Large Shared Memory System Used for “Really Cool” Science

Ease of Processing Massive Data Enables World-Changing Decision Analysis

Key Facts

Organization:
National Institute for Computational Sciences

Primary Location:
Oak Ridge, TN

Application:
Higher Education & Research

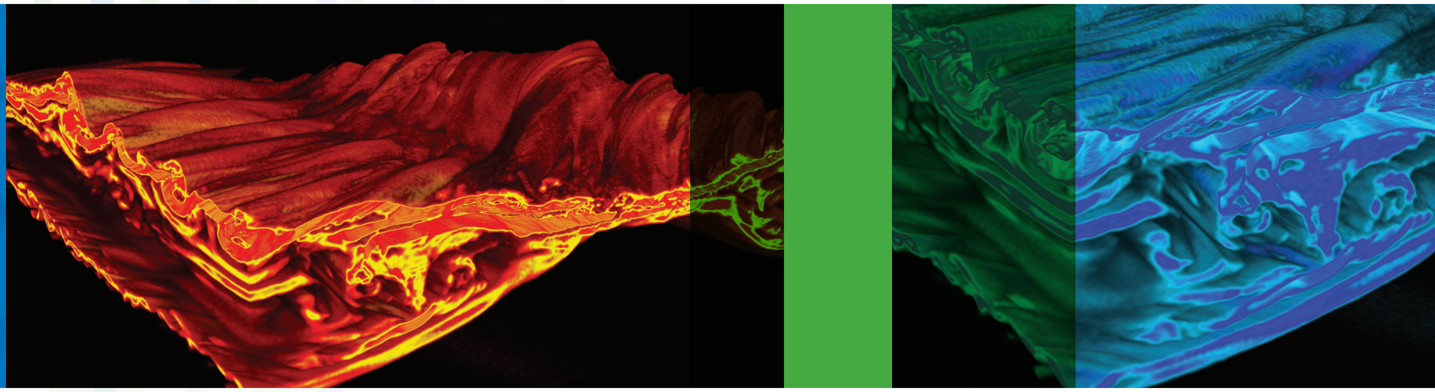

The National Institute for Computational Sciences (NICS) is a partnership between the University of Tennessee (UT) and Oak Ridge National Laboratory that is funded by the National Science Foundation (NSF) to support research at many U.S. universities. NICS resources provide a conduit for scientific discovery, including enabling very large scale data analysis as well as visualization of simulation data. Most NICS resources are provided as part of the Extreme Science and Engineering Discovery Environment (XSEDE) allocation process, and approximately 400 research studies are being undertaken at any given time using these NICS resources.

In 2010 NICS installed an SGI UV shared memory high performance computing (HPC) system. As a shared resource, the system is available for researchers from a variety of institutions who may want to take advantage of its power. A wide range of technical research in the physical and social sciences is being conducted on the system, and members of NICS leadership are witnessing firsthand the amazing insights coming out of the system’s use.

Nautilus: The Heart of the Remote Data Analysis and Visualization Center

Nautilus, an SGI UV 1000 system, is the centerpiece of the NICS Remote Data Analysis and Visualization Center (RDAV). Nautilus consists of 1,024 cores of Intel® Xeon® X7550 processors, 4TB of global shared memory, and eight NVIDIA® Tesla™ GPUs in a single system image, as well as a 960TB Lustre® file system. Nautilus currently has a CPU speed of 2.0GHz and a peak performance of 8.2TF.

The primary purpose of Nautilus is to enable data analysis and visualization of data from simulations, sensors or experiments. Nautilus is intended for serial and parallel visualization and analysis applications that take advantage of its large memory, multiple computing cores and multiple graphics processors. Nautilus allows for simultaneous utilization of a large number of processors for distributed processing, as well as the execution of legacy serial analysis algorithms for very large data processing by large numbers of users.

System Flexibility Facilitates Handling a Wide Variety of Research Projects

The initial idea for the installation of the system was for data analysis, though some simulation work is also being done. NICS leadership wants the system to be used “to turn data into knowledge and insight — to get understanding out of a large bank of data, whether that be from sensors, simulations, news articles or what have you,” according to Sean Ahern, Director of RDAV. “We have several scientists doing widely different things on the system.”

Ahern continues, “MATLAB is, far and away, the most used code on Nautilus. A researcher from the University of Oklahoma runs simulations on other XSEDE systems, producing extreme amounts of data using her simulation code. The data that is gathered is then analyzed on Nautilus.”

Nautilus is also used to analyze ecology data. Ecologists are traveling throughout Great Smoky Mountains National Park, making observations of thousands of species of plant and animal life at various locations throughout the park, and making correlations between these multi-species observations. The team is trying to establish patterns in geographic areas where various species are likely to be found, as well as the relationships between those zones and the various species. Using Java, they've found a way to take advantage of the shared memory of Nautilus to analyze this tracking data in a way they could not effectively do on smaller machines.

