The Pennsylvania State University

CHALLENGE	SOLUTION	BENEFIT
Reduce the cost of high-performance computing in a demanding university environment	A 128-node cluster of Dell™ PowerEdge™ 1550 servers configured with Intel® Pentium® III processors and running the Linux® operating system	Cost savings; new types of research possible because of supercomputer-caliber performance; computer access available to additional researchers; high availability

A new research assistant

Penn State builds a 128-node Dell cluster configured with Intel processors and running Linux to extend high-performance computing services

University researchers are traditionally heavy users of computing power. They can use thousands of MIPS (millions of instructions per second) in modeling everything from the weather to economic theories. At most schools, however, there is never enough computing power to go around. The supercomputers traditionally used for numerically intensive research can cost millions of dollars—and require specialized, expensive care and maintenance.

The Pennsylvania State University (Penn State) is pioneering a better way. They are harnessing the power of hundreds of off-the-shelf Intel® processors inside Dell™ PowerEdge™ servers to provide the power of supercomputers at a fraction of the cost.

Penn State built its first evaluation cluster using Dell Precision™ workstations during the summer of 1998. With scalability confirmed, the school then built a full-fledged cluster with 33 Dell PowerEdge 4350[1] servers. The 64-processor LION-X cluster—named in honor of the school mascot and the cluster's operating system—provided compute power to more than 30 researchers exploring everything from materials design to weather modeling.

> "We are thrilled to be able to offer our academic community greater access to parallel computing resources using off-the-shelf technology."
>
> – **Kevin Morooney**
> Senior Director, Center for Academic Computing
> The Pennsylvania State University

Twice the power in one-fourth the space

In 2001, Penn State moved up the performance curve and down the price curve by building its second Dell-Linux cluster. This cluster included Dell PowerEdge 1550[1] servers, which pack CPUs that are twice as powerful as those in the first cluster and consume one-fourth the space of the older servers.

"Our two LION-X clusters were key to our ability to support parallel computing and enhance the computational researcher's

[1] Newer models are available at http://www.dell.com.

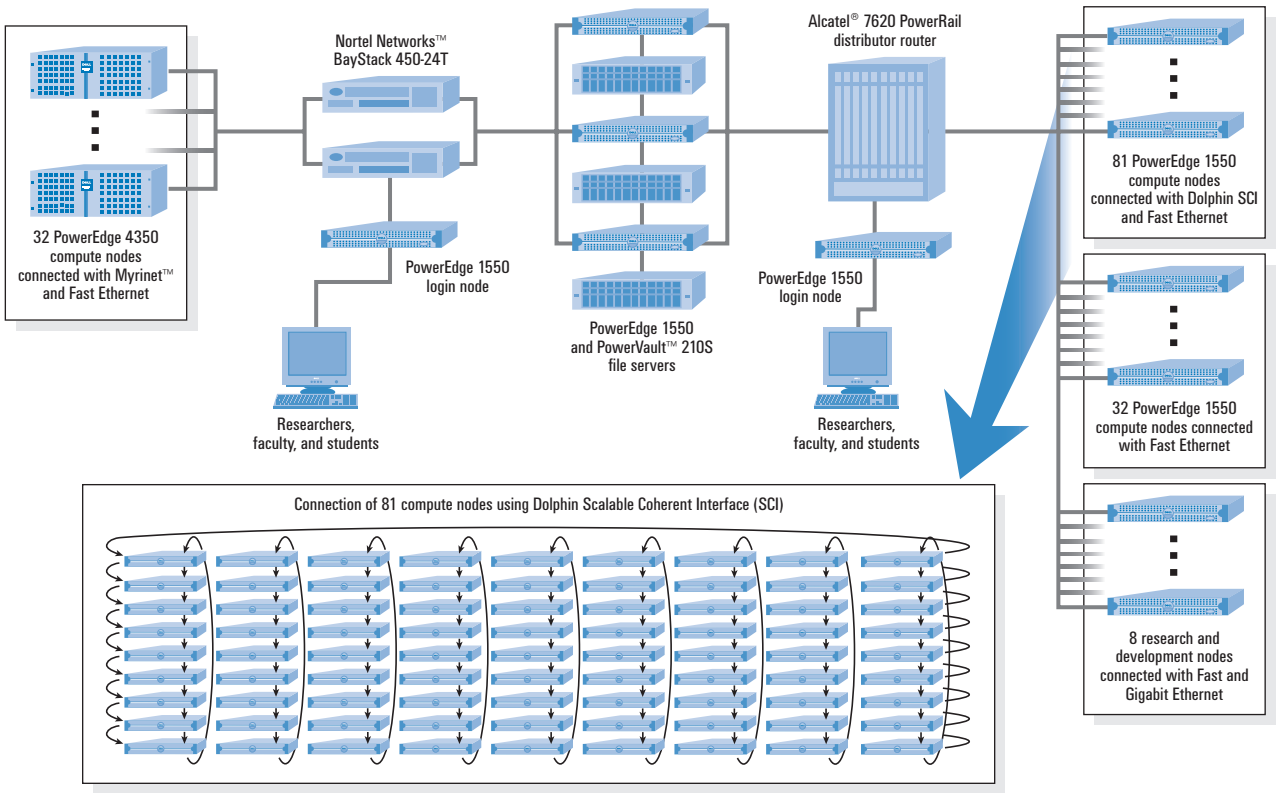

Rack-optimized Dell PowerEdge servers maximize computing power in a small footprint

environment," says Jeff Nucciarone, senior research programmer. "Leveraging the reliability and performance of Dell PowerEdge servers with the cost-effective nature of standards-based computing, we're providing researchers access to a powerful set of hardware to help them solve complex computational problems."

