Success Story

QPSF Data Grid

Queensland Parallel Supercomputer Foundation: SGI® Data Access Solutions Power Research Statewide

"DMF is fantastic, It's the seamless way the whole thing works and the performance you get out of it. It's like disk space is infinite. It just works. It's one of the jewels in SGI's crown."

– Dr. Ian Atkinson, Manager, HPC, James Cook University

Australia's big eastern state of Queensland is renowned for its climate, which varies from temperate to tropical; its natural resources, which range from sugar cane to coal; and the Great Barrier Reef, which parallels nearly the entire Queensland coastline and serves as a global tourist magnet. Visualization of the reef is just one of a multitude of projects that brought six Queensland universities together in a supercomputing consortium with SGI technology at its heart.

The QPSF: Pooled Compute Power for Research

The Queensland Parallel Supercomputer Foundation (QPSF) is based on a simple philosophy: pool funds, acquire major supercomputing and visualization systems, and share these resources over a broadband network. QPSF member institutions include the University of Queensland, James Cook University, Griffith University, Queensland University of Technology, Central Queensland University, and the University of Southern Queensland. These institutions, spread across hundreds of miles, share SGI resources—SGI Reality Center visualization, MIPS® processor-based supercomputing, Intel® processor-based supercomputing, and data storage systems—across a government network.

The Queensland state government's Smart State initiative, which supports the state's knowledge-based industries, has provided QPSF with significant technology funding—an acknowledgement that high-performance computing, visualization, and information sharing will play an important role in building world-class research and innovation capability in Queensland. The Australian national government is supporting QPSF grid access by funding advanced research networks that will take data

The University of Queensland SGI SAN: Cross-Campus File Sharing in an HPC Environment

QPSF resources at the University of Queensland include a wide range of SGI computing systems, all of which are part of its SGI storage area network (SAN). The systems include:

- A 64-processor SGI® Origin® 3000 server, the main computation node for QPSF and the university, and the driver of the SAN
- An SGI® Onyx® 3200 system that drives a room-based immersive SGI® Reality Center™ facility
- Two SGI® Origin® 200 file servers
- An SGI® Origin® 2100 server used for bioinformatics research by the Institute for Molecular Bioscience
- SGI® Altix™ 3700 and SGI® Altix™ 3300 systems, running Linux® on Intel® Itanium® 2 processors

A 48-processor SGI Origin 3000 server, used by the university for earthquake simulation, is also tied into the SGI SAN. This system is scheduled to be replaced by an SGI Altix 3700 system.

Fiber-optic links between buildings provide cross-campus access to 1.5TB of online Fibre Channel disk storage and a 20TB tape storage environment, including a StorageTek® 9310 automated library with five StorageTek® 9840 Fibre Channel tape drives. Researchers use the SGI® CXFS™ filesystem to access and share a single copy of their data on all systems in the SAN. SGI® Data Migration Facility (DMF) provides hierarchical data storage management on all SAN-linked systems. DMF automatically migrates files between disk and tape, providing a massively scalable virtual storage environment that significantly enhances productivity.

"DMF is fantastic," says Ian Atkinson. "It's the seamless way the whole thing works and the performance you get out of it. It's like disk space is infinite. It just works. It's one of the jewels in SGI's crown."

communications within and between QPSF universities to multigigabit levels by late 2004.

"Grid computing is clearly the most desirable approach for us," says Dr. Ian Atkinson, manager of High Performance Computing at James Cook University. "Grid concepts are essentially collaborative. We want to make use of our diverse intellectual resources as well as our compute resources. Pure peer-to-peer computing wouldn't be very effective. We're much more comfortable with large, really capable systems connected to a network."

QPSF is developing an access grid node at each campus to enable student collaboration. The first instance is now operating over a 4Mb network. Atkinson describes the nodes as "a big, full-bore videoconferencing system" that enables scientists at different locations to talk and work together. QPSF has also built compute portals that make remote computing a transparent process for some researchers, freeing them from concerns about where and when their jobs are running.

