



# PBS WORKS 2018 AND BEYOND: MULTI-SCHED, POWER RATE RAMP LIMITING, SOFT WALLTIME, JOB EQUIVALENCE CLASSES, ..., NATIVE MODE

Scott Suchyta | Director, Solutions & Integrations | May 20-24, 2018



Altair

# CRAY AND ALTAIR PARTNERSHIP

Long history of partnering on HPC and CAE

Signed OEM agreement in April 2015 which allows Cray to cost effectively bundle PBS Professional on all Cray systems

Renewed / extended OEM agreement in April 2018 to cover another 3 years

Awarded Cray's supplier of the year award at SC17 (first time given to a software vendor)

## Altair and Cray Extend Their Collaboration in High-Performance Computing PBS Professional® workload manager to continue maximizing utilization on Cray systems

TROY, Mich. – May 22, 2018 – [Altair](#) (Nasdaq: ALTR) Altair announces a 3-year extension to its original equipment manufacturer (OEM) agreement with Cray Inc. Under the terms of this agreement, PBS Professional will remain bundled as the preferred commercial scheduler on new systems manufactured and shipped by Cray.

“Our relationship with Altair continues to benefit our common customers, as we deliver unique, integrated technologies that are ready to meet their highest HPC requirements,” said Fred Kohout, Senior Vice President of Products and Chief Marketing Officer at Cray. “Altair and PBS Professional are established leaders in HPC workload management, and we will continue to leverage this agreement for growth in commercial and other emerging market segments.”



