



Altair

PBS Works™

Enabling On-Demand Computing™



Improving the Accuracy of Weather Prediction

A Cray-Altair Solution for Optimized HPC Resource and Workload Management

Using a Cray supercomputer with Altair’s PBS Professional® scheduler provides weather prediction centers with a **rich solution for optimized forecasting**. The solution is designed to support and leverage the specific features of the Cray architecture, so high-performance computing (HPC) resources are more efficiently scheduled to **increase system utilization, improve application performance and throughput, and enhance quality of service**.

Solution at a Glance

Cray Supercomputer with Application Level Placement Scheduler (ALPS)

- ALPS provides a mechanism to place and launch the applications contained within batch jobs

Altair PBS Professional

- Commercial-grade job scheduling and management
- Policy enforcement
- Standing and advance reservation
- Backfilling
- Pre-emptive scheduling

Challenge

The field of operational numerical weather prediction (NWP) is relied upon nearly as much as oil exploration, in terms of impact to the global economy and society.¹ The numerical forecasts delivered by NWP centers provide guidance to human forecasters and are the basis of all weather forecasts by national meteorological services and the media. In the U.S. alone, funding for systems that support more accurate NWP is in the billions.²

Advances in this area in recent decades have resulted in far more accurate predictions. However, as many NWP centers approach petascale

performance levels – and processor core counts, main memory sizes and I/O requirements continue to grow – significant operational challenges emerge:

- Greater number of concurrent high priority jobs, resulting in an increased likelihood of resource conflicts
- Complexity of mixed operational and research workload
- Unpredictability of emergency or other high-priority jobs

The result is unreliable runtimes, sluggish or unpredictable application performance, cost implications around

Improving the Accuracy of Weather Prediction

“With Cray and Altair we were able to improve the performance and manageability of our numerical weather prediction environment. Their solution, allowing for local customization, allowed us to implement our own custom toolbox for handling production jobs without the need for invasive resource preemption.”

Danish Meteorological Institute

memory requirements, and risk in meeting forecast schedules and maintaining service levels.

Thus, an evolution in system utilization approaches along with workload scheduling and management technologies is needed to exploit the breakthroughs in NWP peak performance capabilities while improving service quality and utilization rates.

Solution

Using a Cray supercomputer in conjunction with PBS Professional from Altair represents a breakthrough in system utilization strategy and high-performance workload scheduling and management. **This solution enables computing facilities to exploit peak performance capabilities while improving service quality and utilization rates.**

Cray supercomputers deliver petascale production and extreme scale multi-core processing through hardware and

software technologies such as the Cray custom-designed interconnect.

Cray systems such as the Cray XK6 and Cray XE6 massively parallel supercomputers are purpose-built for HPC applications, delivering high-performance, balanced capability computing. Only systems designed for capability oriented HPC allow scientists and forecasters to produce results in the shortest amount of time possible while investigating increasingly complex phenomena.

Cray systems incorporate the Application-Level Placement Scheduler (ALPS), a rich scheduling environment that provides application placement, launch and management functionality for all applications. This environment works cooperatively with PBS Professional which makes scheduling decisions regarding when and where to run jobs and enforces policy. PBS communicates with ALPS which then sets up the environment and executes the job run.

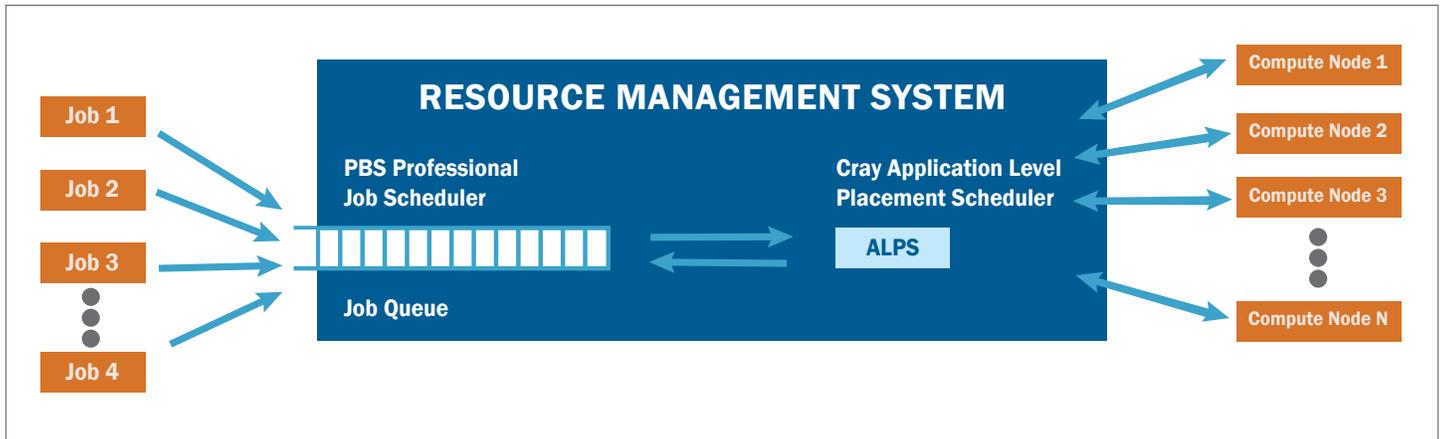
This tight integration ensures extremely high system utilization levels and increases throughput while providing quick turnaround of high-priority jobs.

