UNIVA

DATASHEET

Univa Grid Engine

Enterprise-class workload scheduling and optimization solution

OPTIMIZE WORKLOAD AND MAXIMIZE DISTRIBUTED DATA RESOURCES

Univa® Grid Engine® manages workloads automatically, maximizes shared resources and accelerates the execution of any container, application or service. The solution can be deployed in any technology environment: on-premise, cloud, hybrid cloud or Cloud-native HPC. By using Univa Grid Engine, 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 use of 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. Coupled with Navops Launch, Grid Engine can dynamically schedule workload to a hybrid cloud while Navops Command enables HPC workloads and microservices to run on a shared Kubernetes cluster. The solution also monitors any resource or software license and schedules applications ensuring they are automatically matched to the appropriate licenses and machines.

Host-based resources can be selected at a highly granular level. This includes support for special resources such as GPUs and complex server topologies with multiple CPU sockets, multiple GPUs and multiple network interfaces for which optimal resource allocation is crucial to deliver best workload performance and resource utilization. Univa Grid Engine's container support encompasses the usage of GPU-enabled servers from within Docker or Singularity containers.

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
Univa has proven extreme scalability with Univa Grid Engine and has scaled to 200,000 cores in a single, on-premise environment and to 1 million cores in the cloud. A single Grid Engine cluster can contain more than 10,000 nodes and run more than 200 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, like Docker, Intel KNL, NVIDIA GPUs, can run through Univa Grid Engine.

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

Univa® Grid Engine® Highlights

- Scalability in Cloud to 1M cores
- Improve workload throughput
- Increase utilization
- Accelerate time-to-results
- Cloud-ready and Cloud proven
- Lower total cost of ownership
- Enable large scale Machine Learning

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

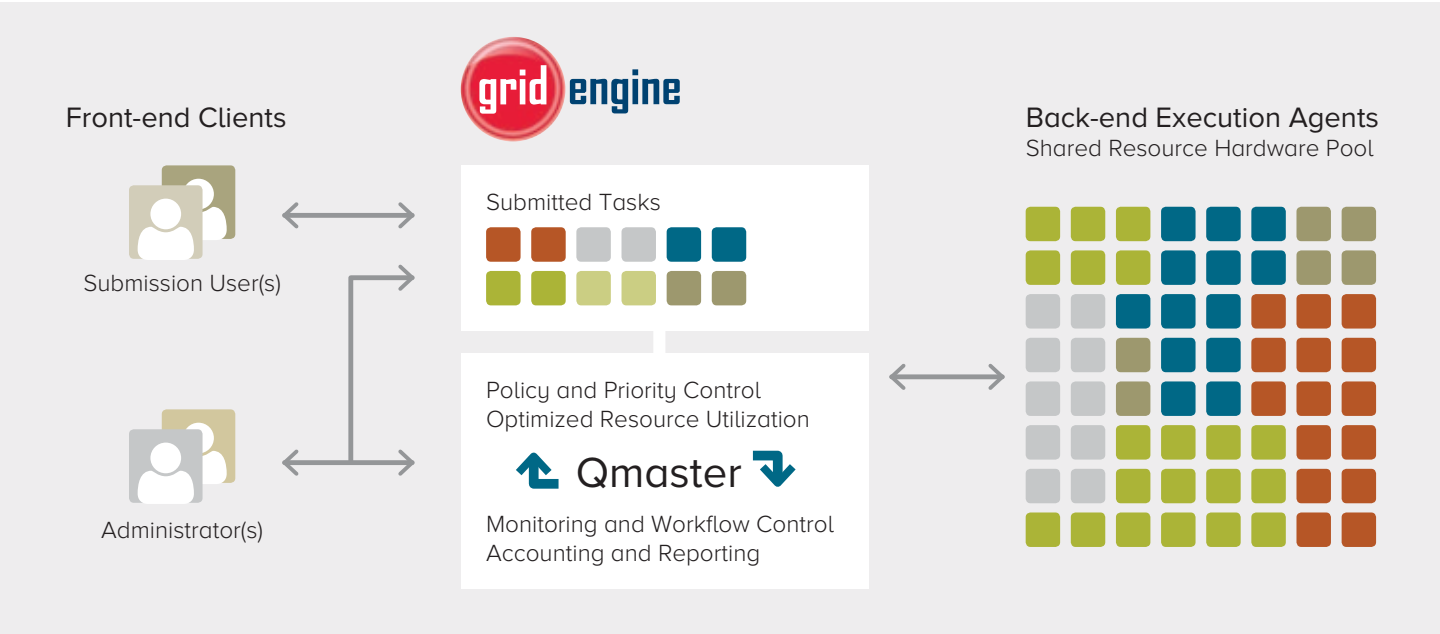

Univa Grid Engine Features:

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

Monitoring and Reporting
Track and measure resource utilization in workload managed clusters with the most comprehensive monitoring and reporting solution on the market.

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.

Univa Grid Engine Add-on:

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

About Univa

Univa is the leading independent provider of software-defined computing infrastructure and workload orchestration solutions. Univa's intelligent cluster management software increases efficiency while accelerating enterprise migration to hybrid clouds. Millions of compute cores are currently managed by Univa products in industries such as life sciences, manufacturing, oil and gas, transportation and financial services. We help hundreds of companies to manage thousands of applications and run billions of tasks every day. Univa is headquartered in Chicago, with offices in Toronto and Munich. For more information, please visit **www.univa.com**.

UNIVA | **Univa Corporation** 2300 North Barrington Road, Suite 400, Hoffman Estates, IL, 60195 USA
Tel: +1.647.478.5901 www.univa.com