M U S T   R E A D

The Secret Lives of Supercomputers, Part 2

By Walaika Haskins
TechNewsWorld
08/08/08 8:30 AM PT

↗ Back to Online Version
✉ E-Mail Article
▤ Reprints

Supercomputers that can help with everything from molecular research to auto design are fast, powerful and incredibly expensive. Organizations in need of the sort of data only a supercomputer can provide sometimes face steep costs. However, less costly alternatives that approach the power of supercomputers are popping up.

Part 1
of this two-part series explores how industries employ supercomputers and the cost of using these extremely powerful systems. Part 2 examines alternatives for companies that can't afford custom-designed supercomputers.

Despite the declining cost of supercomputing, the technology still remains out of reach for many businesses and universities, some of which have found alternative solutions.

Bringing supercomputing to industry is just what the Blue Collar Computing program at the Ohio Supercomputer Center (OSC) was designed to do. Launched in 2004 with the support of the Ohio Board of Regents, the collaborative program seeks to provide easy and affordable access to advanced computing 🖳 technology. It offers resources, hardware, training, software and expertise to business clients so they can be more competitive, Jamie Abel, a center spokesperson, told TechNewsWorld.

Under the program, advanced computer technologies provide companies with innovative tools that allow virtual development of products such as cars, pharmaceuticals and financial instruments. Virtual modeling and simulation provide companies with a competitive edge through improved manufacturing processes that can reduce the time, labor and cost needed to bring products to market.

Simulation simplifies the choice of alternative processing methods. It offers better analysis and documentation of capabilities that boost efficiency, while improved factory and workflow layouts increase productivity 🖳, Abel explained.

For example, OSC recently developed an online weld simulation tool with Edison Welding Institute, Abel said. It uses OSC supercomputers to let welding engineers evaluate the changes in temperature profiles, material microstructures, residual stresses and welding distortion to reduce the extent of experimental trials during the design of welded joints.

Another case in point is Honda's design of the Accord through computer simulations without constructing a single physical prototype, Abel added.

Let's Set Priorities

Most supercomputing centers prioritize usage and have some sort of allocation process for time on the computers, known as resource units, he said.

"At OSC, our Statewide Users Group meets bimonthly to make decisions regarding allocations in general and account applications in particular. Standard and major accounts must undergo a review process," he said.

"Principal investigators provide the names of recommended reviewers, and OSC administrators identify other potential reviewers. These reviewers look at account applications and submit their recommendations to an allocations committee. At national labs, their funding organizations maintain an even more complex and competitive national allocation review and approval process."

Power From the People

What makes computers super, anyway? The definition is "somewhat fluid," Abel said.

"The supercomputer 🔍 of today is in many cases the personal laptop of tomorrow. Therefore, there may always be a limited number of 'supercomputers.' Certainly advances made in the topmost tier of supercomputing filter down quickly and make corporate and personal computing more and more powerful, allowing businesses and individuals to accomplish tasks that were out of their reach only a few short years or even months earlier," he explained.

Clustered supercomputing is an option typically employed by researchers. Rather than purchasing custom-built supercomputers, academics turn to the public and create a supercomputer using a network 📁 of regular PCs that users volunteer for use. This technique pools the powerful resources within today's systems in order to execute high-performance tasks.

The Berkeley Open Infrastructure for Network Computing was originally developed to support SETI@home, part of the Search for Extra-Terrestrial Intelligence project. Now it's used as a platform for distributed applications in mathematics, medicine, molecular biology, climatology and astrophysics. In March, it recorded a processing power of more than 960 teraflops (one trillion floating point operations per second) using a combined 550,000 active, networked computers.

The original SETI@home project logged a processing power of more than 450 teras across a network of 350,000 computers.