Brand New From the Ground Up

Nautilus was newly integrated into the NICS infrastructure along with other machines. Although NICS had a very modest visualization cluster before, Nautilus serves a much more ambitious purpose and is distinguished by its very large, coherent shared memory as its key feature. The system is allocatable for new users; it is not being used for any continuation projects from other systems. In fact, its use is increasing. "We are getting the word out to let researchers nationwide know that it's available. Marketing is a bit of a chicken and egg problem. Utilization was not as high initially. People may have learned about it through other venues," states Amy Szczepanski, NICS RDAV Outreach Coordinator. Allocations have generally come through the peer review process of Teragrid, the precursor to XSEDE. "In this process, different machines are recommended to researchers for different processes. Approximately 50% of the system's time is allocated in this manner, so we've also been providing discretionary allocations and approaching researchers directly," Szczepanski adds.

Ahern states that Nautilus is certainly the largest memory footprint he's ever touched, though there are larger SGI UV systems at other national supercomputing sites. "I expect that fact to remain true for a while longer. 4TB is outside the spectrum for what you can really buy commodity-wise. This was the most difficult machine we've ever stood up, but it's delivered some really cool science we would not have been otherwise able to deliver," concludes Ahern.

Szczepanski adds, "Because of the architecture of the UV, it really opens up large data analysis from outside the typical HPC world. Researchers who have been using desktop workstations can now use these large shared resources, which is very cool. They can take existing workflows to a large system, and they typically 'just work.' We do have staff available to help researchers use the system appropriately, but it's much easier than throwing someone on a cluster." RDAV has a staff of about ten people who support visualization, data analysis, I/O and workflows, as well as training and development, front line support and system administration.

Nautilus Harnessed for Humanities Research, Future Prediction[1]

One amazing use of Nautilus was the recent project called *Culturomics 2.0: Forecasting Large-Scale Human Behavior Using Global News Media Tone in Time and Space*. This study was undertaken to find a way to use tone and geographical analyses methods to yield new insights about global society. This project aims to provide opportunities for societal research at a global scale, and to enable the existence of a "warning bell" before crises occur.

Typically, scientists run code on a huge number of processor cores for a long period of time to get their results. In the *Culturomics 2.0* study, Kalev H. Leetaru, Assistant Director for Text and Digital Media Analytics at the Institute for Computing in the Humanities, Arts and Social Science at the University of Illinois, focused on the large coherent shared memory feature of the SGI UV, in addition to its processor power. Leetaru used standard analysis packages along with Perl scripts to piece together data from differing sources.

Leetaru claims that his analytics experiment had allowed him to successfully forecast recent revolutions in Tunisia, Egypt and Libya. Leetaru also reported that he had been able to foresee stability in Saudi Arabia (at least through May 2011), and retroactively estimated Osama Bin Laden's likely hiding place to within a 200 kilometer radius.

Leetaru says that using a large shared memory system like Nautilus was the key to achieving his research goals. "A system like Nautilus allows researchers more flexibility as they seek to take advantage of vast computing power to analyze 'big data' in innovative ways. My goals with this project represent a perfect example of data-intensive problem solving in research," he states.

Leetaru went on to note, "With a large shared memory machine, you don't have to worry about memory. I never had to worry about writing MPI code to distribute memory across nodes. It's like it's infinite. With a quick script I could grab all locations that mentioned 'Bin Laden' since he first started to appear in the news around 10 years ago, and map it over time or in different ways. It boiled down to writing easy Perl scripts, running in a matter of minutes. If I didn't have all that memory, it would have taken weeks or months with each iteration, so one benefit is that leveraging so much hardware allows you to do things quickly."

[1]http://www.hpcwire.com/hpcwire/2011-09-09/nautilus_harnessed_for_humanities_research,_future_prediction.html

About XSEDE

The Extreme Science and Engineering Discovery Environment (XSEDE) is a five-year, $121 million project supported by the National Science Foundation (NSF), and is the most advanced, powerful and robust collection of integrated advanced digital resources and services in the world. It is a single virtual system that scientists can use to interactively share computing resources, data and expertise. Scientists and engineers around the world use these resources and services to make us all healthier and safer, and lead better and more comfortable lives. XSEDE lowers technological barriers to the access and use of computing resources. Using XSEDE, researchers can establish private, secure environments that have all the resources, services and collaboration support they need to be productive.

The XSEDE partnership is led by the National Center for Supercomputing Applications at the University of Illinois.

XSEDE replaced Teragrid in July 2011, and the NSF has been making awards to various institutions through this program since then. XSEDE is based upon designing an interconnected computing infrastructure, aiming to unify all processes, such as user accounts, security, high speed networking and more, all sitting on top of the services each center provides. XSEDE is much more than a portal — it is the sum of infrastructure plus computer systems.

About NICS

The National Institute for Computational Sciences (NICS) at the University of Tennessee is the product of a $65 million award by the National Science Foundation.

The mission of NICS is to enable the scientific discoveries of researchers nationwide by providing leading-edge computational resources and education, outreach and training for under-represented groups. NICS is a joint project of the University of Tennessee and Oak Ridge National Laboratory. For more information, please visit www.nics.tennessee.edu.

Global Sales and Support: sgi.com/global