PBS Pro™
Altair Engineering

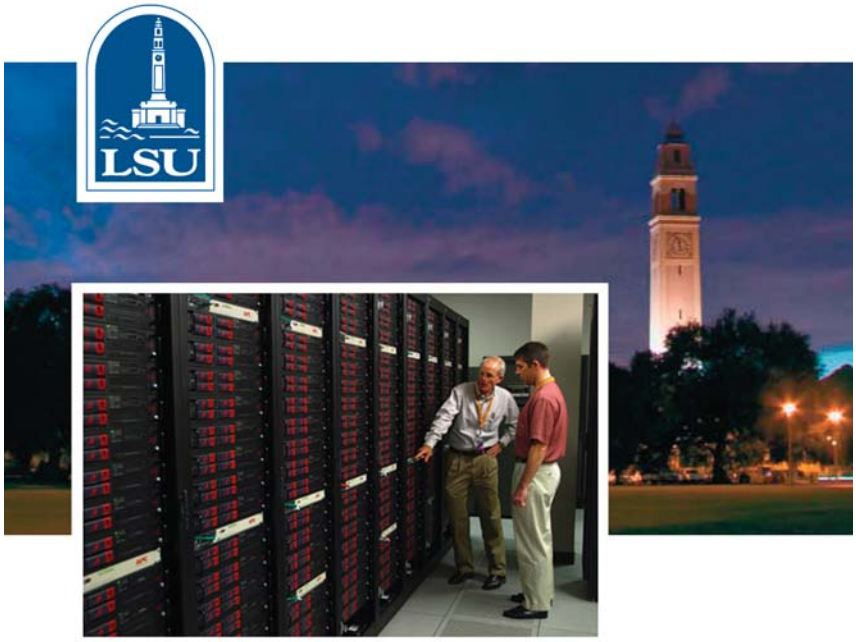

PBS Pro at LSU's Center for Computation & Technology:

Workflow Management for A Statewide Supercomputing Infrastructure

When then-governor Mike Foster came out with his Louisiana: Vision 20/20 program in 1998, he set forces in motion that are turning the state into a research and scientific high performance computing (HPC) powerhouse. Part of the plan was to drive economic development with information technology through higher education. As a result, five of the state's universities share an annual state appropriation of $25 million as funding for a statewide HPC infrastructure.

Part of the plan was to drive development with information technology through higher education.

Louisiana State University (LSU) began receiving nearly $7 million annually from this fund. Through its Center for Applied Information Technology and Learning (CAPITAL), interim director Dr. Joel Tohline from the Department of Astronomy and Physics started funding key technology initiatives across campus. The campus network backbone was migrated to Gigabit Ethernet; an information technology residence hall was created; the College of Business created a simulated stock trading floor with near-real-time feeds from the NYSE; and in 2002 the university invested in its first large-scale computational cluster.

James River Technical, Inc., 4439 Cox Road, Glen Allen, VA 23060
Phone: +1.804.935.0150 • Fax: +1.804.935.0165 • www.jrti.com • altair@jrti.com

Welcoming SuperMike and A New Director

SuperMike, a 512-node, 1024-processor Beowulf cluster, made a fast-track arrival at LSU. It took just 10 months from its conception to make it an operating reality. The system went operational in August 2002, benchmarked at 2.2 teraflops, and was ranked 17th among the world's fastest computational systems by Dr. Jack Dongarra, who helps maintain the Top500 list.

LSU then focused on attracting faculty members and researchers that could raise the bar in terms of research. Dr. Ed Seidel was recruited in August 2003 from the Max-Planck-Institut fur Gravitationsphysik (Albert-Einstein-Institute) to be the full-time director of the research facility, which was renamed the Center for Computation & Technology (CCT). Many of his research staff came with him from Germany. Seidel had interesting qualifications: A physics degree from Yale, eight years experience at the National Center for Supercomputing Applications (NCSA), and a need for massive computational power to pursue his own interest in black hole collision research.

When SuperMike was originally provisioned, LSU wasn't interested in investing in commercial workload management software. They wanted an open workload manager, and turned to Open PBS.

"That was a pretty quick decision, and it was based mostly on PBS's reputation," says Brian Ropers-Huilman, Assistant Director, High Performance Computing and Computation at CCT. "We had prior experience with it. We had looked at Sun GridEngine and LSF, and we were much happier with Open PBS. It was easy to install and easy to configure, and it was going to suit our needs handily."

