

Scalable Cube

Scalable Cube is an enterprise ready, supported distribution of Cubeman (cubejobscheduler.net), an open-source, workload management that supports a wide variety of HPC and analytic applications. Whether deployed on-site, on virtual infrastructure, or in the cloud, customers can take advantage of top-quality support services from HPC Scalable, helping ensure the success of managing their complex and mission critical HPC workloads.

Unlike other firms supporting open-source offerings, the HPC Scalable team has decades of experience developing production quality software for large distributed production environment. HPC Scalable employees and associates are active committers to the Cubeman community meaning that support is coming from people with in-depth knowledge of the code who are in a position to provide expert advice and solve problems quickly. Regardless of your workload management requirement, HPC Scalable can help you get productive fast with Scalable Cube software.

Highlights

- Robust, scalable, and high performance
- Compatible with mainstream commercial workload management software
- Leverage hundreds of existing integrations
- Advanced enterprise features
- Deploy in the cloud or on premises
- Top quality commercial support

Commercial Grade

Scalable Cube offers a variety of standard and available features for business and research, including license-aware scheduling pre-emption. It is production proven in environments with thousands of compute nodes, hundreds of thousands of cores, and millions of jobs. Features like fair-share, pre-emptive scheduling, job dependencies, NUMA support, reporting, and an easy-to-use web interface for job submission, monitoring, and management, make Scalable Cube the right choice for large, complex, and mission critical HPC environment.

Open-source components

- Avoid restrictive licensing arrangements
- Run on-premises or in the cloud
- Preserve flexibility, reduce license costs

Compatible with mainstream commercial workload management software

- Leverage hundreds of existing integrations
- Preserve investments in scripts & expertise

Enhanced scheduling features

- Pre-emptive scheduling to ensure SLAs
- Multi-level, Fairshare scheduling
- Application license-aware scheduling
- License / resource-based pre-emption
- Job Arrays
- Queue dispatch and run windows
- Exclusive scheduling, interactive job support
- Job limits per user, queue and host
- Job dependencies for multi-step workflows
- Built-in GPU support
- Docker container job support

Management features

- Easy-to-use web interface
- Flexible reporting subsystem
- Enhanced support for NUMA environments

Outstanding scalability & reliability

- Millions of jobs, thousands of hosts
- Dynamic master selection, n-way failover
- Low-latency, high-throughput scheduling
- Parallelized event handling for faster startup
- Additional tuning parameters for scalability
- Enhanced inter-daemon communications
- Pluggable authentication
- Never lose a job - job restart on host failure

Top-quality service and support

- Support team with decades of workload management experience
- Standard and premium support options
- Installation, configuring & tuning assistance
- Optional on-site implementation services
- Emergency patches, custom extensions
- Migration services

HPC Scalable

Case in point: Delivering outstanding efficiency gains for a tier-one semiconductor design firm

Working with a multi-national design firm with centers in the U.S., Europe and Asia, HPC Scalable assisted in deploying Scalable Cube workload management across 40,000 cores to help provide strategic flexibility for the client. Along with migration services and a global support agreement, HPC Scalable developed custom scheduling enhancements to meet the client's evolving business needs, while allowing the client to retain ownership & control of the resulting IP.

Benefits

- Faster time to deploy new scheduling features
- Reduced TCO (software licensing and support costs)
- Deployment flexibility, ownership of IP for competitive advantage

"By using Scalable Cube, we've been able to provide our clients with the high levels of support they need, while simultaneously reducing license costs for workload management"

Flexible Deployment

Enterprise customers can download and use open-source Cubeman either on-premises or with their choice of cloud provider for free. Customers needing enterprise features (e.g. license-based pre-emption, reporting capability), or requiring professional support, can deploy Scalable Cube for an annual fee based on the number of cluster hosts.

Independent Software Vendors (ISVs) can optionally bundle Cubeman with their cluster-aware applications, providing end-users with a compelling, supported alternative to costly commercial workload managers.

For clients that don't want to be concerned with infrastructure, Scalable Cube or the open source Cubeman can be easily deployed in a public cloud environment like AWS, ideal for Computer-aided Engineering (CAE), Electronic Design Automation (EDA), Life Sciences, Financial Services, and Machine Learning.

To learn more about Scalable Cube, visit <http://HPC Scalable.com>.



© Copyright HPC Scalable 2018

HPC Scalable Pte Ltd, 10 Anson Road, #26-04, International Plaza, Singapore S079903

<http://hpcscalable.com> E-mail: info@hpcscalable.com

February 2018, All Rights Reserved