Empowering research at the National Center for Supercomputing Applications

- Technological innovation
- HPC with CPUs and GPUs
- Green efficiency

"We've had a very deep relationship with Dell for a number of years, and deployed a lot of their first-of systems."

John Towns, director, Persistent Infrastructure, National Center for Supercomputing Applications

Customer profile

Company:	National Center for Supercomputing Applications
Industry:	Research
Country:	Champaign-Urbana, IL (USA)
Employees:	200+
Website:	www.ncsa.illinois.edu

Business need

Design and deploy high-performance computing (HPC) clusters to meet the extreme computational and visualization needs of researchers

Solution

Dell and NCSA personnel worked in close collaboration to develop Abe and Lincoln, two HPC clusters with a combined speed of 137 teraflops.

Benefits

- **Get IT Faster: Deliver terascale systems to increase researcher productivity.**
- **Run IT Better: Use graphic processing units (GPUs) for better price/performance.**
- **Grow IT Smarter: Work with Dell ® to design systems that overcome current barriers.**
- **Partnering with Dell system design experts to collaboratively explore new platforms.**

The National Center for Supercomputing Applications (NCSA), located at the University of Illinois at Urbana-Champaign, provides powerful computers and expert support that help thousands of scientists and engineers across the country improve our world. With the computing power available at NCSA, both academic and industry researchers simulate how galaxies collide and merge, how proteins fold, how molecules move through the wall of a cell, how tornadoes and hurricanes form, and other complex natural and engineered phenomena.

NCSA—established in 1986 as one of the original sites of the National Science Foundation's Supercomputer Centers Program—is supported by the state of Illinois, the University of Illinois, the National Science Foundation, and grants from other federal agencies.

For NCSA, life is a constant quest to push the limits of supercomputers and the scientific and engineering applications that run on them.

NCSA is currently home to two Dell-based supercomputers, called Abe and Lincoln. Abe, launched in 2007, is a 90-teraflop cluster based on 1,200 Dell PowerEdge 1955 dual socket, quad core compute blades, an InfiniBand interconnect, and 400 TB of storage in a Lustre file system. Lincoln, installed in 2008, consists of 192 compute nodes (Dell PowerEdge 1950 dual-socket nodes with quad-core Intel Harpertown 2.33GHz processors and 16GB of memory) and 96 NVIDIA Tesla S1070 GPU computing systems.

The Challenges of Life at the Bleeding Edge

For John Towns, director, Persistent Infrastructure, National Center for Supercomputing Applications, life is a constant quest to push the limits of supercomputers and the scientific and engineering applications that run on them.

"Typically we're on the bleeding edge, so we're breaking things on a regular basis. That presents a lot of challenges," notes Towns, who is also the elected chair of the TeraGrid Forum. TeraGrid is a project that connects computing infrastructure at 11 partner sites funded by the National Science Foundation.

NCSA needs to build systems that are able to scale up dramatically to meet the needs of researchers who need extreme amounts of computational and visualization resources. NCSA needs to keep those systems operating at optimum performance levels, which requires system management software that scales to supercomputing sizes. And it needs to support researchers in their quests to scale their software applications—many of which are homegrown—up to a point that they can utilize massive compute clusters.

In addition, when it comes to supercomputing, one size or type of cluster doesn't fit all application needs, so the NCSA must maintain multiple computing architectures.

"We have a variety of different architectures on site, and they are targeted at different types of applications," Towns notes. These include a large shared-memory system that supports various applications from independent software vendors and high performance computing (HPC) clusters based on central processing units (CPUs) and graphics processing units (GPUs) that are deployed for general purpose computing.

CPUs are highly versatile processors with large, complex cores capable of executing all routines in an application. They are used in the majority of servers and desktop systems. GPUs are more narrowly focused processors with smaller, simpler cores and limited support for I/O devices. Recent generations of GPUs have specialized in the execution of the parallel computing-intensive portions of applications.

Exploring Scientific Frontiers with Abe and Lincoln

To meet the needs of researchers working at bleeding edge of scientific investigations, NCSA needs computing clusters that are at the bleeding edge of HPC technology. That's where Dell comes into the NCSA story.

"We've had a very deep relationship with Dell for a number of years, and deployed a lot of their first-of systems," Towns notes. "And for the last several years we've had some of the largest Dell systems out in the field." Even after three years of operation, Abe is still the 3rd largest Dell system in the field, moving to that position only due to two system deployments in the last year.

