



# Univa Grid Engine

## Enterprise-class workload scheduling and optimization solution

### Univa® Grid Engine® Highlights

- Improve workload throughput
- Increase utilization
- Accelerate time-to-results
- Decrease management costs
- Lower total cost of ownership

#### In the news:

*"Using Univa Grid Engine cut hardware costs in half and reduced the time needed to process large data sets and perform calculations."*

— **eWeek**

*"A product that has the features, capability, and performance that are not only better than the competition, but exceeding what customers require."*

— **Financial Post**

*"Univa Grid Engine creates efficiencies that otherwise wouldn't be possible, which shaves a great deal of time from the process."*

— **Urgent Communications**

*"Univa's efforts to incorporate containers into the Grid Engine scheduler and dispatch system have significant implications."*

— **HPC Wire**

*"Without the correct optimization and management tool in place, enterprises risk losing the benefits of containers."*

— **EnterpriseTech**

### OPTIMIZE WORKLOAD AND MAXIMIZE DISTRIBUTED DATA RESOURCES

Univa® Grid Engine® manages workloads automatically, maximizes shared resources and accelerates deployment of any container, application or service. The solution can be deployed in any technology environment, on-premise or in the cloud. By using Univa Grid Engine software, enterprises and organizations can deliver products and results faster, more efficiently, and with lower overall costs.

With Univa Grid Engine, workloads are efficiently shared across machines in a data center to optimize the computing infrastructure. Scheduling policies can be applied to all work submitted to the cluster, ensuring high-priority jobs are completed on time while simultaneously maintaining maximum utilization of all cluster machines. The solution also monitors any resource or software license and schedules applications ensuring they are automatically matched to the appropriate licenses and machines.

### KEY FEATURES AND CAPABILITIES

#### Priority and Utilization Policies

Univa Grid Engine software delivers multiple scheduling policies for matching workload in the cluster to business objectives such as maximizing utilization across all machines, reducing turnaround time for jobs in the cluster, and prioritizing workloads according to group or department.

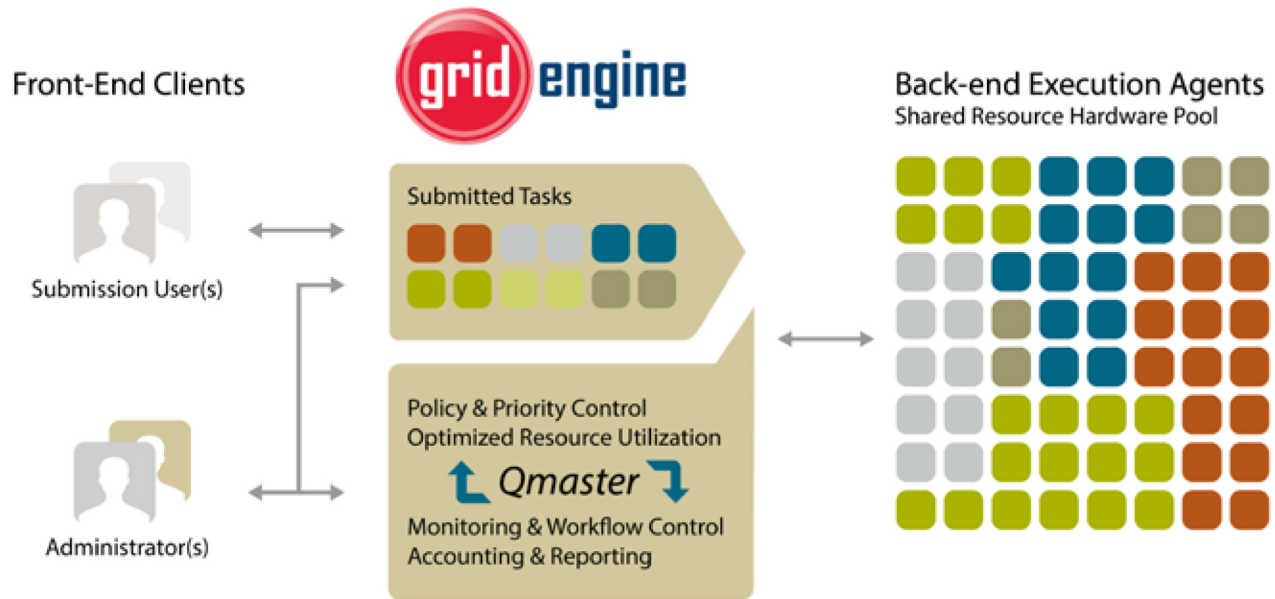
#### Scalability

Can scale to a cluster of 120,000 cores in a single managed environment. A single Grid Engine cluster can contain more than 10,000 nodes and run 100 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.





Univa has developed the industry's most powerful and integrated solution suite for dynamically shared enterprise clusters. Select from a variety of Univa Grid Engine add-ons to create a customized solution:



### Container Edition

Run Docker containers in a Univa Grid Engine cluster at scale and blend containers with other workloads supporting heterogeneous applications and technology environments.



### Short Jobs

A framework to run massive quantities of tasks in microseconds and with near-zero overhead to Univa Grid Engine.



### UniSight Reporting

An analytics tool that allows organizations to measure, track and chargeback usage on dynamic and shared clusters.



### UniCloud

Ensure infrastructures are created and configured properly — local or remote in a hybrid cloud, bare metal or virtual.



### License Orchestrator

A manager for the allocation of licensed applications and application features shared across Univa Grid Engine clusters.



### Universal Resource Broker

Create a dynamically shared compute pool out of distributed resources by hosting popular Big Data frameworks and data center services on top of clusters.



### Web Services

An API that maximizes developer efficiency while making it simple to on-board applications into a cluster.