PBS Professional Manages Workload for NCI Raijin, Largest Supercomputer in Southern Hemisphere

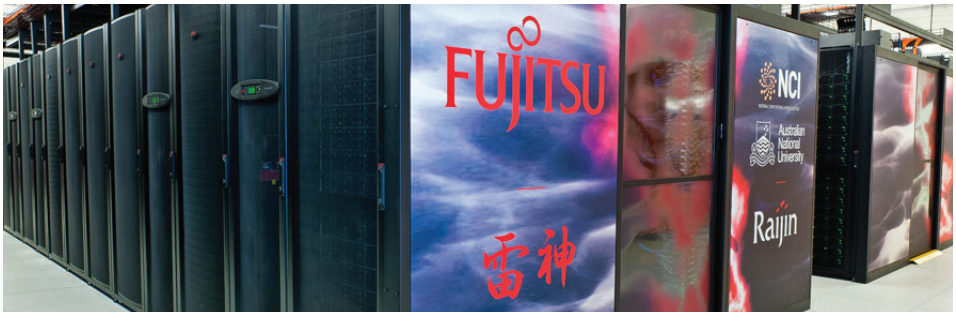

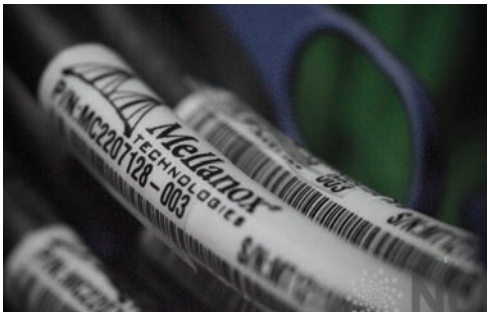

Key Highlights

Industry
Government and Research

Challenge
Meeting complex system requirements with a flexible, high-performance product

Altair Solution
PBS Professional

Benefits
- Flexibility, scalability and extensibility
- Proven performance and functionality
- Services and support excellence
- Support for growth

Customer Profile

As Australia's national research computing service, the National Computational Infrastructure (NCI) provides world-class, high-end services to Australian researchers, including access to advanced computational and data-intensive methods, support, and high-performance infrastructure. NCI supports computationally-based research with a focus on the environment, climate and earth system science in particular. Since 2007, NCI's infrastructure investments, which exceed $80M, have been provided by the Australian Government under its National Collaborative Research Infrastructure Strategy (NCRIS) and Super Science Initiatives.

The Challenge: Delivering a Highly Customized and Scalable HPC System

NCI's supercomputing system, a Fujitsu PRIMERGY cluster, is named Raijin after the Japanese God of Thunder and Lightning. In production since June 2013, Raijin is the largest supercomputer in Australia and is among the top 30 systems in the world. One of the distinguishing requirements of this system is the wide spectrum of job types, including jobs that vary both in scale (from very large-scale jobs using a substantial portion of the cluster, to memory-intensive jobs with high I/O requirements) and in completion time (from long jobs on the order of several days, to very short jobs requiring fast turn-around).

Overall system balance is important to ensure strong, scalable performance and deliver a high-quality user experience. Architecture and subsystems (such as memory, interconnect and filesystem) must scale as the software and hardware scale out, to protect application performance. In short, all elements of the system – application, hardware, subsystems and middleware – need to scale without impacting each other.

Thus, for Raijin's workload management system, NCI needed a highly scalable, flexible and reliable product that could handle both the size and complexity of its computing requirements.

NCI previously operated an in-house OpenPBS system with a locally customized scheduler and associated accounting system to manage its resources. However, it was becoming increasingly difficult to maintain

Keywords: PBS, Cluster, PRIMERGY, Raijin, NCI, Fujitsu

NCI Success Story

"PBS Professional proved its superiority as the most flexible and reliable workload manager, and Altair gave us the confidence and comfort level we need in a long-term partner. Altair really knows HPC – they understand the challenges of maintaining complex systems and know how to deliver reliable solutions that work."

Dr. Ben Evans, Associate Director,
NCI

the development and support for this system, so NCI decided to investigate new options for Raijin's workload manager.

The Selection Process: Proving Performance and Flexibility

In considering a replacement system, NCI conducted a full system "bake-off" between several workload management and cluster management products. This bake-off took place under a demanding time schedule, so the new system could be turned into production and made available to users as quickly as possible.

The bake-off was a critical part of the decision-making process – potential vendors had to demonstrate their solutions could operate according to rigorous requirements and be customized to meet specific usage scenarios. Vendors had to perform extensive customizations to deliver the specific functionality required and integrate with other NCI systems and facility components.

Dr. Ben Evans, Associate Director at NCI, led assessment of the workload managers and the associated accounting systems. "Given the demand for the new system, we set a fairly extensive list of requirements for the new workload manager to meet; including performance, functionality, robustness as well as speed to implement the software on the system," explained Evans.

The Solution: Altair's PBS Professional for Flexible, Scalable Workload Management

NCI engaged several vendors to work onsite with the NCI team to deploy trial solutions before selecting Altair's PBS Professional® for their workload management system. "PBS Professional beat the competition in both performance and flexibility. Altair provided a superior level of functionality, and with their conscientious team of experts proved they had the capability to meet our requirements better than any other vendor in the trial," said Evans.

"NCI's users have grown to expect a lot from our computing environment: high performance, support for diverse systems management functionality, and a seamless upgrade to new systems," added Evans. "In PBS Professional, Altair delivered a platform that exceeded our previous system, and by incorporating our requirements, PBS Professional provided better scope for improvement and growth than the other scheduling software available to us."

