

Release Notes for Open Grid Scheduler/Grid Engine

Version: Grid Engine 2011.11



New Features

Berkeley DB Spooling Directory Can Be Located on NFS

The Berkeley DB spooling framework has been enhanced such that any file systems can be used as the Berkeley DB backing store. Starting from the Grid Engine 2011.11 release, sites that require highly available qmaster hosts are recommended to place the spooling DB on a shared file system (e.g. NFS, including NFSv3 and earlier versions), such that standby masters (shadow hosts) can access the spooled configuration and job information written by the qmaster.

Additionally, the Berkeley RPC spooling server is now deprecated and support will be dropped in Q3 2013.

Portable Hardware Locality Library

The *hwloc* topology library is now used for the Job to Core Binding Feature. This version of Open Grid Scheduler/Grid Engine ships with *hwloc* v1.2.2, which has fixes that are important to AMD Opteron 6100-series (Magny-Cours) and 6200-series (Interlagos) processors.

In addition, this version of Open Grid Scheduler/Grid Engine extends the Job to Core Binding support to the following Operating Systems: AIX, BSD, and HP-UX.

CUDA GPU load sensor

A new load sensor for CUDA GPUs is available to Open Grid Scheduler/Grid Engine for GPGPU clusters. The load sensor uses the NVIDIA Management Library (NVML) to report the GPU device temperature, power usage, fan speed, ECC error, memory utilization, and load.

User Notification Mails Can Be Sent From A Configurable User ID

An optional prefix *user@* specifies the user under which the mailer is to be executed. The default user that invokes the mailer is still the root user.

Job Exit Status Available to the epilog Procedure

The epilog can get the exit status of the job from the environment variable `$SGE_JOBEXIT_STAT`.

ARM Linux port available

The ARM platform is now supported by Open Grid Scheduler/Grid Engine.

qmake upgraded to version 3.82

Distributed parallel make has been upgraded to version 3.82 from 3.78.1.

Support for Linux 3.0

Modifications were made to the arch string reporting and arch handling code to handle newer Linux kernel versions, including 3.0 & 3.1. Further, the architecture name of x86-64 has been changed from AMD64 to x64.

Perfstat Library for System Load on AIX

The standard AIX Perfstat API is now used for fetching processor utilization and memory usage on AIX. Open Grid Scheduler does not use non-portable way of fetching system information from the AIX kernel.

Support for Newer Versions of AIX

On AIX, as the Perfstat Library is now the default, Open Grid Scheduler/Grid Engine can now take advantage of the IBM AIX binary compatibility guarantee for AIX 7. The arch string of AIX has been adjusted from aix51 to aix to reflect operating system version independence.

Tango icons for qmon

This version of Open Grid Scheduler/Grid Engine ships with the new icons from the Tango Desktop Project.

Static Program Verifier

Static program analysis has been added to the Quality Assurance tests in Open Grid Scheduler/Grid Engine. Static analysis improves code quality by performing data-flow analysis on code and catch bugs in code that is exercised by traditional software test cases.

Technology Preview

Parallel Environment Queue Sort (PQS) Scheduler API

The PQS API, originally developed by Shannon V. Davidson at Raytheon, is now available as technology preview. Sites that have special network topologies and locality scheduling requirements can develop scheduler plugins to customize and fine tune scheduling decisions made by the default Grid Engine scheduler.

Pre-compiled binaries with the PQS API enabled will be available for download at a later date. To enable the PQS API from source compilation, the `-DSGE_PQS_API` flag can be used to enable the code.

Deprecated Features

Support for the following features will be dropped (EOL) in Q3 2013:

Berkeley RPC server

The installation procedure has been modified to remove the Berkeley RPC server as a spooling backend option. All fresh installations are required to use either the classic spooling backend or the Berkeley DB backend.

However, sites already using Berkeley RPC server spooling can keep their choice of spooling backend, but are recommended to migrate to the Berkeley DB backend before the final EOL date.

ibm-loadsensor

The ibm-loadsensor is not used by this version of Grid Engine to retrieve system load and memory utilization information.

AIX 4.3 (and below)

Support for AIX 4.3 or older was dropped by IBM.

PLPA

PLPA is deprecated by the newer hwloc topology library that is developed and maintained by the same open source project (Open MPI).

Solaris Processor Set binding

Queue Processor Binding method inherited from SGE 5.3

Linux with glibc 2.1

Supported Operating Systems

Operating System	Version
AIX	5.1 or above
HP-UX	11i (64-bit kernel)
Linux	Kernel 2.4, 2.6, 3.0 with glibc 2.2 or above
Mac OS X	10.6 or above
Solaris	10, 11 (x64 & SPARC, 64-bit kernel)
Windows	XP, Windows 7, Windows Server 2003 or above

Support Options

Open Grid Scheduler / Grid Engine is free and open source software, and will remain so in the foreseeable future. Optional commercial support is available under the *Scalable Grid Engine Support Program* from *Scalable Logic*.

Note that there are no special binaries used for the commercially supported version. The package downloadable from the Open Grid Scheduler project website or the Scalable Logic homepage can be used freely.

Further information is available at: <http://www.scalablelogic.com/>

Acknowledgments

We would like to thank the following people and open source projects who have contributed to this release of Open Grid Scheduler/Grid Engine:

- The hwloc project: especially Brice Goglin, Jeff Squyres
- Prakashan Korambath, University of California, Los Angeles
- Alain Champenois (for the new qmon icons)

And we would like to thank Sun Microsystems for releasing Sun Grid Engine under the SISSL open source license.