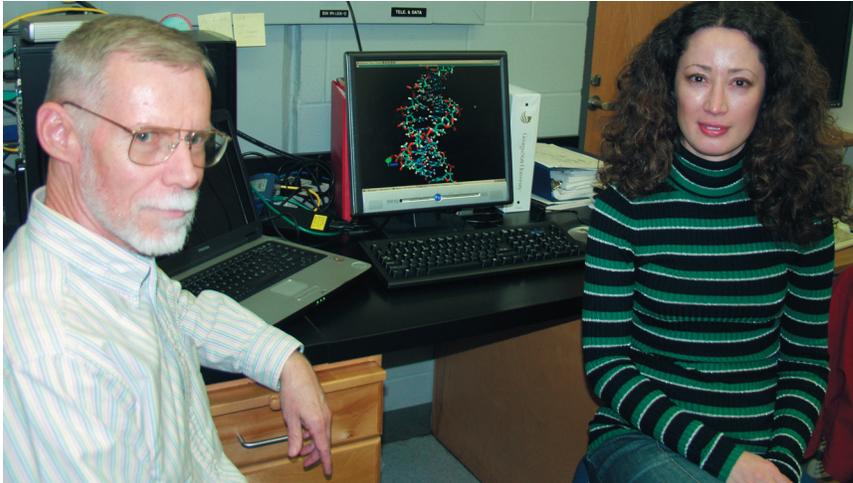


## IBM helps Georgia State University contribute a powerful building block to a supercomputing grid



In Dr. David Wilson's lab, SURAggrid supports Ph.D. student Maryam Rahimian in molecular dynamics simulation on bio-molecules.

With more than 60 member schools, the Southeastern Universities Research Association (SURA) is committed to using technology to foster collaborative research in a variety of fields, from science and engineering to the liberal arts. In 2003, Georgia State University was among a group of pioneering SURA members to build SURAggrid—a shared computing infrastructure that leverages a rich set of distributed resources for faculty research. Located in the heart of downtown Atlanta, Georgia State University is a leading urban research institution with more than 50,000 students enrolled annually. Today, Georgia State University is among 28 SURA member schools that have created a supercomputing system that shares processing power from heterogeneous systems to conduct research that at Georgia State includes neuroscience, proteomics, risk management, sub-atomic structure and more.

### Overview

#### ■ Challenge

Provide researchers from a collection of universities throughout the southeastern United States with increased processing power through a supercomputing grid infrastructure

#### ■ Solution

Deploy IBM System p5™ 575 cluster nodes, an IBM System Storage™ disk system and tape library, the IBM AIX 5L™ operating system and a full array of IBM system administration tools, including IBM Tivoli® Workload Scheduler LoadLeveler® software

#### ■ Key Benefits

- The p5-575 system enables numerous faculty members to conduct diverse projects simultaneously while helping to dramatically reduce the time to complete complex tasks
- The DS4800 offers faculty fast access to data and a reliable, security-rich platform for backing up research

*“We’re confident that by integrating the p5-575 into the SURAggrid, we can provide SURA member schools with the added processing power they may need.”*

— Art Vandenberg  
Director, Advanced Campus Services  
Georgia State University

Previously, researchers' work was constrained by the resources available at their own university or a national super-computing facility. "Before SURAGrid was established, a physicist might have to book time at and travel to the Thomas Jefferson National Accelerator in Virginia," says Art Vandenberg, director of Advanced Campus Services at Georgia State. "Otherwise that researcher would have to try to find or buy resources locally. Though some faculty found what they needed, many more had difficulty gaining access to the computational power they required for their research. By combining resources with SURAGrid, researchers across the southeast now can have greater access to the right resources."

Each member of the SURAGrid network offers a percentage of its computing cycles to the grid. As a result, the entire network benefits when a school deploys a new, more powerful system. "Even if we commit 20 percent of our cycles to the grid," says Vandenberg, "we stand to receive more than 200 percent back if we access 10 systems from other universities."

### **IBM makes a strong commitment to SURAGrid**

IBM realized the important research potential of SURAGrid and began to work with Georgia State and other SURA member schools on ways to bolster the power of computing resources available with IBM products. "The opportunity to work with IBM came at a very fortuitous time for us," says Victor Bolet, analyst programmer and p5-575 systems administrator. "Our commodity

cluster was reaching its end of life, and our SMP architecture-based machine was ancient by most standards."

