



UNIVA CORPORATION

GRID ENGINE DOCUMENTATION

---

# Grid Engine Release Notes

---

*Author:*  
Univa Engineering

*Version:*  
8.4.4

November 9, 2016



# Contents

<b>1 License</b>	<b>1</b>
<b>2 Supported Operating Systems, Versions and Architectures</b>	<b>6</b>
<b>3 Fixes and Enhancements</b>	<b>7</b>
3.1 Summary . . . . .	7
3.1.1 Define behaviour if epilog script dies because of signal . . . . .	7
3.1.2 Docker integration . . . . .	7
3.1.3 Host resolving and host_aliases file . . . . .	8
3.1.4 Scheduler specific changes . . . . .	8
3.1.5 Execution daemon specific changes . . . . .	9
3.1.6 Intel® Xeon Phi™ x200 (Knights Landing) integration . . . . .	9
3.2 Full List of Fixes and Enhancements . . . . .	10
<b>4 Upgrade Notes</b>	<b>20</b>
4.1 Upgrade Requirements . . . . .	20
<b>5 Compatibility Notes</b>	<b>21</b>
<b>6 Known Issues and Limitations</b>	<b>21</b>

# 1 License

## TERM SOFTWARE LICENSE AND SUPPORT AGREEMENT

This agreement is between the individual or entity agreeing to this agreement and Univa Corporation, a Delaware corporation (Univa) with its registered office at 2300 N Barrington Road, Suite 400, Hoffman Estates, IL 60195.

1. SCOPE: This agreement governs the licensing of the Univa Software and Support provided to Customer.
  - Univa Software is defined as the Univa software described in the order, all updates and enhancements provided under Support, its software documentation, and license keys (Univa Software), which are licensed under this agreement. This Univa Software is only licensed and is not sold to Company.
  - Third-Party Software/Open Source Software licensing terms are addressed on the bottom of this agreement.
2. LICENSE. Subject to the other terms of this agreement, Univa grants Customer, under an order, a non-exclusive, non-transferable, renewable term license up to the license capacity purchased to:
  - (a) Operate the Univa Software in Customer's business operations and
  - (b) Make a reasonable number of copies of the Univa Software for archival and backup purposes.

Customer's contractors and majority owned affiliates are allowed to use and access the Univa Software under the terms of this agreement. Customer is responsible for their compliance under the terms of this agreement.

The initial term of this license is for a period of one year from date hereof to be automatically renewed at each anniversary unless a written notification of termination has been received 60 days prior to each anniversary.

3. RESTRICTIONS. Univa reserves all rights not expressly granted. Customer is prohibited from:
  - (a) assigning, sublicensing, or renting the Univa Software or using it as any type of software service provider or outsourcing environment or
  - (b) causing or permitting the reverse engineering (except to the extent expressly permitted by applicable law despite this limitation), decompiling, disassembly, modification, translation, attempting to discover the source code of the Univa Software or to create derivative works from the Univa Software.
4. PROPRIETARY RIGHTS AND CONFIDENTIALITY.
  - (a) Proprietary Rights. The Univa Software, workflow processes, designs, know-how and other technologies provided by Univa as part of the Univa Software are the proprietary property of Univa and its licensors, and all rights, title and interest in and to such items, including all associated intellectual property rights, remain only with Univa.

The Univa Software is protected by applicable copyright, trade secret, and other intellectual property laws. Customer may not remove any product identification, copyright, trademark or other notice from the Univa Software.

- (b) Confidentiality. Recipient may not disclose Confidential Information of Discloser to any third party or use the Confidential Information in violation of this agreement.
- (c) Confidential Information means all proprietary or confidential information that is disclosed to the recipient (Recipient) by the discloser (Discloser), and includes, among other things:
  - any and all information relating to Univa Software or Support provided by a Discloser, its financial information, software code, flow charts, techniques, specifications, development and marketing plans, strategies, and forecasts
  - as to Univa the Univa Software and the terms of this agreement (including without limitation, pricing information).
- (ii) Confidential Information excludes information that:
  - was rightfully in Recipient's possession without any obligation of confidentiality before receipt from the Discloser
  - is or becomes a matter of public knowledge through no fault of Recipient
  - is rightfully received by Recipient from a third party without violation of a duty of confidentiality
  - is independently developed by or for Recipient without use or access to the Confidential Information or
  - is licensed under an open source license.

Customer acknowledges that any misuse or threatened misuse of the Univa Software may cause immediate irreparable harm to Univa for which there is no adequate remedy at law. Univa may seek immediate injunctive relief in such event.

5. PAYMENT. Customer will pay all fees due under an order within 30 days of the invoice date, plus applicable sales, use and other similar taxes.
6. WARRANTY DISCLAIMER. UNIVA DISCLAIMS ALL EXPRESS AND IMPLIED WARRANTIES, INCLUDING WITHOUT LIMITATION THE IMPLIED WARRANTY OF TITLE, MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE UNIVA SOFTWARE MAY NOT BE ERROR FREE, AND USE MAY BE INTERRUPTED.
7. TERMINATION. Either party may terminate this agreement upon a material breach of the other party after a 30 day notice/cure period, if the breach is not cured during such time period. Upon termination of this agreement or expiration of an order, Customer must discontinue using the Univa Software, de-install it and destroy or return the Univa Software and all copies, within 5 days. Upon Univa's request, Customer will provide written certification of such compliance.
8. SUPPORT INCLUDED. Univa's technical support and maintenance services (Support) is included with the fees paid under an order. Univa may change its Support terms, but Support will not materially degrade during any paid term. More details on Support are located at [www.univa.com/support](http://www.univa.com/support)

9. **LIMITATION OF LIABILITY AND DISCLAIMER OF DAMAGES.** There may be situations in which, as a result of material breach or other liability, Customer is entitled to make a claim for damages against Univa. In each situation (regardless of the form of the legal action (e.g. contract or tort claims)), Univa is not responsible beyond:

- (a) the amount of any direct damages up to the amount paid by Customer to Univa in the prior 12 months under this agreement and
- (b) damages for bodily injury (including death), and physical damage to tangible property, to the extent caused by the gross negligence or willful misconduct of Univa employees while at Customer's facility.

