthquake," recalls Professor Nakashima.

Furthermore, nuclear power plants throughout Japan stopped operating as a result of the earthquake, and power prices rose roughly 50% at peak times in the Kansai region. As usage charges for the HPC system also include electricity fees, there has been a demand for improved power efficiency rates to reduce the cost to the user.

However, the pace of improvement in power efficiency rates has been slow in relation to the improvements seen in processing performance rates resulting from the evolution of processors. Because of this, power management is an important issue and Professor Nakashima's office has been performing research related to code optimization technologies to maximize application performance, and software for adaptive control of power management mechanisms.

Aware of the necessity for more intricate power control amid increasing demands for power consumption capping and improved power efficiency rates, ACCMS introduced Intel® Data Center Manager (Intel® DCM) to monitor and control power and temperature in real time during the October 2016 HPC refresh. Planning and Information Management Department IT Service Division's Supercomputing Section Leader Junichi Hikita says, "Although we had been measuring power performance at the level of individual racks and nodes in previous environments, we decided to introduce Intel® DCM to monitor individual servers in even more detail."

Checking the temperature and power consumption of each server as well as CPU, memory, and I/O usage rates

Figure 1. Intel® DCM console screen

Intel® DCM is currently being used at ACCMS for monitoring HPC power consumption and temperature. Because there is no possibility of the total power consumption of the HPC system exceeding the rated power of the entire facility, the power cap functions are essentially safeguards and actual caps are not implemented. However, for regular monitoring of power consumption and temperature, the Intel® DCM console screen is used to check the temperature and power consumption of each server and CPU, memory, and I/O usage rates (Figure 1). "Although we don't check the console screen daily other than when an error has occurred, we output and save log information and provide it to Professor Nakashima's office as data for research as necessary," explains Hikita.

Reducing power consumption by allocating jobs from nodes with high power efficiency

Professor Nakashima's office conducts a variety of research to improve HPC power efficiency based on the log data output from Intel® DCM with the aim of feeding the research back into actual operation. Tests were carried out with three specific measurements: (1) changing power consumption times when running applications, (2) variations based on node power consumption rankings and different applications, and (3) power consumption based on allocating jobs from nodes with high power efficiency.

In the tests for (1), after ascertaining that there were variations in power consumption depending on CPUs when the same application was run, it was confirmed that power consumption changed as a result of CPU heat etc. over time. In the tests for (2), benchmark programs for memory load (STREAM*) and CPU load (HPCG) were run on each node, measuring the power consumption when the program was running. The results showed that even on the same node, there were differences in power consumption depending on the program. Keiichiro Fukazawa, an Associate Professor at ACCMS, had the following to say regarding the results.

"We were able to confirm that there were variations in power consumption caused by individual differences in CPUs with the same specifications, and that there were some with high power efficiency and some with low power efficiency. With the low-performance CPUs, the hotter they become, the more power they consume. Improving power efficiency requires correct monitoring, and we expected that allocating jobs from nodes with high power performance would reduce power consumption."

In practice, schedules were set based on monitoring values, and a comparison of cases where jobs were allocated from nodes with high power efficiency and cases where jobs were allocated randomly confirmed a 2-4% reduction in power consumption, even with a 70% node usage rate (Fig. 2). In addition, when compared to cases in which jobs were allocated from nodes with the worst power efficiency, there was a 5-8% reduction in power consumption at a 70% node usage rate. "A 2-4% decrease in power consumption may seem small, but it is a huge result for ACCMS, where the yearly electricity fee reaches about 150 million yen. This result is not only beneficial to an HPC environment, but to general data centers as well," says Professor Nakashima.

The University predicts that by placing CPU power consumption restrictions on programs with no CPU load as well as memory power consumption restrictions on programs with low memory usage rates, it will also be possible to tune the system in the future so that power consumption is reduced without affecting performance.

Fukazawa intends to continue performing research to reduce power consumption, saying that, "we also want to grasp application characteristics through machine learning to automate job allocation schedules." Furthermore, he is examining strengthened computational performance for the next HPC refresh and is expecting further processor evolution. Intel will continue to evolve Intel® Xeon® processor and management tool technology to contribute to the advanced IT platforms that ACCMS aims for.

Estimation of Power Saving

Allocating jobs from nodes with high power efficiency

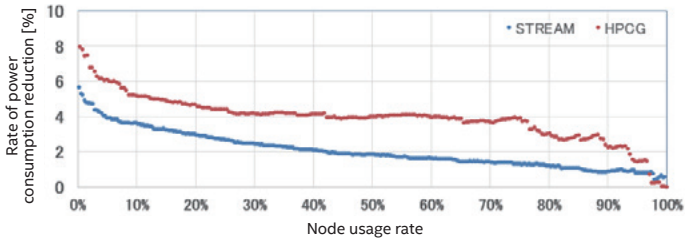

A 2-4% reduction in power consumption compared to normal job allocation (when node power efficiency is ignored), even with a 70% node usage rate, can be expected.
An enormous figure for supercomputer centers with high electricity charges

Figure 2. Rate of power consumption reduction when allocating jobs from nodes with high power efficiency (Blue: STREAM* (memory prioritized), red: HPCG (CPU prioritized))

For more information on Intel® Data Center Manager, please visit the following site.
http://www.intel.com/dcm

JPN/1807/PDF/SE/SSG/MT 337847-001US