

Altair Accelerator™ is the industry's fastest enterprise job scheduler, with high-performance hierarchical scheduling and scalable, event-driven architecture for the highest throughput. Accelerator maximizes capacity utilization and is capable of running millions of jobs per day.

Accelerator is an enterprise-grade job scheduler designed for distributed high-performance computing (HPC) environments. It's a highly adaptable solution capable of managing compute infrastructures from small, dedicated server farms to complex, distributed HPC environments.

The Plus package, Altair Accelerator Plus™, includes a hierarchical scheduler architected to offload the base scheduler, enabling greater throughput, better license and resource utilization, and more flexible scheduler usage models.

Maximizing Throughput

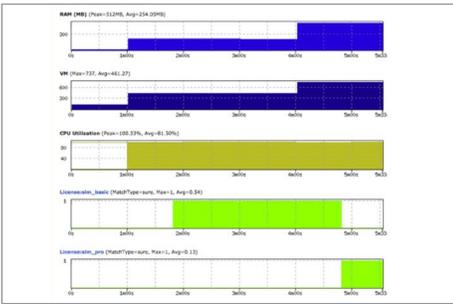
Accelerator is a full-featured, high-performance job scheduler designed to manage complex workload environments. Its scalable, event-driven architecture combined with a small memory footprint and sub-second dispatch latency accelerates design throughput and maximizes capacity utilization – and it's capable of running billions of jobs per day.

Equipped with a comprehensive set of policy management features including fairshare and preemption, Accelerator provides the essential controllability and observability tools to set and alter policies, optimizing the workload environment as needed. Accelerator enables the implementation of fairshare policies for both hardware resources and software licenses to ensure that business-critical work is completed on time.

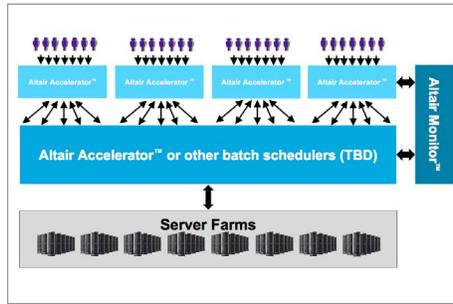
Job preemption and reservations can be used to allow urgent jobs to take precedence over less important jobs. Accelerator can suspend running jobs and resume them once higher-priority jobs have been completed, eliminating the need to set aside licenses for urgent jobs. Accelerator can achieve 100% license utilization while servicing all classes of jobs.

Accelerator Plus includes a hierarchical scheduler and millisecond dispatch latency, which yields the most benefits for short jobs. It serves as a “personal scheduler” director, allowing users to queue short jobs sequentially on their own. This reduces the scheduling burden on the lower-level scheduler by presenting a batch of short jobs as one longer job while maintaining full visibility into each individual job – a common need for users. It can also handle all the user query, job submission, and reporting functions. The offload from the

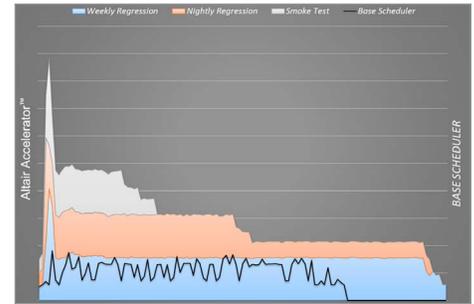
Learn more:
pbsworks.com



Visually monitor server farms



Scale out scheduler loads with Accelerator Plus



Coelesce jobs with Accelerator Plus

base scheduler is substantial, since 80% of typical scheduler loads come from these non-dispatch functions.

For highly parallel, short jobs with consistent license and computing resource needs, Accelerator Plus deployments have realized a 6-10X increase in throughput with no added compute capacity.

Greater Productivity

With Accelerator’s visualization and failure identification capabilities, system administrators get a big-picture view of jobs that have completed, failed, and are waiting – so they can analyze relevant report files to quickly identify the root cause of a failing job.

Accelerator can be easily configured to send automatic alerts to users or administrators when situations occur.

Enhanced Visibility

Accelerator Plus is equipped with a comprehensive set of policy management features including fairshare and preemption. Cross-organizational resources can be assigned and managed as users prefer – by workflow, by project, or on demand. A full-featured GUI enables full visibility into each job and compute farm statistics for queued jobs, running jobs, and allocated resources. User interfaces include web browsers, command line, and an extension API.

For example, in a semiconductor verification application, under a typical high-load situation of two regressions and a smoke test, Accelerator Plus manages the peak loads with ease. When combining all three workloads, the effective load on the base scheduler stays stable at a level that is less than any of the individual workloads.

Product Highlights

Designers

- Sub-second latency for best performance and user experience
- Full visibility into key job resources (CPU, memory, and license)

Added benefits with Plus package:

- Exclusive access to private scheduler; allows self-directed flows
- High throughput for highly parallel workloads of uniform short jobs

Infrastructure

- Scalable, small-footprint architecture for best resource utilization
- Customizable alerts and notifications
- Fast restart without impacting running jobs

Added benefits with Plus package:

- Flexible usage models (on-demand, workload, project)
- 6-10X higher throughput without adding compute capacity

Management

- High-capacity, scalable architecture capable of scheduling tens of millions of jobs per day
- Manage organizational priorities with comprehensive policy management features

Added benefits with Plus package:

- Workload isolation to support difficult workloads without impacting other users
- Scale-out mechanism for job dispatch