NCSA's Abe and Lincoln supercomputers are definitely in the class of first-of systems. They power a diverse range of scientific and engineering applications that are designed to take advantage of the capabilities of supercomputers.

"A large number of our researchers are able to make use of clusters, and thus we

Technology at work

Abe
1,200 Dell™ PowerEdge™ 1955 dual socket, quad core compute blades
InfiniBand interconnect
400 TB of storage in a Lustre file system

Lincoln
192 Dell PowerEdge 1950 dual-socket nodes with quad-core Intel Harpertown 2.33GHz processors and 16GB of memory
96 NVIDIA ® Tesla™ S1070 GPU computing systems

have Abe, which is our current largest resource—a 90 teraflop Dell cluster," Towns explains. "The largest number of users, and the most computing time we deliver, is through that resource."

To get an idea of the compute capacity of Abe, it helps to reflect on the meaning of 90 teraflops. A teraflop is one trillion floating point operations per second. So, in simple terms, Abe can perform 90 trillion calculations in a single second.

"And then the more interesting and on-the-edge sort of resource is Lincoln, which is our GPU cluster based on Dell nodes," Towns says. "There's a smaller community of users and applications that can effectively make use of that. That's an example of an area where we are trying to reach a little bit farther forward with technology in order to enable science and engineering research."

Lincoln, which is a 47-teraflop peak machine, is based on NVIDIA GPU units and conventional Intel CPUs. In its mixing of GPUs and CPUs, Lincoln is breaking new ground in the use of heterogeneous processors for scientific calculations. This combination allows NCSA to take advantage of the cost economies and extreme performance potential of GPUs.

"Currently the early adopters for GPU-based computing have been in the fields of biomolecular dynamics and high-energy physics," Towns notes. These researchers put Lincoln to the test on complex simulations that explore, for example, functions such as protein docking, which is involved in drug design, and the inner workings of biomolecular systems, such as viruses. The simulations that run on Lincoln help researchers understand the physical characteristics and behavior of their research targets at an extremely detailed level.

"A few years ago, a researcher on this campus, Dr. Klaus Schulten, did the first full simulation at a molecular level of a life form," Towns says. "The life form is very simple—it's the satellite tobacco mosaic virus— but it is actually a complete life form."

The program calculated how each of the million or so atoms in the virus and a surrounding drop of salt water interacted with almost every other atom every fermtosecond or millionth of a billionth of a second. With researchers turning to supercomputers to simulate from the world's simplest to the most complex life organisms, there's easily a need for exascale computing. Dr. Schulten directs the Theoretical and Computational Biophysics Group.

A team from Temple University's Institute for Computational Molecular Science, meanwhile, is using Lincoln to work to model surfactants, which are used in detergents, shampoo and other household products. They are also investigating another class of surfactants for its potential to control the delivery of drugs in the body and improve their impact.

This is no easy task. Mixtures of surfactants, water, and other molecules are exceptionally challenging to model computationally, according to NCSA literature. They self-assemble themselves into structures that can trap materials in a process that takes place at the micrometer scale over hundreds of nanoseconds, and this process often includes hundreds of millions of atoms.

Changing the Game with GPU Computing

Looking ahead, Towns sees a promising future for GPU computing. NCSA is investigating ways that GPU technology could be used to create an exascale computing system—or a supercomputer that could perform one quintillion calculations per second. That's the equivalent of one million teraflops. In round numbers, an exascale computer would be about 10,000 times faster than NCSA's Abe supercomputer.

"It seems very clear that traditional microprocessor-based systems are not going to be able to scale to that level," Towns notes. "So we need to look at new technologies to do that. GPUs are one possible path—we think a very promising path—toward exascale computing. It's also much more cost-effective for applications that have been either ported or developed directly for GPUs."

That was a conclusion that NCSA came to in the course of exploring its options for upgrading Lincoln. "What we have found is that for applications whose performance on a per-GPU basis is the equivalent or better than approximately 10 CPU cores, it's more cost-effective to have a GPU system," Towns says. "That's the breakeven point on cost."

The price advantage gets even better when applications have been developed to take advantage of the general purpose computing on GPU approach, he notes.