Folding@home, one of the best publicized cluster projects, made an appeal to owners of the PlayStation 3 🔍 video game console shortly after the device's release. The program called for gamers willing to add their PlayStation 3s to a distributed network in order to leverage the machine's Cell processor, using the Web-connected gaming machines to do computational work when not being used for play. Folding@home estimated that it would be able to achieve performance on a scale of 20 gigaflops per system, but it passed expectations. In September 2007, researchers reported processing power of 1.3 petaflops, of which nearly one peta was due to the use of PlayStation 3 systems. By March, more than one

million PlayStation 3 users had signed on to the project, Sony (NYSE: SNE) 🔍 reported.

Google (Nasdaq: GOOG) 🔍 is yet another example of the power of many. With an estimated 450,000 servers housed in the walls of its Googleplex, the search engine giant has a network that musters between 126 and 316 teras of computing power.

Gaming the Supercomputer

Clustered computers need not be massive undertakings, however. Frank Mueller, an associate professor of computer science at North Carolina State University, created a cluster using eight PlayStation 3 consoles. It was academia's first PlayStation supercluster.

Mueller configured the gaming machines to give him 64 logical processors -- more than enough power for high-level number processing.

Meanwhile, the eight PlayStations can still run the latest games. Their combined power is equivalent to that of a small supercomputer, at a fraction of the cost: US$5,000, according to Mueller, who announced his achievement last year.

"Scientific computing is just number-crunching, which the PS3s are very good at given the Cell processor and deploying them in a cluster," Mueller said.

"Right now one limitation is the 256 megabyte [random access] memory constraint, but it might be possible to retrofit more RAM. We just haven't cracked the case and explored that option yet. Another problem lies in limited speed for double-precision calculations required by scientific applications, but announcements for the next-generation Cell processor address this issue," he said.

The fastest supercomputers contain hundreds of thousands of processors. That's the case with IBM's (NYSE: IBM) 🔍 BlueGene/L, which was ranked No.1 on the Top 500 list from November 2004 to June 2008, when Roadrunner took the crown. BlueGene/L has 130,000 processors. Mueller estimates that a cluster of some 10,000 PlayStation 3s would surpass the computing power of BlueGene/L, with one caveat -- the cluster would have limited single-precision capabilities and networking constraints.

The Future of Supercomputing

As chipmakers continue to push the boundaries of processing power and develop new computing architectures, they extend the life of Moore's Law, which postulates that the number of transistors on a processor will double every 18 to 24 months.

Intel's (Nasdaq: INTC) 🔍 new multi-core Larrabee chip line, which will offer tens and eventually hundreds of cores on one processing unit, is an example of the innovation coming out of the microprocessor industry. Larrabee, dubbed a "graphics-capable processing unit" by Jon Peddie, president of Jon Peddie Research, is not a traditional graphics processing unit (GPU), but rather a "gang of x86 cores that can do processing," he wrote on the Peddie Research blog.

"This is a focus of new programming models and new computer architectures," Jim McGregor, an analyst at In-Stat, told TechNewsWorld. "It is great to see this type of innovation in the market now, because where we're going in the next few years will be phenomenal.

"For the first time, we really have to be innovative. Before, it was, 'We can use all the transistors we can just to build a better mousetrap, a better core.' Then it was like, 'We've kind of exceeded it, but the more memory we put on it gives it more performance.'

"Now we can do just about anything in silicate, putting all kinds of heterogeneous solutions on a single chip with different types of memory and

everything else. It's a great time in the industry for innovation," McGregor said.

As these advances trickle down, making their way from high-end computers to everyday systems, they leave the realm of supercomputers in flux.

"The definition of 'supercomputer' always changes," said the University of Tennessee's Dongarra. "Machines that were supercomputers 10 years ago are the kinds of things that are today's laptops." **ECT**

The Secret Lives of Supercomputers, Part 1

Social Networking Toolbox: ShareThis
Next Article in Must Read:
The Secret Lives of Supercomputers, Part 1

 Related Resources
  Telelogic Focal Point: Product Portfolio Management Solution
  Business Agility in a World of Brick and Mortar
  Accounts Payable has a New Hero
  Accelerating the Delivery of On-Demand Applications to Meet the Escalating Expectations of …