

CREATING CYBERINFRASTRUCTURE

The TeraGrid is the world's most comprehensive distributed cyberinfrastructure for open scientific research. As a major partner in this National Science Foundation program, PSC plays a leadership role in shaping the vision and progress of the TeraGrid.

Big Ben Goes To Work



PSC's Cray XT3 system "BigBen" became a TeraGrid production resource in October 2005. It was the first Cray XT3 anywhere and remains the only one available to NSF researchers. It comprises 2,090 processors with an overall

peak performance of 10 teraflops: 10 trillion calculations per second.

On a per-processor basis, BigBen is 2.4 times faster than LeMieux, PSC's six-teraflop HP Alphaserver system that pre-dates it and is, like BigBen, "tightly-coupled" — designed to optimize inter-processor bandwidth, the speed at which processors share information. More than sheer processor speed, BigBen's primary technological advance is its superior inter-processor bandwidth.

This is a large advantage for projects that demand hundreds or thousands of processors working together. Over the past year, because of this capability, BigBen has demonstrated performance 10 times or more better than LeMieux on a number of applications. Because of this capability also, BigBen has proven itself to be a champion at "scaling" — the ability to use a large quantity of processors without seriously reducing the per-processor performance.

Several research groups, including Klaus Schulten's group at the University of Illinois, Urbana-Champaign (p. 18) and Michael Klein's group at the University of Pennsylvania (p. 22), have found that BigBen scales to twice as many processors as the same applications on LeMieux, an improvement that, along with faster processors, represents a big gain in capability, and has led to many research successes (pp. 18, 22, 26, 30, 34, 43, 45, 46).

Researchers with large-scale parallel projects quickly caught on to BigBen's advantages. Over its first year as a production resource, nearly half of BigBen's usage has been for projects that use 1,024 processors or more, and at the last national allocation meeting, it was the TeraGrid resource most "oversubscribed" — demand in excess of available time.

BigBen's predecessor as PSC's lead system, LeMieux, remains an actively used TeraGrid resource, and PSC fulfills a unique function within the TeraGrid in providing two tightly-coupled systems. Between them, during the first half of 2006, LeMieux and BigBen provided 40 percent of overall TeraGrid usage.



TeraGrid™



TeraGrid Resource Providers

- Indiana University
- National Center for Supercomputing Applications
- National Center for Atmospheric Research
- Oak Ridge National Laboratory
- Pittsburgh Supercomputing Center
- Purdue University
- San Diego Supercomputer Center
- Texas Advanced Computing Center
- The University of Chicago/Argonne National Laboratory



Jim Kasdorf, PSC director of special projects



PSC staff whose work contributes to TeraGrid include (l to r): Laura McGinnis, Michael Schneider (seated), Nathan Stone, Josephine Palencia, Sergiu Sanielevici, Shandra Williams, Kathy Benninger, Rob Light, David O'Neal (on floor), R. Reddy, Derek Simmel, Rich Raymond, Jim Marsteller.

TeraGrid 2006: PSC Rolls Up Its Sleeves

In several projects, PSC staff have helped to advance the technological infrastructure of the TeraGrid:

- In a major TeraGrid effort, PSC staff applied security controls for jobs run on LeMieux, PSC's terascale system, from the "portal" of the NanoHub Science Gateway. The solution involved adapting "community shell" software developed at NCSA. This PSC project is the first effort within TeraGrid to implement a security model that reconciles the community-wide reach of a Science Gateway with the secure environment of a large-scale system like LeMieux, experience we can apply to other Gateways.
- PSC network staff deployed a TeraGrid version of PSC's NPAD diagnostic service (see p. 14) on three of TeraGrid's network monitoring computers, and users have found it to be an effective tool.
- PSC staff developed software to automate the execution of "speedpage," a TeraGrid routine to measure file-transfer performance among TeraGrid sites with Globus middleware. As a result, speedpage now functions effectively to give users advance information on the file-transfer rate they can expect with these middleware routines.
- A reliable, versatile wide-area-network (WAN) file system is an important TeraGrid objective. Highly experienced in file systems across a range of architectures, PSC's systems & operations staff have implemented two separate WAN file-system projects this year. The GPFS-WAN (Global Parallel File System) developed at SDSC is now operational across TeraGrid. In a forward-looking project, PSC has led efforts to create a testbed for the flexible, open-source Lustre-WAN file system, used on PSC's TeraGrid resources. The Lustre-WAN testbed is now operational at PSC, Oak Ridge National Laboratory, Indiana University and NCSA.

Support For TeraGrid

PSC is actively involved in directing the progress of TeraGrid. Scientific director Ralph Roskies serves on the executive steering committee of the GIG, the Grid Infrastructure Group that guides TeraGrid, and scientific director Michael Levine is the principal contact for PSC as one of nine "resource provider" sites.

Several of PSC's key staff are committed heavily to TeraGrid work. Sergiu Sanielevici, PSC director of scientific applications and user support and one of PSC's most experienced computational scientists, is TeraGrid Area Director for User Support. He manages the TeraGrid-wide user-support team and coordinates TeraGrid's ASTA program — Advanced Support for TeraGrid Applications.

Jim Marsteller, who heads PSC's security effort, chairs the TeraGrid Security Working Group. He's responsible for developing sound security practices for TeraGrid and to conduct risk assessments and incident response as necessary.

PSC's manager of systems/software projects, Laura McGinnis, has made significant contributions to the TeraGrid accounting system, which handles the complex task of tracking allocations at nine resource-provider sites with many different systems.

xt3dmon

On a 42-inch, high-definition, liquid-crystal screen, the NSF Office of Cyberinfrastructure displays a live readout from PSC-developed software called xt3dmon. This program, written by interns in PSC's training program, provides an onscreen realtime picture of BigBen's processor activity. The 3D visual interface allows PSC staff to identify patterns in how jobs are allocated among processors to improve job performance and scheduling.