# PBS WORKS



Technology to manage HPC clusters,  
clouds and supercomputing environments

Targets IT administrators and HPC users

Direct sales and support  
Also sold by computer hardware  
companies

Workload Management

Cloud Computing

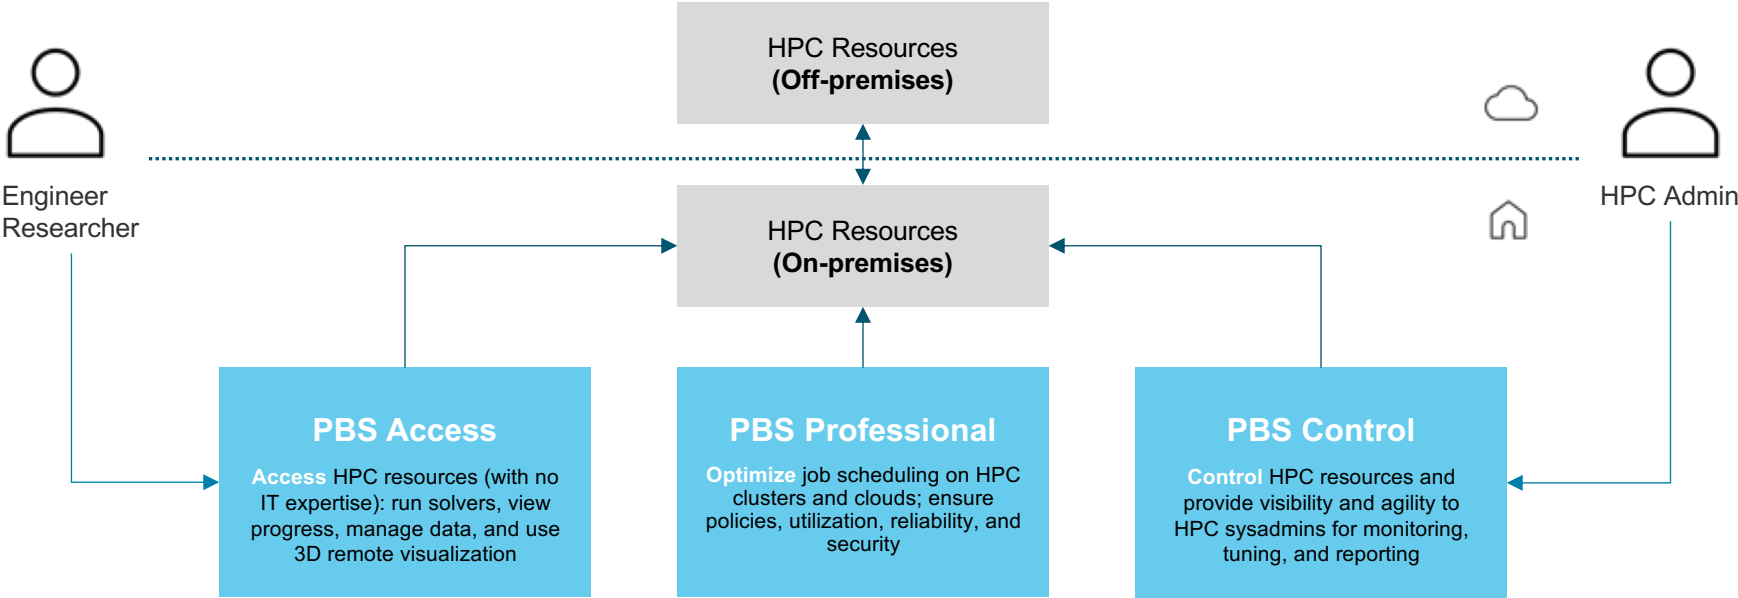
Remote Visualization

IT Analytics

Cross-domain Security



# PBS WORKS SUITE



# PBS ACCESS

A PORTAL FOR ENGINEERS AND RESEARCHERS

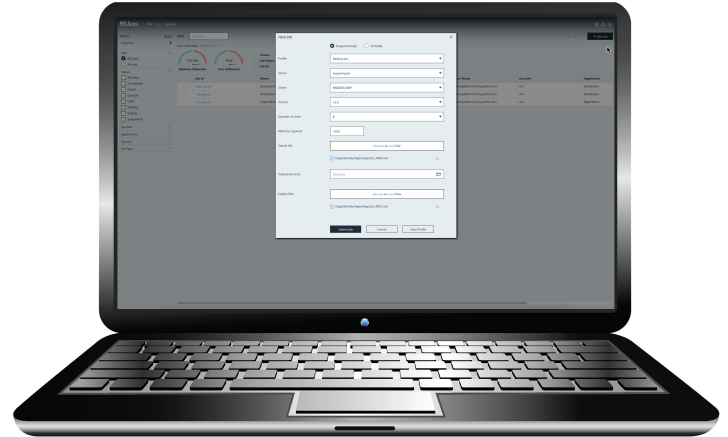


# PBS ACCESS FEATURES

Altair's new PBS Access provides Engineers and Researchers a more simple, powerful and consistent interface for submitting and monitoring jobs on remote clusters, clouds or other resources.

## Highlights

- **Novice to Expert:** simple and powerful
- **Same UX:** desktop and web
- **Secure:** protected access to HPC resources
- **End-to-end:** submit, monitor progress, steer, fix, and rerun jobs
- **Collaborate:** shared 3D analysis
- **3D Remote Visualization**
- **Save time:** Simplify job submission and management thanks to a powerful GUI with smart, simplified interfaces
- **Be more productive:** Spend more time focused on work and not IT tasks
- **Increase ROI:** Consolidate access to applications and optimize license availability
- **Reduce errors and improve consistency:** Embed your company's best-practice "know how" directly into Application Definitions used for job submission



# PBS ACCESS CAPABILITIES

Engineers, scientists, and researchers, can easily access portals that provide access to HPC (no IT expertise needed) to run solvers, view progress, manage data, and use 3D remote visualization via web and desktop.



## Files

Simplify job **file** submission and management thanks to a powerful GUI with smart, simplified interfaces.



