



Case Study

Server

Tokyo Institute of Technology

Tokyo Institute of Technology's cluster-type supercomputer, Tsubame, was launched in 2006 as a 'supercomputer for everyone' for cutting-edge research. Upgraded over the years, the current Tsubame 2.5 is widely used by the Tokyo Institute of Technology, external research institutions, businesses and other organizations. From the beginning, SUSE® Linux Enterprise Server was chosen as the operating system for Tsubame's stable operation.



Overview

Tokyo Institute of Technology (Tokyo Tech) is a world-leading science and technology university that cultivates a wealth of human talent capable of making a mark on the global stage in science and technology. The university makes a positive contribution to society by generating research results that help solve global problems. Established in Tokyo in 1881, it employs around 1,300 academic staff and serves around 10,000 students.

Challenge

Tokyo Tech has undertaken research on cluster-type supercomputers since 1996, as Professor Satoshi Matsuoka of the Global Scientific Information and Computing Center at Tokyo Tech explains:

“SUSE Linux Enterprise Server for High Performance Computing enabled Tsubame 1.0 to gain many users and become the leading large-scale cluster-type supercomputer in Japan.”

PROFESSOR SATOSHI MATSUOKA

Global Scientific Information
and Computing Center
Tokyo Institute of Technology

“The opportunity to use supercomputers was traditionally restricted to specific categories of researchers, but there was growing recognition at that time of a need for ‘a supercomputer for everyone’ that a wider variety of users could benefit from. The challenge was to create a new computing environment while maintaining reliability and usability.

“Our objective with our supercomputer, Tsubame, which we planned to begin operating by 2006, was to build a cluster-type supercomputer with performance superior to that of Earth Simulator—the best performing supercomputer in Japan at that time.”

Tokyo Tech knew that having an operating system that was easy to maintain and manage, while still providing high reliability, would be crucial to Tsubame's development.

“There wasn't really much change in terms of the complexity of the applications and other factors; however, there was a need to maintain reliability at the same level as in the past. It was especially important to use an operating system with a high level of reliability because, if a problem with the operating system occurred, all of the compute nodes would stop functioning,” notes Professor Matsuoka.



Tokyo Institute of Technology at a Glance:

■ Industry and Location

Education and Research, Japan

■ Products and Services

SUSE Linux Enterprise Server for High Performance Computing

■ Results

- + Supports peak performance of 5.7 petaflops (PFlops), the best of any cluster-type supercomputer in Japan
- + Provides interconnectivity with consumer IT equipment as well as support for high-end environments
- + Helped Tsubame become one of the world's top-ranked large-scale cluster-type supercomputers
- + Offers excellent management capabilities

Solution

Tokyo Tech wanted an open-source operating system, as Professor Matsuoka explains, “Since running a supercomputer requires putting together a wide range of software stacks and having them apply patches to themselves as necessary, the operating system needed to be open source. We chose Linux, which at the time was already coming into widespread use.”

After careful consideration Tokyo Tech decided to adopt SUSE Linux Enterprise Server for High Performance Computing.

Tokyo Tech had already used SUSE Linux Enterprise Server on 32-core IBM POWER architecture servers. “About two years before Tsubame began operation, while we were still in the research stage, SUSE Linux Enterprise Server already had a record of successful implementation, with stability that was just as good as the IBM AIX UNIX operating system,” says Professor Matsuoka.

The first version of Tokyo Tech’s supercomputer, Tsubame 1.0, was a relatively large-scale supercomputer with 5,000 CPU sockets and over 10,000 cores. With the adoption of an architecture that included 32 GB of memory (16 cores on 8 sockets) in each compute node, creating a “fat node,” as well as enhanced parallelism within each node, multi-threading was required to utilize shared memory. SUSE Linux Enterprise Server made it possible to make effective use of the highly reliable hardware and system architecture that had been adopted. The end result—Tsubame 1.0—overtook Earth Simulator in the ‘Top 500’ supercomputer performance rankings in June 2006, achieving a global ranking of seventh place.

“SUSE Linux Enterprise for High Performance Computing enabled Tsubame 1.0 to

gain many users and become the leading large-scale cluster-type supercomputer in Japan,” comments Professor Matsuoka.