SGI® Altix™ Systems: Adding Linux Clusters to the Mix

Most of QPSF's compute power resides at the University of Queensland, where SGI high-performance systems have provided computing muscle for research programs since 1995. In 2001, QPSF acquired a 64-processor SGI® Origin® 3400 shared-memory supercomputer with 64GB of memory—the first 600 MHz Origin® system delivered anywhere. It is accessed across the QPSF network for heavy-duty number crunching in bioinformatics, computational biology, computational chemistry, earth systems simulation, computational fluid dynamics, and other areas of research. A shared data storage solution using SGI CXFS

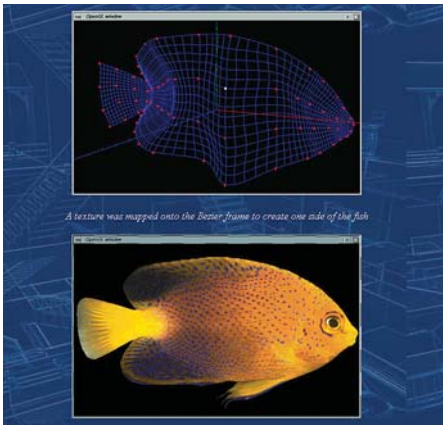

Image courtesy of Queensland Parallel Supercomputer Foundation

serves both the Origin 3400 and the Silicon Graphics® Onyx2® systems in the nearby Reality Center facility.

When SGI announced the availability of the SGI® Altix™ 3000 open-source Linux clusters based on the Intel Itanium 2 processor, QPSF moved quickly to complement its MIPS processor-based SGI system. The foundation ordered two SGI Altix 3000 shared-memory systems for installation at the University of Queensland: a 64-processor system with 64GB of memory for batch processing and a 16-processor system with 16GB of memory primarily for interactive use by students. "QPSF is delighted to partner with SGI in bringing the world's second Itanium2-based Altix system into a production environment." says Professor Bernard Pailthorpe, CEO of QPSF.

"We had research groups lined up waiting to use the Altix systems," says Atkinson. "We've continued to expand our SGI supercomputing capacity based on the MIPS processor, for which there is a strong demand, but we know the SGI open-source Altix concept will be highly successful here." QPSF research groups that have expressed interest in using the Altix architecture include the bioinformatics group at the University of Queensland's Institute for Molecular

"We had research groups lined up waiting to use the Altix systems. We've continued to expand our SGI supercomputing capacity based on the MIPS processor, for which there is a strong demand, but we know the SGI open-source Altix concept will be highly successful here."

– Dr. Ian Atkinson, Manager, HPC, James Cook University

Bioscience, a systems modeling group, and a computational physics group, as well as a number of computational chemists and engineering groups.

"The approach that SGI has taken with their new Linux systems is very exciting," says Atkinson. "The benchmark results for the Altix architecture were amazing. SGI Altix 3000 systems give us a complete range of benefits, including fast processor performance and the SGI NUMAlink flat memory space that gives individual users a great deal of memory when they need it. The economics of this architecture for grid computing are extremely important to us as well. A single employee can manage big 60- and 80-processor systems, which is impossible when you tie many small Linux machines together."

The Altix systems form part of the QPSF SAN. The heterogeneous capability of CXFS enables researchers to share files with the other SGI HPC and visualization resources on the SAN. "People can be on the visualization system, the IRIX systems, or the Altix systems, with their servers attached to SAN, seeing that as contiguous disk space," says Atkinson. "Backing up and all those infrastructure issues just disappear. It's really cost-effective—not only in terms of good performance, but it enables better productivity because of its simplicity. It's a

safety issue, too, because the data's all in one place and can be backed up more effectively and inexpensively in terms of staffing costs."

Visualizing the Reef Remotely with Reality Center and VAN Technology

Given Queensland's geographically dispersed researchers and institutions and the length of the Great Barrier Reef, it is inevitable that QPSF computing technology would include remote access to the SGI Reality Center system at the University of Queensland for visualization studies of marine ecology. The three-projector, curved-screen Reality Center system, driven by an SGI® Onyx® 3000 visualization supercomputer, is used by reef biologists at Queensland universities to enhance the understanding of coral habitat by viewing simulations over a Visual Area Network (VAN).