Through an agreement between IBM and SURA, Georgia State was able to acquire IBM System p5 575 cluster nodes that it could use locally and share with the grid. The school's system includes eight nodes, each node with eight dual-core 1.9 GHz IBM POWER5+™ processors. The eight-node system has a total of 256 GB of memory, 2 GB per core. With 64-bit processor cores, simultaneous multithreading capabilities and access by each processor core to 1.9 MB of L2 and 36 MB of L3 dedicated cache, the p5-575 solution can provide the power for high-performance computing and business intelligence applications. The p5-575 is also available in an 8-core configuration in which only one core of each dual-core processor is active.

The ultra-dense 2U rack-space configuration of the p5-575 enables Georgia State to use 128 CPUs without even filling up the single rack frame. "The smaller footprint will definitely help us reduce power and cooling consumption," says Vandenberg. "At the same time, the p5-575 is simple to access and service." As part of its innovative design, the p5-575 includes a highly efficient DC power distribution module that is integrated into the lid of the node. The lid opens easily for access to the processor cores and memory. The front-end cooling module has two air-intake ventilation grids and two custom-designed blowers designed for extended life and straightforward serviceability.

*"For some projects—like large proteomics data sets or videos of student teaching from the college of education—the faculty needs more and more disk space. With the DS4800, we now have a scalable solution."*

– Art Vandenberg  
Director of Advanced Campus Services  
Georgia State University

The p5-575 is also easily scalable within the system or within the cluster, giving Georgia State the option to expand its processing power in the future. The school could add up to 256 GB of memory per node and up to 600 GB of disk capacity with the hot-swappable internal disk drives. Housing up to 12 nodes in a system frame, the school could add frames up to 128 nodes, or 2,048 processors. In March 2007, Georgia State acquired two more compute nodes for the frame, with planning underway to further expand the processing capability of the cluster.

Georgia State is running its applications on the IBM AIX 5L operating system for its reliability and flexibility. "The IBM AIX 5L offering is a mature operating system," says Vandenberg. "We can run the same applications that we previously ran on Linux®. But with AIX 5L, researchers and administrators can also take advantage of many system capabilities that were just not available through Linux."

*“We know that we can contact IBM people and get help whenever we need it. That kind of support gives us the confidence to build on the architecture. We have a strong, long-term relationship.”*

– Art Vandenberg  
Director of Advanced Campus Services  
Georgia State University

Georgia State also uses IBM Cluster Systems Management (CSM) software and IBM Tivoli Workload Scheduler LoadLeveler software to help manage the p5-575 system. CSM is designed to help reduce the complexity of systems management by providing management capabilities from a single point of control. LoadLeveler is a job-scheduling system that can help users to run more jobs in less time by matching each job's processing needs and priority with available resources. “CSM and LoadLeveler have been huge assets in simplifying systems management,” says Bolet.

#### **IBM System Storage DS4800 offers fast and reliable access to research data**

Georgia State selected the high-performance IBM System Storage DS4800 to provide fast, reliable and unified enterprise storage for the school's diverse storage needs. Designed to deliver up to 1,600 MBps throughput with up to

550,000 input/output per second (IOPS) and supported by an IBM eServer™ pSeries® High Performance Switch and two IBM TotalStorage® SAN16B-2 16-port switches, the DS4800 offers researchers rapid access to data. The DS4800 is also easily scaled up to 224 Fibre Channel (FC) or 224 Serial ATA (SATA) disk drives, for a total capacity of up to 67.2 TB of FC storage or 112 TB of SATA storage, so Georgia State IT administrators can expand capacity as requirements grow. “For some projects—like large proteomics data sets or videos of student teaching from the college of education—the faculty needs more and more disk space,” says Vandenberg. “With the DS4800, we now have a scalable solution.”

In April 2007, Georgia State plans to add 4 TB of FC storage to support the research of Dr. Xiaochun He, Physics and Astronomy. Dr. He's research uses large data sets and generates complex cosmic ray simulation data, requiring terabytes of storage.

The DS4800 also provides the IT group a way to help faculty members back up data in a reliable location. “Previously, faculty would just back up their own data,” says Vandenberg. “With the DS4800, we can provide backup services so that research is less susceptible to loss. In the future, we could even use the DS4800 to institute a distributed storage model and establish a disaster recovery model among University System of Georgia schools for even greater data protection.”