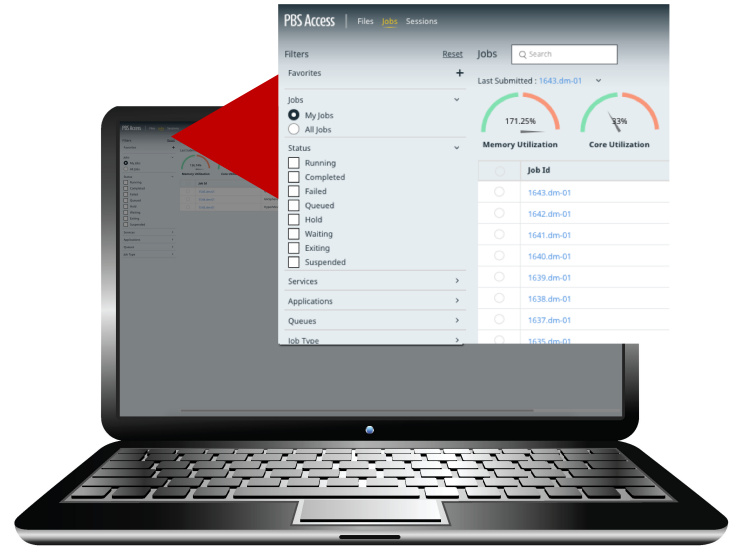
## Jobs

Easily define and configure applications for **job** submission via PBS Application Services (PAS).



## Sessions

Manage 3D remote visualization **sessions** and data directly from the secure web portal.



# PBS CONTROL

HPC ADMINISTRATORS CONTROL CENTER FOR MANAGING,  
OPTIMIZING AND FORECASTING HPC RESOURCES



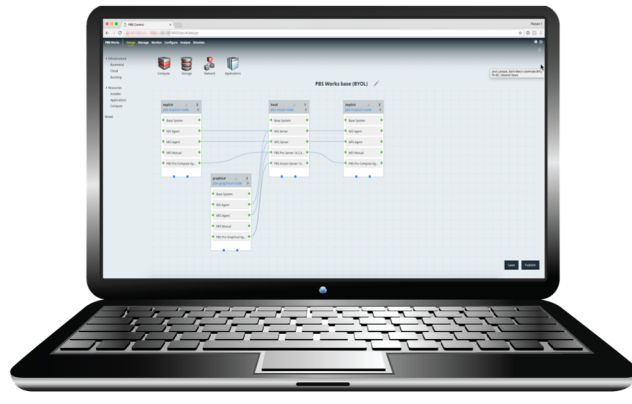


# PBS CONTROL FEATURES

Altair's new PBS Control is an easy-to-use web application for monitoring and managing jobs and nodes of a High-Performance Computing (HPC) cluster with advanced analytics to support data-driven planning and decision making.

## Highlights

- **Single pane of glass:** configure, deploy, monitor, burst, manage, troubleshoot, simulate, analyze and tune
- **Real-time monitoring:** simplify troubleshooting and maintenance
- **Reporting:** PBS Analytics powered by Carriots Analytics™
- **Workload simulator:** simulate and optimize infrastructure sizing
- **Multi-cloud bursting:** burst to any cloud for peak loads
- **Modern UX:** drag-and-drop simplicity



# PBS CONTROL CAPABILITIES

A platform to easily manage, monitor, analyze, and simulate HPC clusters



## Design

**Design** an HPC model so that an HPC cluster can use Cloud bursting to public and private clouds adding capacity exactly when needed.



## Manage

Deploy an HPC model on the cloud. **Manage** Cloud bursting by updating HPC details.



## Monitor

**Monitor** and manage an HPC cluster's nodes and jobs. Monitoring is divided into sub-tabs for easy viewing.



## Configure

**Configure** default WLM server and security settings and manage queues, resource defaults and limits.



## Analyze

Easy to use job accounting and reporting solution with advanced **analytics** to support data-driven planning and decision making.



