

ALTAIR® GRID ENGINE® FOR ENTERPRISE-CLASS SCHEDULING & OPTIMIZATION

Altair® Grid Engine® manages workloads automatically, maximizes shared resources, and accelerates the execution of containers, applications, and services. The solution can be deployed in any technology environment: on-premises, cloud, hybrid cloud, or cloud-native high-performance computing (HPC). By using Altair Grid Engine, enterprises and organizations can deliver products and results faster, more efficiently, and with lower overall costs.

Optimizing Distributed Resources

Altair Grid Engine efficiently shares workloads across machines to optimize the use of computing infrastructure. Scheduling policies can be applied to all work submitted to a cluster, ensuring high-priority jobs are completed on time while maintaining maximum utilization of all cluster machines. Coupled with Altair Navops Launch™, Altair Grid Engine can dynamically schedule hybrid cloud workloads. Host-based resources can be selected at a highly granular level, so GPUs can be mapped to Docker containers and then used by GPU-enabled applications that run on any Docker-enabled host in the cluster.

Priority and Utilization Policies – Altair Grid Engine delivers multiple scheduling policies for matching workloads in the cluster to business objectives such as maximizing utilization across all machines, reducing job turnaround time, and prioritizing workloads per group or department.

Scalability – It can scale to a cluster of 350,000 cores in a single, managed environment. A single Altair Grid Engine cluster can contain more than 10,000 nodes and run 225 million jobs per month.

Resource Management – The solution continuously collects metrics from all cluster nodes, then uses scheduling strategies configured by the administrator to evaluate all pending workloads and match specific job requirements to available resources.

Multiple Workloads – Any type of application or accelerator workload, including NVIDIA GPUs and Docker, can run through Altair Grid Engine.

Quotas and Limits – The solution can configure flexible quotas on users, projects, and groups to control how much workload is run in the cluster and by whom, ensuring customers achieve their business SLAs.



[Altair Grid Engine is] a product that has features, capability, and performance that are not only better than the competition, but exceed what customers require.

Financial Post



Try Altair Grid Engine:
[Get a Free Trial](#)



Better Results, Faster

With Altair Grid Engine you can obtain higher-quality results faster with industry-leading throughput — plus share resources more effectively and improve resource utilization. Reduce hardware, software, and data center costs with efficient workload management that boosts user productivity with reduced wait times, higher throughput, and less downtime.

Because we believe in the broadest application, OS, and hardware support, Altair Grid Engine supports thousands of commercial and open-source applications in fields from life sciences and manufacturing to energy, machine learning, and AI. Use your choice of Linux, Windows, and other OS distributions on x86, Power, and Arm systems.

With advanced GPU support, proven cloud-ready features, and extreme scale and performance, Altair Grid Engine delivers everything you need to manage the most demanding workloads.

Key Features

Native Docker Support – Run Docker containers at scale and blend containers with other workloads supporting heterogeneous applications and technology environments.

GPU Support – Scale GPU-based frameworks from containers and servers to clusters and clouds.

Cloud-native HPC Support – Run HPC workloads and microservices on shared Kubernetes clusters.

Industry-leading Customer Support – Altair’s experts are always here to help.

Altair Grid Engine helps organizations improve ROI and deliver better results faster by optimizing throughput and performance while maximizing shared compute resources.

350K

CORES IN A SINGLE ENVIRONMENT

10K

NODES PER CLUSTER

225M

JOBS PER MONTH

Altair Grid Engine improves workload throughput, increases compute resource utilization, and accelerates time-to-results.