A second Linux cluster, with hardware identical to SuperMike except in the number of processors, was purchased by LSU's College of Basic Sciences. This 128-node, 256-processor system, named SuperHelix, also became the responsibility of CCT, and was also running Open PBS. Both clusters were available to any on-campus researchers with a need for massive processing power. Collaborators from CalTech, Stanford, and other schools also had cluster accounts.

Moving to PBS Pro: Scalability and Professionalism

In December 2003, Ropers-Huilman officially joined CCT to oversee a sweeping hardware and software upgrade that makes SuperMike a totally new machine. Processors were upgraded from Xeon 1.8 to Xeon 3.06. Myrinet adapters were migrated from B to D cards. Red Hat Enterprise Linux 3.0 replaced Red Hat 7.3, and the compiler infrastructure was upgraded.

Dr. Gabrielle Allen, who had worked in Germany with Dr. Seidel and was recruited to CCT, contacted Altair Engineering about potential licensing agreements for PBS Pro.

"One reason we went from Open PBS to PBS Pro was to get better scalability," says Ropers-Huilman. "We were having trouble deploying jobs on anything larger than 128 nodes. We had massive timeouts, and it was taking 15 or 20 minutes just to deploy some of these bigger jobs. And depending on a node's history state, we sometimes were unable to get a job deployed at all, because we had too many open network sockets as the jobs were being deployed from the master node.

"We knew specifically that PBS Pro had addressed those issues. We were pleased to see that PBS Pro had the scalability, reliability, and resiliency features we needed to run a large environment. The migration to PBS Pro went smoothly and easily, and Altair's educational grant program allowed LSU to get commercial-quality support at a very reasonable price."

Even before the upgrade was complete and the machine became publicly accessible

Putting SuperMike to Work

Before the upgrade, 200 users had accounts on SuperMike. Between 50 and 60 were active at any given time, and 15-20 were core users who were on the system almost daily. Even before the upgrade was complete and the machine became publicly accessible, a dozen accounts were active in Friendly User mode.

What kind of work is SuperMike doing? LSU Departments such as Physics, Civil Engineering, and Chemistry are frequent users, as well as the on-campus Hurricane Center. Subjects range from sawmill log-cutting simulations (using CT scan data) to fluid flow in the eye to the effects of combustion-generated nanoparticles on mucous membranes. Applications range from home-grown code to NWChem, used for molecular dynamics studies, and Cactus, an applications framework for high-performance computing. And Dr. Seidel's research team uses all SuperMike's 1024 processors to do black hole research.

Statewide Clusters, 40 Gigabit Networks, and 5000 PBS Pro Licenses

LSU's licensing agreement with Altair covers 5000 PBS Pro licenses. SuperMike, SuperHelix, and MiniMike (a 32-processor development cluster) account for only 1578 processors. Small clusters in the Civil Engineering Department, the Mathematics Department, and other campus locations were also licensed. CCT had other destinations in mind for the rest of the licenses.

In 2005, each of the five other Louisiana universities that have been receiving funds from the Vision 20/20 plan will install a cluster similar to SuperMike. Together they will account for the rest of the 5000 PBS Pro licenses. Two other LSU initiatives will bring all this processing power together.

LSU, through grants from the Board of Regents and efforts of the Southeastern Universities Research Association (SURA), bought a seat on the National Lambda Rail (NLR), the new premier national research network, which will operate over a 40 Gbit fiber optic core. The network route was extended to Louisiana, which triggered interest from other universities across the southern U.S. and resulted in the extension of the network. The link will enable LSU scientists to share remote and potentially large resources across an extremely high-speed network with colleagues across the country or around the world.

LSU followed this up by raising initial funding for the Louisiana Optical Network Initiative (LONI), aimed at deploying a similar 40Gbit network across the state. In early September 2004, Governor Kathleen Blanco surprised an audience of researchers during the LONI Forum by announcing full support of the network with $40 million over 10 years. The phased deployment, which begins late in 2004, will first link CCT to the Louisiana universities where clusters will be installed. Corporations and government agencies are being invited to invest in the network and participate in its use. The result will be a networked virtual supercomputing resource with 4000 processors – connected to the National Lambda Rail network.

As CCT continues to expand its role in supercomputing, it is changing its mode of operation, and PBS Pro will play a key role.

"CCT is migrating to an allocations model and an external committee will consider requests and award hours on the machine," says Ropers-Huilman. "PBS Pro will give us more accurate tracking of usage hours for project accounts. It will enable CCT to operate in the same professional way as other large U.S. supercomputing centers."