## Simulate

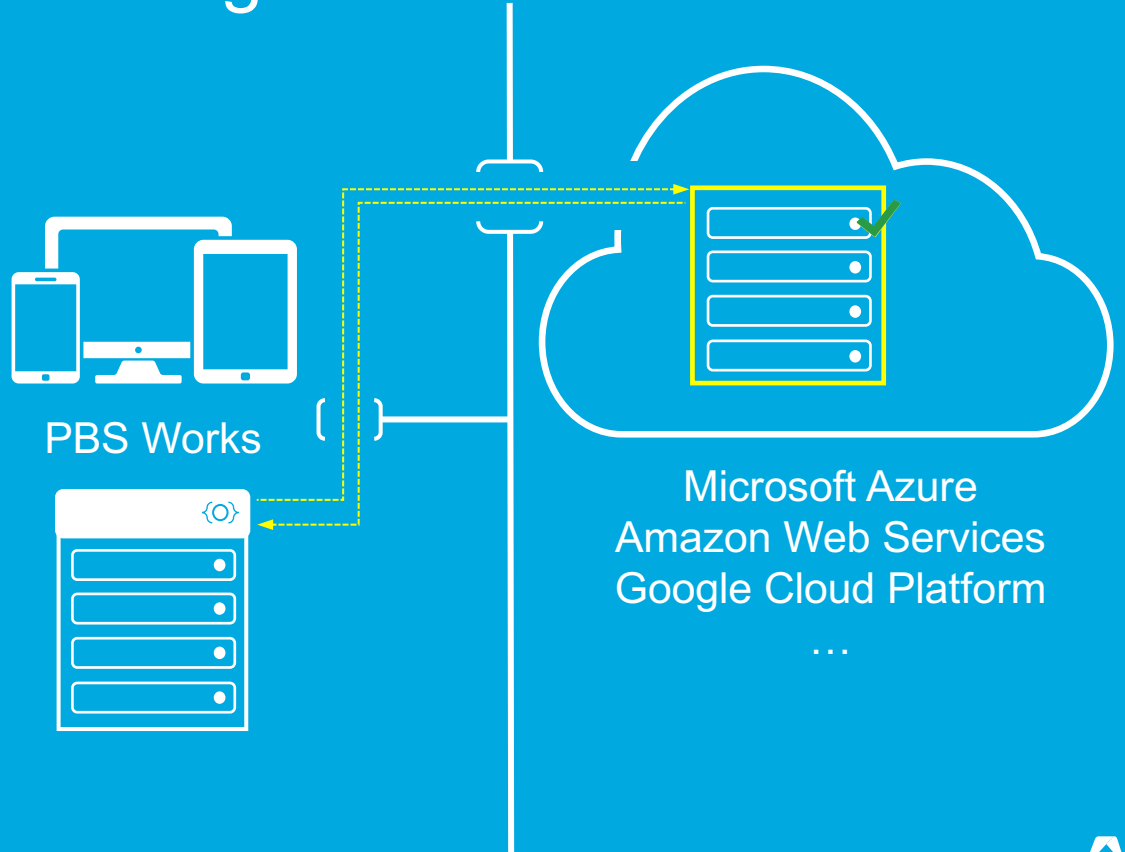
Run a **simulation** to perform what-if analysis to determine the most productive way to scale an HPC's resources.



# PBS Control: Cloud Bursting

On-demand use of cloud resources to maximize efficiency

- Improve responsiveness, adding capacity exactly when needed
- Automatic governance and cost controls via site-defined policy and quotas
- Understands on-premise utilization, ensuring bursting only when cost-efficient
- Vendor-agnostic: no lock-in
- Fast: +1,000 nodes in minutes



# PBS PROFESSIONAL

PBS PROFESSIONAL® IS ALTAIR'S INDUSTRY-LEADING WORKLOAD  
MANAGER AND JOB SCHEDULER  
FOR HIGH-PERFORMANCE COMPUTING (HPC) ENVIRONMENTS.



