



ALTAIR® LIQUID SCHEDULING™ – FAST, FLEXIBLE WORKLOAD MANAGEMENT

Liquid Scheduling takes Altair® PBS Professional® to the next level. It's faster, more scalable, and more resilient than anything previously seen in high-performance computing (HPC). Liquid Scheduling catapults the Altair® HPCWorks® platform into a bigger, more connected world of HPC, artificial intelligence (AI), and cloud. It can handle larger workloads and more compute than ever before by breaking down silos and uniting multiple HPC clusters, with improved user experience and system utilization, and without changing users' applications or workflows.

The Future of Hybrid HPC and AI

HPC is rapidly becoming more complex, and emerging AI workloads are changing the face of computing. Data analytics is prevalent, machine learning is often used in place of or in addition to solvers and traditional algorithms, users are deploying large language models (LLMs), and AI is training and testing self-driving cars. AI jobs are often reliant on containers and GPUs, as well as high-throughput computing with lots of short, bursty jobs — bursts of activity often followed by long dormant periods — with tight deadlines.

Today's administrators need to support a wide range of new workloads that mix AI and machine learning with HPC while pulling data from various sources, and compute environments that consist of assorted structures including GPUs, CPUs, and new accelerators. PBS Professional is widely used as a platform for AI. It delivers full-spectrum scheduling for both high-throughput and HPC workloads, as well as cloud bursting and cluster scaling.

Liquid Scheduling

Liquid Scheduling — which works with PBS Professional — was designed to meet the challenges of the changing world of HPC and AI. It delivers scalability for next-generation supercomputers at exascale and beyond, with flexibility to meet the demands of today's diverse, distributed workloads. It's a leap in performance and scalability that improves utilization and access to resources, breaking down silos by connecting multiple HPC clusters. Liquid Scheduling uses a distributed streaming database that enables multiple scheduling domains to manage resources and prioritize workloads based on available resources and capacity.



Altair Liquid Scheduling is the biggest change to HPC in 30 years. The best part is that your users don't need to change a thing about their applications or workflows; they simply flow to the right place with Liquid Scheduling.

Dr. Rosemary Francis,
Chief Scientist HPC, Altair



A New Level of Scalable Scheduling

Liquid Scheduling distributes workloads automatically across multiple PBS Professional clusters, improving performance and scalability even on a single cluster. Jobs are moved from Liquid Scheduling to the workload manager (WLM) that will run them using our patented just-in-time scheduling mechanism. This means that with Liquid Scheduling, WLMs will only address jobs they can handle based on resource requirements and capacity, queuing them only when they're ready to run. Individual WLMs don't need to manage the entire queue. This is one of the main reasons Liquid Scheduling is highly scalable, and it's very different from past approaches of pre-queueing jobs and moving them around between clusters based on capacity. Its global policy engine makes fairshare decisions across an entire distributed compute environment to ensure the right workloads run in the right place at the right time.

Liquid Scheduling is built for future scalability using modern, web-scale technologies and streaming architectures. PBS Professional users can use Liquid Scheduling via familiar CLI and GraphQL interfaces, and with no change to existing workflows or applications. They're freed from having to choose where to run their jobs. Liquid Scheduling is so fast because it connects distributed resources and separates scheduling policies from resource management. System monitoring and accounting is unified. Liquid Scheduling's architecture makes resources more elastic by supporting growing and shrinking scheduler domains between platforms, on-premises and in the cloud.

Liquid Scheduling Users

The National Supercomputing Centre (NSCC) Singapore maintains multiple PBS Professional complexes, each designed for a specific type of workload, some optimized for traditional HPC and others for machine learning and AI. NSCC collaborates with stakeholder institutions in Singapore and with partners worldwide, including the National Computational Infrastructure (NCI) in Australia. [NSCC deployed Liquid Scheduling](#) to enable users to run workloads using a global pool of resources. Liquid Scheduling extends NSCC's existing Altair HPC stack, giving users fast, scalable access to a wider compute pool without changing their workloads.

PERFORMANCE SPOTLIGHT

100M

JOBS RUNNING ON

14

PBS CLUSTERS
= MORE THAN

250K+

JOBS PER SECOND

Liquid Scheduling catapults the HPCWorks platform into a bigger, more connected world of HPC, AI, and cloud.