

CHALLENGE	SOLUTION	BENEFIT
Provide supercomputing facilities for all research projects at an academic institution with a limited budget	A high-performance cluster using standard components that include Dell™ PowerEdge™ servers, Intel® Pentium® III processors, and the Linux® operating system	Better allocation of supercomputing resources, increased real estate within the computer center, lower costs, less maintenance, and high performance previously available only on expensive, proprietary servers

A lesson in supercomputing

The School of Computing at the National University of Singapore gains supercomputing performance by building a cluster of Linux-based Dell servers running Intel processors

High-performance computing has always required significant resources and investment. Buying and running a traditional number-crunching supercomputer calls for huge sums of money and uncommon expertise in maintaining proprietary platforms.

The School of Computing at the National University of Singapore, however, had a compelling need—to make sure that its 100 professors and lecturers, 2,500 undergraduate students, and more than 500 graduate students had the necessary computing resources to perform their numerous research projects. To build such a conducive learning environment, the School of Computing had to invest in high-performance computing.

The National University of Singapore is a leading global university with 10 departments offering undergraduate and graduate programs. The university also boasts 12 affiliated national-level and 14 university-level research centers. Established in 1998, the School of Computing is an international center of excellence for multidisciplinary learning and research in information technology. This multidisciplinary teaching approach distinguishes the School of Computing from other departments within the university.

To support these research efforts, the School of Computing had to find more scalable and affordable high-performance computing (HPC) resources. The school wanted to break the vicious cycle of buying and upgrading expensive proprietary platforms, so it decided to deploy an open platform that ran open source software powered by industry-standard processors. The Linux® operating system was

“We wanted to provide the smallest footprint, rack-optimized server with the fastest Intel CPUs—these requirements were critical to our needs.”

— Professor Teo Yong Meng
School of Computing, National University of Singapore

the obvious choice for software; Intel® processor-based Dell™ PowerEdge™ servers solved the hardware part of the equation.

Linux opens up the world of supercomputing

Linux and open source software have made parallel processing and supercomputing on Intel processor-based platforms readily available to the masses. Instead of staking large investments on expensive systems, organizations can now run their number-crunching software on less expensive hardware and software. In fact, studies have shown that Linux clusters can meet the needs of 70 to 80 percent of supercomputing tasks, making them the preferred platform for high-performance technical computing.

The School of Computing at the National University of Singapore wanted to leverage the advancements in this market so that all its research programs could tap into the same level of

computing power and the HPC resources would not be limited to a privileged few.

To accomplish this goal, the School of Computing had to build a supercomputing architecture that would provide far more processing capacity at a cost 10 times lower than previous levels. “We had been using an older Fujitsu® SunSPARC™ cluster,” says Teo Yong Meng, a professor in the School of Computing. “It had performed well, but we needed to upgrade it to a faster and more open architecture to address the needs of our students.”

Space was also a constraint. “We wanted to provide the smallest footprint, rack-optimized server with the fastest Intel CPUs—these requirements were critical to our needs,” says Professor Meng. For this reason, the school chose Dell servers.

Dell standardizes HPC components

Working with Singapore Computer Systems (SCS), Dell built a Linux-based HPC cluster. This type of cluster relies on commercial off-the-shelf (COTS) components—Intel processor-based servers, commodity networking equipment, and open source software—to build a high-performance system for parallel processing tasks and computations. These clusters offer another advantage over proprietary supercomputers: The investment in setting up and maintaining Linux-based clusters is transferable to future systems with a similar architecture.

“Since the SCS Linux Competency Center (SCS-LCC) team has a strong track record of implementing clusters in Singapore since 1999, we were comfortable that their experience with Linux and cluster computing meant that our clusters would be implemented correctly and would be able to meet our technical requirements,” says Professor Meng.

The Linux-based HPC cluster installed in May 2002 uses 64 nodes of dual symmetric multiprocessing (SMP) Dell PowerEdge servers. The cluster’s engine comprises 128 Intel Pentium® III processors, each running at 1.4 GHz. These processors power the Red Hat® Linux version 7.2 operating system and Rocks Cluster Distribution software from the National Partnership for Advanced Computational Infrastructure (NPACI). All hardware was promptly delivered to the school and it took only three engineers to get the cluster up and running within three days.

A Gigabit Ethernet¹ network serves as the interconnect for cluster administration and for computing tasks not using the Message Passing Interface (MPI). For high-performance, MPI-based parallel computation, the cluster uses a Myricom™ Myrinet™-2000 network. The Myrinet-2000 cluster fabric is a high-performance packet-communications and switching technology that is widely used to interconnect HPC clusters.

Dell delivers performance/peace of mind with a smaller price tag

Once the cluster was up and running, the School of Computing immediately witnessed the benefits of choosing the Dell implementation.

“An acceptance test based on the parallel Linpack benchmark was performed to ensure that all the cluster hardware was performing properly,” says Professor Meng. The cluster achieved a Linpack result of 100 GFLOPS, which exceeded the university’s expectations.

The Dell-Linux cluster also offers significant scalability. “Scalability is vital as we deploy supercomputers based on clusters of Dell SMP servers. Today, you can buy 8 servers; in six months when demand increases, you can purchase another 8 or 16 servers and easily plug them in. This incremental and modular scalability of the Dell-Linux clusters makes them very attractive from a performance, manageability, and financial perspective,” says Laurence Liew, manager of the SCS Linux Competency Center.

Training, key to any implementation success, was not forgotten. The SCS engineers held a training seminar for the staff of the School of Computing to teach them how to use and maintain the cluster.

“The total costs of the cluster including all networking equipment, software, and services came to about \$500,000²,” says Professor Meng. “A 128-CPU supercomputer that costs \$500,000 would not have been imaginable a few years ago.”

In addition to providing cost-effective supercomputing, the Dell-Linux infrastructure minimized the space used and eliminated the need for expensive professionals to administer the systems.

“The Linux cluster at the School of Computing has three advantages that other colleges should emulate: NPACI Rocks software has made the Linux cluster easy to manage; Dell provides the essential, reliable server-class building blocks; and SCS-LCC has simplified the operation and maintenance of the project,” advises Dr. Philip Papadopoulos, program director for the San Diego Supercomputer Center, Grid and Cluster Computing.

Not all the benefits were tangible, however. The Dell-Linux cluster gave one important intangible advantage to the School of Computing: peace of mind. “We prefer to concentrate on our research and leave the maintenance of the cluster to the SCS-LCC engineers who work on such clusters daily,” comments Professor Meng.

That benefit alone—the ability to focus on the research that the HPC cluster enables instead of on the cluster itself—makes the Dell-Linux infrastructure truly valuable. By saving time, money, space, and effort, the Dell-Linux cluster allows everyone at the School of Computing to get out of the data center and back into the classroom. ☺



¹ Gigabit Ethernet indicates compliance with IEEE® 802.3ab and does not connote speeds of 1 Gbps.

² Singapore dollars (U.S. \$283,000)