In addition to fulfilling NCI's first stage of workload functionality and performance benchmarks, Altair was required to provide a replacement grant management and accounting system to flexibly integrate with both the Raijin system and other NCI resources – tracking everything from scheduling decisions to overall resource management for the facility (including CPU time, local scratch and global storage resources and interconnecting cloud systems). The system would replace a home-grown system that had been continuously developed and extended over many years to meet increasing and changing demands.

Altair and NCI developed this new accounting system cooperatively during the bake-off period, and the first stage of delivery provided NCI the necessary confidence that the solution could be extended to meet production requirements. "We were impressed with how quickly Altair was able to understand our needs and collaborate to provide a working initial system," says Evans. "They also demonstrated that their solution was designed to be flexible and sustainable, so it could be extended as necessary to meet future needs once the system was in production."

PBS Professional also had to integrate with the OneSIS cluster management software chosen by NCI to manage cluster nodes and other Fujitsu hardware. Once the system was deployed, a large range of jobs were loaded for PBS Professional to schedule across the entire system -- ensuring that PBS Professional could meet high performance standards and integrate with the continuous operating system and hardware maintenance regime which is common to such large-scale HPC systems. Altair Professional Services were engaged to write and integrate most of

Project Summary

Australia's National Computational Infrastructure (NCI) chose PBS Professional to manage workload for its 'Raijin' cluster from Fujitsu, the largest supercomputing system in Australia with over 1000 users. Through Altair's expert services and development support, NCI and Altair collaborated to design and implement a highly complex system capable of meeting NCI's performance and growth requirements.

Above: *The Raijin cluster at NCI, a high-performance distributed-memory cluster with peak performance of approximately 1.2 PFlops*

the basic replacement functionality, which was then tested under simulated load, systems management and component failure conditions to ensure the system would be viable for production.

NCI also engaged Altair Professional Services to perform several additional tasks, including: writing an extensive series of plug-in to further integrate NCI's system requirements; working collaboratively with NCI local system developers and administrators to port functionality to the new plugin system; writing utilities expected by NCI system managers and users to effectively manage and track the system; and expanding the functionality of the granting and allocation system to be suitable for initial users. "The plugin architecture was simple to use and allowed us very effectively to customize the system for our unique environment," said Evans.

Altair and NCI collaborated throughout the process to reconsider design decisions of the core PBS Professional software and NCI's long-held methods for managing systems in the previous in-house software. "Altair provided invaluable expertise to work through this complex and dynamic process," said Evans. "They approached the effort as a partnership, which elevated both the quality of our solution and the working experience."

Once the system was made generally available, Altair delivered on-site training for basic PBS functionality, advanced features, customized environment delivered under the contract, and the framework for plugins. The training provided an important transition point to allow NCI system developers to take ownership for ongoing development. During this phase, NCI also ported its in-house system for managing software licenses and providing key functionality from the previous system. This final step allowed NCI to migrate all users from the previous system.

The Result: Completion of a Highly Customized, Massive-Scale System

PBS Professional is in production on Raijin's 57,472 Intel Sandy Bridge cores, connected by Mellanox FDR and a 9 PByte Lustre filesystem for scratch space. The system is regularly accessed by over 1000 users, with applications that comprise a very broad range of scientific application areas and packages, including both open source and licensed products.

PBS Professional manages workload for these applications and, via the plugin extension system, provides functionality such as local and distributed job scratch spaces, software licenses and resource placement. In addition, the new PBS Professional accounting system is being used to manage computational, storage and cloud resources across the entire NCI facility. "Working closely with Altair's team was essential -- from Altair's CTO to system developers, professional services and level 2 support, the entire PBS Professional team was actively engaged to isolate, resolve and manage issues across several time zones to ensure decision making processes were streamlined," said Evans.

PBS Professional has also been installed on NCI's OpenStack Cloud system, to be available in the future for a broader range of use cases that may not fit the standard time-shared/centrally managed system model for the current clusters.

"Thanks to PBS Professional and our partnership with Altair, we have a solution we can be confident in to meet our needs -- both now and in the future."

"PBS Professional is the right choice for any production systems with complex and changing requirements, from departmental to large-scale computing systems. The product's flexibility, and the company's expertise and support, is simply unparalleled."

Naoya Tamura, General Manager, Global HPC
Fujitsu

Visit the PBS Works library of

Success Stories

at www.pbsworks.com

About Altair

Altair empowers client innovation and decision-making through technology that optimizes the analysis, management and visualization of business and engineering information. Privately held with more than 1,800 employees, Altair has offices throughout North America, South America, Europe and Asia/Pacific. With a 27-year-plus track record for high-end software and consulting services for engineering, computing and enterprise analytics, Altair consistently delivers a competitive advantage to customers in a broad range of industries. Altair has more than 3,000 corporate clients representing the automotive, aerospace, government and defense, and consumer products verticals. Altair also has a growing client presence in the electronics, architecture engineering and construction, and energy markets.

About PBS Works

PBS Works™, Altair's suite of on-demand cloud computing technologies, allows enterprises to maximize ROI on existing infrastructure assets. PBS Works is the most widely implemented software environment for managing grid, cloud, and cluster computing resources worldwide. The suite's flagship product, PBS Professional®, allows enterprises to easily share distributed computing resources across geographic boundaries. With additional tools for portal-based submission, analytics, and data management, the PBS Works suite is a comprehensive solution for optimizing HPC environments. Leveraging a revolutionary "pay-for-use" unit-based business model, PBS Works delivers increased value and flexibility over conventional software-licensing models.

www.pbsworks.com

Altair

Altair Engineering, Inc., World Headquarters: 1820 E. Big Beaver Rd., Troy, MI 48083-2031 USA
Phone: +1.248.614.2400 • Fax: +1.248.614.2411 • www.altair.com • info@altair.com