Georgia State uses IBM Tivoli Storage Manager (TSM) software to help automate data backup and restore functions from a centralized console. Tivoli Storage Manager is designed to protect computers running a variety of different operating systems, on hardware ranging from notebooks to mainframe computers and connected together through the Internet, wide area networks, local area networks or storage area networks. Georgia State also uses TSM to integrate the DS4800 with its System Storage TS3310 tape library. “With the DS4800, Tivoli Storage Manager and the TS3310, we have a robust, reliable and integrated IBM storage solution,” says Vandenberg.

#### **IBM support gives Georgia State confidence to take on new challenges**

IBM helped Georgia State deploy the system rapidly. “IBM delivered the equipment, helped us set it up and made sure we had it up and running quickly,” notes Vandenberg. “They also gave us the initial training we needed for the AIX 5L operating system. Only two months from the time we signed the paperwork, we had the first applications running. That could not have been done without outstanding IBM service.”

Now that the system is deployed, the Georgia State IT group knows it can rely on IBM for post-sales support. “We know that we can contact IBM people and get help whenever we need it,” says Vandenberg. “That kind of support gives us the confidence to build on the architecture. We have a strong, long-term relationship.”

The Georgia State team also works with IBM to acquire new applications for its system through the IBM Academic Initiative, which provides educators and academic institutions an easy way to keep up with the latest technologies. "The IBM Academic Initiative makes ordering software simple," says Vandenberg. "We can just log onto the IBM Academic Initiative site, order what we need and a few days later we have stacks of CDs and licenses there for us. It is pretty remarkable."

### **The p5-575 is poised to power SURAggrid**

The Georgia State IT team is finding that the p5-575 has been very successful at accommodating a wide range of research applications at once. Georgia State research applications have been running around the clock since December 2006 on the IBM System p5 server—and the school's researchers are already taking advantage of the new capabilities for research that improved processing provides.

In Dr. David Wilson's lab, for example, chemistry Ph.D. student Maryam Rahimian uses Amber software for molecular dynamics simulation on biomolecules. While a 10 nsec dynamic calculation of 10 bp DNA took almost six days on Georgia State's previous system, it now only takes one and a half days on the p5-575. Furthermore, Dr. Martin Grace, James S. Kemper Professor in the Department of Risk

Management, is using the p5-575 to run SAS software for risk modeling based on insurance data. Initial tests have shown a significant increase in computation speed after moving to the new infrastructure.

The benefits of SURAggrid extend outside Georgia State as well. "The multi-threading capabilities are very welcome here," says Bolet. "With the p5-575, two classes of problems can run at the same time, all in a single machine. That capability is a huge benefit for us since this is essentially a general-purpose computing system that will be used by researchers both here and across the grid." Adds Vandenberg, "A job that could typically take 4 days was able to run in just 18 hours, even without specially tuning the system. We realize that our researchers will quickly push the limits again, but we're confident that by integrating the p5-575 into SURAggrid, we can provide SURA member schools with the added processing power they may need."

### **For more information**

For more information, contact your IBM sales representative or IBM Business Partner, or visit us on the Web at:

- [ibm.com/storage](http://ibm.com/storage)
- [ibm.com/systems/p](http://ibm.com/systems/p)



© Copyright IBM Corporation 2007

IBM Systems and Technology Group  
Route 100  
Somers, NY 10589  
U.S.A.

Printed in the United States of America  
May 2007  
All Rights Reserved

IBM, the IBM logo, AIX 5L, eServer, LoadLeveler, POWER5+, pSeries, System p5, System Storage, Tivoli and TotalStorage are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries or both.

Linux is a registered trademark of Linus Torvalds in the United States, other countries or both.

Other company, product or service names may be trademarks or service marks of others.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program or service is not intended to imply that only IBM's product, program or service may be used. Any functionally equivalent product, program or service may be used instead. Offerings are subject to change, extension or withdrawal without notice.

All client examples cited represent how some clients have used IBM products and the results they may have achieved. Performance data for IBM and non-IBM products and services contained in this document was derived under specific operating and environmental conditions. The actual results obtained by any party implementing such products or services will depend on a large number of factors specific to such party's operating environment and may vary significantly. IBM makes no representation that these results can be expected or obtained in any implementation of any such products or services.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS-IS" WITHOUT ANY WARRANTY, EITHER EXPRESSED OR IMPLIED.