Clusters are rapidly gaining momentum as a new model for scalable, enterprise computing and an affordable, cost-competitive approach to academic and scientific computing. A networked cluster of Dell servers, combined with scheduling software to divide and distribute the processing tasks across the network, allows standards-based, commodity equipment to offer computation speeds comparable to specialized, more expensive, high-performance computers.

Eight times the power per square inch

Penn State considered several factors—initial price, ease of maintenance, reliability, footprint, and form factor—in the evaluation process before choosing Dell as the supplier. "We already had a strong relationship with Dell, so we were confident we would be well supported," says Vijay Agarwala, director of Penn State's High Performance Computing and Visualization Group. "Dell pricing was very competitive, and their range of servers met our specifications. Their systems have sufficient redundancy to meet our needs for uptime and recovery. Finally, we needed a rack-mountable system to match our physical space requirements, and the PowerEdge was well suited to that type of setup."

The new LION-XE cluster consists of 128 PowerEdge 1550 servers configured with dual Intel Pentium® III processors running at 1 GHz and equipped with 2 GB RAM and a 36 GB SCSI hard drive. Other features include a fast 64-bit, 66 MHz Peripheral Component Interconnect (PCI) bus that connects to two gigabit[2] network fabrics, and dual Ethernet adapters built into the motherboard.

Its ultra-thin 1U form factor is another advantage of the PowerEdge 1550. "Because the PowerEdge 1550 is so slim and the processors are more powerful than

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

those in our first cluster, we're getting eight times the power per square inch," Agarwala says. "You just don't see those kinds of rapid price/performance advances in the proprietary world."

The freedom of open source Linux

The Linux® operating system brings similar freedom, low costs, and rapid innovation on the software side. "Because Linux is open source software, we can tap into the Linux community," says Jason Holmes, research programmer. "We're not constrained by a proprietary system. Openness gives us the ability to collaborate with others by both contributing and taking advantage of solutions developed by others. We're not waiting for vendors to develop applications that we need."

> "Because the PowerEdge 1550 is so slim and the processors are more powerful than those in our first cluster, we're getting eight times the power per square inch. You just don't see those kinds of rapid price/performance advances in the proprietary world."
>
> **— Vijay Agarwala**
> Director, High Performance Computing and Visualization Group
> The Pennsylvania State University

Since they came online, the LION-X and LION-XE clusters have performed beyond expectation. Benchmarks run by Penn State's research computing staff have consistently shown excellent performance. Those benchmarks are available at the LION-XE Web site (http://gears.aset.psu.edu/hpc/systems/lionxe/performance/).

Strong demand

Demand from academic departments across the campus has been so strong that Holmes implemented a scheduling system to accommodate everyone's projects. The LION-XE cluster was actually purchased by 13 different research groups at Penn State and the Center for Academic Computing. Other researchers without funding also use it.

"These groups realized they could pool their money to buy one big, shared computer instead of each group buying a small machine," Holmes says. "Now they all have access to a far more powerful system. Because clusters can experience idle time between jobs, we have implemented a scheduling system to allow shorter jobs to run as well as to enable other groups to use the cluster when it's not in use by the priority groups."

All about data

All in all, about 80 researchers keep the LION-XE cluster churning around the clock. The LION-XE cluster supports a wide range of research.

Professor of Entomology and Biology Ottar N. Bjornstad uses the LION-XE cluster to study the rise and fall of epidemics of contagious diseases, with emphasis on learning the spatial spread of infection and prediction of local epidemic sizes. His research uses the cluster extensively to estimate spatial networks and perform random simulations of spatiotemporal models. The work has important practical consequences for design of intervention strategies and control of human diseases and bioterror agents.

The Garrison group, on the other hand, uses the LION-XE cluster to model fast energy deposition processes by performing molecular dynamics (MD) simulations. One such process is laser ablation. The group has developed a novel breathing sphere model that allows them to expand the time and length scales in the simulations. Applications of laser ablation include mass spectrometry of large biological molecules, thin film deposition, and medical procedures such as the eye surgery LASIK.

"We are thrilled to be able to offer our academic community greater access to parallel computing resources using off-the-shelf technology," says Kevin Morooney, senior director for the Center for Academic Computing. "With the LION-XE cluster, we have taken a balanced approach, providing an environment not only with very fast parallel computing cycles but also with high reliability."

Penn State is evaluating a Dell server based on Intel's new 64-bit architecture. "It's definitely in our future," Agarwala says. "A 64-bit architecture is ideal for many large-scale calculations. The rate at which these Dell servers are advancing is amazing."

FOR MORE INFORMATION

http://www.dell.com

http://www.intel.com