The joint Cray-Altair solution guarantees the availability of resources reserved in advance, optimizes the usage of non-reserved resources, and improves scheduling predictability so users have greater insight into when jobs will run. The system is highly configurable, permitting sites to easily implement custom scheduling behaviors.

Features

The Cray-Altair solution offers the following key features:

- **Integrated application placement, launch and management functionality** provided for all applications whether interactive or batch
- **Minimal to no system intervention** required for job execution
- **Unparalleled scalability** from compute nodes to software to I/O



The Cray-Altair Solution for Optimized Resource Management

- **Increased reliability at scale**, thanks to redundant power supplies and voltage conversion modules
- **Gemini interconnect** for highly scalable performance of communication-intensive algorithms
- **Standing and advance reservation features** to proactively provide scheduler with data about future high priority resource requirements – so the scheduler can guarantee QoS for this workload
- **Backfilling** to optimize space sharing by examining available resources against outstanding job requests, so operational and research workloads can be managed together
- **Pre-emptive scheduling**, so unplanned critical jobs can be launched by allowing the scheduler to free up resources used by other jobs
- **Job array** providing a mechanism for grouping related work, so users can submit, query, modify and display work as a single unit

Benefits

High-performance computing facilities enjoy a range of benefits from using a Cray-Altair solution, including:

Improved QoS for operations

- Greater availability of resources
- Reliable and repeatable runtimes
- Improved ability to forecast when a job will run
- Reduced execution times to meet production schedules
- Fast turnaround of high-priority jobs
- Support for long-range scheduling needs

Optimized performance

- Optimized application throughput and overall system efficiency
- Optimized resource availability
- Avoidance of system thrashing

Improved manageability

- Reduced need for process-level management
- Configurability to easily implement custom scheduling behaviors

User Spotlight: DMI

Researchers at the Danish Meteorological Institute (DMI) have been using the Cray-Altair solution to support operational NWP for over three years. The highly customizable solution has allowed DMI to develop a framework that extends the functionality of the system for their unique environment and behavioral policies. Operating mostly unattended, this solution has successfully fulfilled DMI's operational scheduling requirements.

This implementation is described in the technical paper "Producing Weather Forecasts on Time in Denmark Using PBS Professional" by Thomas Lorenzen et al.³

Download the technical paper:

This solution is described in more detail in the Cray-Altair white paper "Operational Numerical Weather Prediction Job Scheduling at the Petascale" including:

- Solution strategy details
- Architecture deep dive
- Multiple code demos for scheduling use cases

Download at: pbsworks.com/resources

Visit the PBS Works library of
Technical Papers
at www.pbsworks.com/resources



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



Global supercomputing leader—Cray Inc. provides innovative systems that enable scientists and engineers in government, industry and academia to meet existing and future computational challenges. Building on expertise in developing, marketing and servicing the world's most advanced supercomputers, Cray offers a comprehensive portfolio of high performance computing (HPC) systems delivering unrivaled sustained performance on a wide range of challenging applications.

www.cray.com

¹ NOAA Celebrates 200 Years: www.celebrating200years.noaa.gov/foundations/numerical_wx_pred/welcome.html

² National Oceanic and Atmospheric Administration: www.noaa.gov/newsroom/stories/2012/pdfs/FY13_Briefing_021612_final.pdf

³ Producing Weather Forecasts on Time in Denmark using PBS Professional.™ Thomas Lorenzen (Danish Meteorological Institute), Thor Olason (Danish Meteorological Institute), Frithjov Iversen (Cray Inc.), Paolo Palazzi (Cray Inc.). www.cug.org/1-conferences/CUG2011/index.php (select Final Program to access .pdf – requires login)



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