# PBS PROFESSIONAL FEATURES

PBS Professional automates job scheduling, management, monitoring and reporting. It is the trusted solution, not just for complex systems, but also for smaller clusters.

## Highlights

- **Ubiquity:** Open Source and commercial licensing
- **Performant:** million core scalability, fast throughput
- **Secure:** EAL3+ security certification and MLS (SELinux) support
- **Intelligent:** policy-driven and topology-aware scheduling
- **Scalable:** proven to run millions of jobs per day
- **Green Provisioning™:** power management and control
- **Robust:** known in the industry for stability and support
- **Open Architecture:** implement virtually any new policy
- **Modern:** containers, cloud bursting, parallel scheduling



NASA's workload manager of choice  
for all NAS HPC resources

250k+ cores scheduled by one PBS Pro



# PBS PROFESSIONAL CAPABILITIES

A fast, powerful workload manager designed to improve productivity, optimize utilization & efficiency, and simplify administration for HPC clusters, clouds and supercomputers.



## Policy

Meet enterprise-wide scheduling and use policies with resource-based scheduling to ensure the right jobs run at the right times on the right resources



## Utilization

Maximize ROI and minimize waste, ensuring the best use of hardware, software licenses, energy, and clouds resources



## Performance

Fast and scalable execution to meet the needs of even the most demanding enterprises with thousands of users and millions of cores



## Resilience

Secure and hardened to minimize downtime: health checks and automatic failover ensures jobs get done despite faults and failures



## Open

Extensible plugin framework to meet each site's unique needs. Dual-licensed: commercial & Open Source

# AGGRESSIVELY OPEN & COMMUNITY ORIENTED

- Opening the full core of PBS Professional
  - not just a weak subset or older version
- Reorganizing to behave as one of the (hopefully many) contributors
- Using community-accepted practices
  - OSI-approved license
  - GitHub, JIRA, open roadmaps, ...
- Focusing on longevity
  - For a viable, sustainable community



[www.pbspro.org](http://www.pbspro.org)



## PBS PROFESSIONAL – 18.X

- **Docker:** automatically run jobs in per-job specified containers, with added security that mitigates the Docker “as-root issue”
- **MultiSched:** multiple schedulers each handle separate partitions, with per-partition scheduling policies, speeding scheduling while retaining a single global view of the whole cluster (first seen in PBS-Cray 13.0.40x branch!)
- **Cgroups:** provide kernel-level resource isolation, usage limits, and reporting, so jobs sharing nodes run as fast as possible without interfering with each other
- **Job equivalence:** job class optimization improves scheduling speed by an order of magnitude for workloads with many thousands of small jobs





## PBS PROFESSIONAL – 18.X... THERE'S MORE!

- **Soft walltime resource:** improve backfill calculations and increase utilization
- **Jobs can shrink:** release no-longer-used nodes, e.g., during post-processing or stage-out, effectively increasing utilization
- **High Throughput:** 10k jobs per minute!!
- **New platform support:** ARM64
- **Cray 13.0.40x Features Merged:** all Cray PBS-ALPS features have been ported to master and available in open source
  - Intel Xeon Phi – provision hbmem and NUMA config
  - Support IMPS (CLE 6.0)
  - Preemption w/ Suspend/Resume
  - Power-awareness – on/off nodes, power profiles, energy accounting, and power ramp rate limiting
  - Alter PBS [standing|advanced] reservations



## PBS PROFESSIONAL – WHAT'S NEXT?

- **Cray's Next Generation "Shasta"**: even stronger relationship with Cray
- **Supporting Cray's "Native Mode"**: preparing for the next generation
- **PBSPro.org**: The community website!
- **PBSPro.atlassian.net**: Developer/Contributor Portal includes roadmaps, processes, how to articles, coding standards, release notes, etc.



# THANK YOU!

Technical Session 20C (Wednesday 4:00)

PBS Professional - Optimizing the "When & Where" of Scheduling  
Cray DataWarp Jobs



THANK YOU!