In 2010, Tsubame 1.0 was upgraded to Tsubame 2.0, which delivered a 30-fold improvement in performance. This marked the birth of an entirely new kind of supercomputer, one that used graphics processing units (GPUs) with their outstanding parallel processing capabilities. Not only was Tsubame 2.0 ranked fourth in the world in performance, it also ranked number one in the world among production supercomputers in the Green500 energy efficiency rankings.

The process of evolution continued with Tsubame 2.5, a cluster-type supercomputer running SUSE Linux Enterprise Server for High Performance Computing that features further improvement in GPU performance and the ability to run a wide range of software in line with user needs.* Tsubame 2.5 has 4,200 GPUs installed in its 1,442 compute nodes, giving it peak performance of 5.7 petaflops (PFlops)—the best performance of any cluster-type supercomputer in Japan. Tsubame 2.5 uses an industry-standard hybrid computing architecture that combines CPUs (for scalar operations) with GPUs (for vector operations). The network used to link the compute nodes is the fastest fiber optic network in the world, with a bisection bandwidth of 200 TB/s. The architecture guarantees that there is no disruption even if all compute nodes are communicating at the same time.

Results

SUSE Linux Enterprise Server plays a critical role in the overall reliability of Tokyo Tech’s supercomputer.

“Adopting a hybrid computing approach that combines different types of

processors to achieve dramatic improvements in hardware performance means that peripheral systems must become more complex,” says Professor Matsuoka.

Managing all of this complexity properly places an immense burden on software, including the operating system. “The excellent management that SUSE Linux Enterprise Server provides, including the management structure itself, is one of the key factors behind Tsubame 2.5’s success. And of course, regarding reliability, the fact that SUSE Linux Enterprise Server is a distribution of Linux means that a wide range of management and other tools are available. From that perspective, SUSE Linux Enterprise Server plays a very important role in the overall reliability of our supercomputer,” comments Professor Matsuoka.

At Tokyo Tech, SUSE Linux Enterprise Server provides interconnectivity with ordinary consumer IT equipment, as well as support for more high-end environments.

“There are many cases of Tsubame being used with PCs and workstations, with students performing testing in their own computing environments and then transferring the results to Tsubame. Some Tokyo Tech students run multiple virtual machines (VMs) on their laptops, with client-version openSUSE running on the VMs; after using this testing environment, they move the results to Tsubame. The way that SUSE provides interconnectivity with ordinary consumer equipment, as well as support for more high-end environments, is very important,” notes Professor Matsuoka.

* In the latest Top500 Supercomputers list, published in June 2016, Tsubame was ranked #31 in the world. www.top500.org/lists/2016/06/

The next iteration of the supercomputer, Tsubame 3.0, is scheduled to be launched in 2017, and envisages the integration of cloud computing with supercomputers. It will require an environment in which different kernel versions can be deployed and coexist within the same operating system.

“The operating system will play a crucial role in ensuring that a complex array of hardware and software elements work together effectively,” says Professor Matsuoka. “There will need to be a framework for providing a runtime environment

that can be tailored flexibly to meet user needs, and a lot will be expected of the operating system. The operating system is the core element, not just for the kernel, but also for the file systems, schedulers, virtualization layer, and so on. Our choice of SUSE Linux Enterprise Server will help us to achieve high performance and usability while connecting to a wide range of middleware.

Although the evolution to Tsubame 3.0 will inevitably lead to the software stacks

becoming more complex, the Global Scientific Information and Computer Center’s role as a leading-edge research institution makes it vital for the center to take on the challenge of developing a new kind of supercomputing environment. Professor Matsuoka says, “The center is committed to moving forward with the creation of next-generation supercomputers in close collaboration with leading HPC vendors such as SUSE and with other research institutions.”

“The excellent management that SUSE Linux Enterprise Server provides, including the management structure itself, is one of the key factors behind Tsubame 2.5’s success.”

PROFESSOR SATOSHI MATSUOKA

*Global Scientific Information and Computing Center
Tokyo Institute of Technology*

www.suse.com



**Contact your local SUSE Solutions
Provider, or call SUSE at:**

1 800 796 3700 U.S./Canada
1 801 861 4500 Worldwide

SUSE
Maxfeldstrasse 5
90409 Nuremberg
Germany