"What we're seeing, for the applications

> "GPU platforms can be anywhere from 5 - 50 or more times more cost-effective than a CPU-only based computing platform."
>
> *- John Towns, Director, Persistent Infrastructure, National Center for Supercomputing Applications*

that have emerged on GPUs, are applications that on a per-GPU basis have an equivalent performance of anywhere from 30 to 40 CPU cores all the way up to over 200 CPU cores. So this makes GPU platforms anywhere from 5 to 50 or more times more cost-effective than a CPU-only based computing platform."

In addition, GPU-based systems have some distinct total cost of ownership advantages over CPU-based system, Towns notes. These savings stem from reduced power and cooling consumption.

"The compute power density is a lot higher with the GPUs," he says. "They also have much greater heat density. The advantage is a smaller footprint and an attained performance per watt that is much greater than that of traditional CPUs. While there are some challenges in being able to cool and provide power, GPUs are more cost effective because the total power per FLOP and total cooling per FLOP are less."

Towns offers the example of a molecular dynamics application named AMBER, which many academic researcher use for biomolecular simulations. "For that application, researchers are realizing in the neighborhood of 5 to 6 gigaflops per watt. The thing to keep in mind is that most of the time when you're talking about this with respect for CPUs, you're talking about megaflops per watt. And that's realized performance, not peak performance. So the realized performance for this application on CPU cores is more on the order of 300 to 400 megaflops per watt, as opposed to 10 to 20 times that on a GPU. So it makes a big difference when it comes to considering total cost of ownership in delivering resources to a broad research community."

A Long Relationship with Dell

To meet the needs of researchers who continually push the limits of current technologies, NCSA works closely with its technology vendors to design and configure systems that break through current barriers in high-performance computing.

"This is where deep relationships with vendors like Dell have been very beneficial," Towns says. "We have been able to have some level of influence on the products that are being provided based on the requirements that we have from our user communities. Clearly, as an academic environment, it is a small part of the market for a vendor like Dell. However, they do have some specialized divisions like DCS who provide specialized solutions."

DCS, Dell's Data Center Solutions team, offers solution consulting, engineering, supply chain, and project management expertise for high-performance computing environments, hyper-scale data centers, and cloud computing environments.

"We've been working with DCS quite a bit lately to define platforms that will be very interesting for us going forward," Towns notes. "These platforms typically then trickle down to the more commodity products that vendors provide in future generations of their products. Often the academic community breaks new ground in systems design and configurations, and then those designs and configurations are migrated into the commercial market for servers."

This work to create new high-performance systems is far from straightforward. As much as anything, it's a search for answers, according to Towns. Working in a collaborative manner, personnel from NCSA and Dell have developed a series of supercomputers, each of which breaks new ground in high-performance computing. Its earlier-generation Dell clusters included Tungsten for academic research and T2 for industrial research, which were followed first by Abe and then by Lincoln.

"As we were looking at the work we had been doing over several years investigating GPUs, we recognized early on that it was a promising technology," Towns notes. "Building Lincoln was something that neither we nor Dell had done before. So it took quite a bit of working together to come to a configuration that would satisfy the needs as best we could. Dell was willing to do that."

Driving Lincoln Forward

In the ongoing quest to take their computing resources to new levels, Towns and his NCSA colleagues are looking closely at newer technologies and how they can be applied for extreme scale computing.

"We are investigating how we can secure funding for an upgrade to Lincoln," Towns says. "Lincoln is currently about a 47-teraflop peak machine. I'd like to put on site a system that is in excess of 300 teraflops based on GPUs. I'm planning to do an upgrade to that system to put a Fermi GPU-based system on site."

"Fermi", the codename for NVIDIA's next-generation CUDA GPU architecture, is "the most advanced GPU computing architecture ever built," according to NVIDIA. With more than three billion transistors and many hundreds of parallel processing cores, "Fermi delivers supercomputing features and performance at 1/10th the cost and 1/20th the power of traditional CPU-only servers," the company says.

In the push to upgrade Lincoln to a new-generation GPU architecture, Towns and his colleagues are not just trying to build a better supercomputer. They are working to create an engine that will propel science and innovation forward and help the United States stay competitive in a global economy.

"Centers like NCSA live in a highly competitive environment to secure funding to deliver capabilities to the science and engineering community. We must constantly be looking to push the limits to continue to provide leadership in the community and remain competitive," he says. "So you can't go back, and if you don't go forward, you're in trouble. We have to move forward. Lead, follow, or get out of the way."

For more information, visit NCSA's Web site at: http://www.ncsa.illinois.edu/

View Dell case studies at: dell.com/casestudies