"Traditionally, reef biologists think qualitatively," says Atkinson, "which is why visualization is very important to them. It gives them a way to interface qualitative skills with quantitative information. Coral is extraordinarily hard to visualize, but we're starting to get somewhere. And when we have something to show, we process it on the Onyx system and use SGI OpenGL Vizserver technology to let people in separated, locations look at it in real time together over the access grid using any PC."

Looking Ahead: A Faster Network, Faster Processing, and Open-Source Development

New SGI systems, faster networks, and developing technologies will keep QPSF member universities at the cutting edge for the foreseeable future. For example, the University of Queensland recently tested a remote visualization link with a nearby hospital that enables doctors to process MRI images quickly on the Onyx2 system and use OpenGL Vizserver™ remote serving technology to view the results. In another example, major drug company has financed a University of Queensland research program aimed at helping schizophrenic patients. Researchers interview patients, then help them describe their symptoms by re-creating their hallucinations on the screen of the Reality Center system at QPSF. Seeing the re-created images helps therapists and scientists analyze and treat patients' conditions and provides valuable training for psychiatric practitioners.

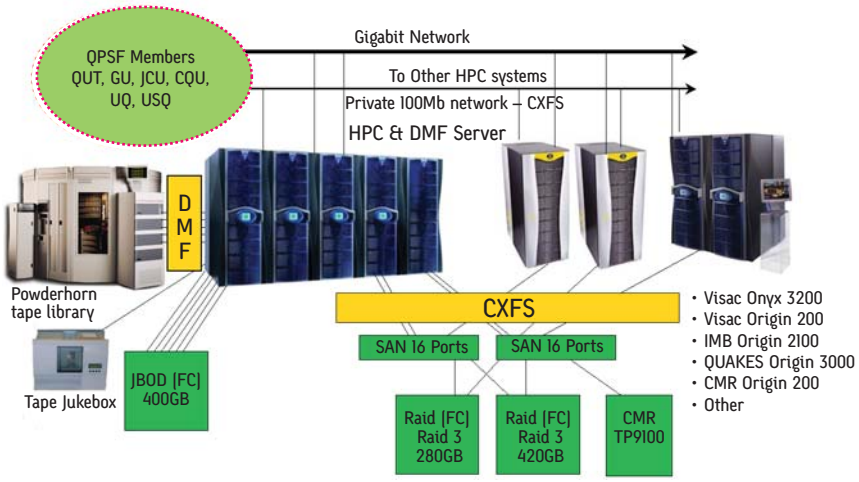

These projects are being carried out with the help of the Advanced Computational Modeling Centre (ACMC) at the University of Queensland. "There is a very strong desire by the research community here to support the 'Smart State' approach of the Queensland government," said ACMC Director Kevin Burrage. "To do that, we feel it is very important to have world-leading infrastructure in critical areas such as HPC, visualization, and data management, and that is what we have been able to achieve."

Meanwhile, interest has grown swiftly on QPSF member campuses in using the shared computing resources at the University of Queensland for biological research. "It's really interesting," says Atkinson. "We had very few people from the biological sciences using the systems, and now they're the predominant users. It's just a great wave of interest from people in the biological sciences who were running things on PCs and Macs for months and months. Then there'd be a power outage or a crash and they'd have to start all over again. Those people are really, really happy that we can provide a high level of reliability and capability and this big data I/O. In cost-effective terms, we think we win big-time."

Quality of service support was another significant factor in QPSF's tender process. "Support is clearly very important to us," said Atkinson. "There's a great deal of work going on here that we view as mission-critical. We're as impressed as ever with SGI support, which remains at an impeccably high level."

Corporate Office
1500 Crittenden Lane
Mountain View, CA 94043
(650) 960-1980
www.sgi.com

North America +1 800.800.7441
Latin America +55 11.5509.1455
Europe +44 118.912.7500
Japan +81 3.5488.1811
Asia Pacific +1 650.933.3000

3484 [01.18.2006]

J15096