Other than for breach of the Confidentiality section by a party, the infringement indemnity, violation of Univa's intellectual property rights by Customer, or for breach of Section 2 by Customer, in no circumstances is either party responsible for any (even if it knows of the possibility of such damage or loss):

- (a) loss of (including any loss of use), or damage to: data, information or hardware
- (b) loss of profits, business, or goodwill or
- (c) other special, consequential, or indirect damages

10. **INTELLECTUAL PROPERTY INDEMNITY.** If a third-party claims that Customer's use of the Univa Software under the terms of this agreement infringes that party's patent, copyright or other proprietary right, Univa will defend Customer against that claim at Univa's expense and pay all costs, damages, and attorney's fees, that a court finally awards or that are included in a settlement approved by Univa, provided that Customer:

- (a) promptly notifies Univa in writing of the claim and
- (b) allows Univa to control, and cooperates with Univa in, the defense and any related settlement.

If such a claim is made, Univa could continue to enable Customer to use the Univa Software or to modify it. If Univa determines that these alternatives are not reasonably available, Univa may terminate the license to the Univa Software and refund any unused fees.

Univa's obligations above do not apply if the infringement claim is based on the use of the Univa Software in combination with products not supplied or approved by Univa in writing or in the Univa Software, or Customer's failure to use any updates within a reasonable time after such updates are made available.

This section contains Customer's exclusive remedies and Univa sole liability for infringement claims.

11. **GOVERNING LAW AND EXCLUSIVE FORUM.** This agreement is governed by the laws of the State of Illinois, without regard to conflict of law principles. Any dispute arising out of or related to this agreement may only be brought in the state of Illinois. Customer consents to the personal jurisdiction of such courts and waives any claim that it is an inconvenient forum. The prevailing party in litigation is entitled to recover its attorney's fees and costs from the other party.

12. **MISCELLANEOUS.**

- (a) **Inspection.** Univa, or its representative, may audit Customer's usage of the Univa Software at any Customer facility. Customer will cooperate with such audit. Customer agrees to pay within 30 days of written notification any fees applicable to Customer's use of the Univa Software in excess of the license.
- (b) **Entire Agreement.** This agreement, and all orders, constitute the entire agreement between the parties, and supersedes all prior or contemporaneous negotiations, representations or agreements, whether oral or written, related to this subject matter.
- (c) **Modification Only in Writing.** No modification or waiver of any term of this agreement is effective unless signed by both parties.
- (d) **Non-Assignment.** Neither party may assign or transfer this agreement to a third party, except that the agreement and all orders may be assigned upon notice as part of a merger, or sale of all or substantially all of the business or assets, of a party.
- (e) **Export Compliance.** Customer must comply with all applicable export control laws of the United States, foreign jurisdictions and other applicable laws and regulations.
- (f) **US Government Restricted Rights.** The Univa Software is provided with RESTRICTED RIGHTS. Use, duplication, or disclosure by the U.S. government or any agency thereof is subject to restrictions as set forth in subparagraph (c)(I)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 or subparagraphs (c)(1) and (2) of the Commercial Computer Software Restricted Rights at 48 C.F.R. 52.227-19, as applicable.
- (g) **Independent Contractors.** The parties are independent contractors with respect to each other.
- (h) **Enforceability.** If any term of this agreement is invalid or unenforceable, the other terms remain in effect.
- (i) **No PO Terms.** Univa rejects additional or conflicting terms of a Customer's form-purchasing document.
- (j) **No CISG.** The United Nations Convention on Contracts for the International Sale of Goods does not apply.
- (k) **Survival.** All terms that by their nature survive termination or expiration of this agreement, will survive.

Additional software specific licensing terms:

Grid Engine incorporates certain third-party software listed at the URL below. These licenses are accepted by use of the software and may represent license grants with restrictions in which Univa is bound to provide. We are hereby notifying you of these licenses.

Unicloud Kits

- Third Party Software is defined as certain third-party software which is provided along with the Univa Software, and such software is licensed under the license terms located at: <http://www.univa.com/resources/licenses/>
- Open Source Software is defined as certain opens source software which is provided along with the Univa Software, and such software is licensed under the license terms located at: <http://www.univa.com/resources/licenses/>

## Grid Engine

- Third Party Software is defined as certain third-party software which is provided along with the Univa Software, and such software is licensed under the license terms located at: <http://www.univa.com/resources/licenses/>
- Open Source Software is defined as certain opens source software which is provided along with the Univa Software, and such software is licensed under the license terms located at: <http://www.univa.com/resources/licenses/>

Rev: August 2014



## 2 Supported Operating Systems, Versions and Architectures

Univa Grid Engine supports various platforms, hardware architectures and versions of operating systems. Find the full list in following table:

Operating System	Version	Architecture
SLES	10,11,12	x86, x86-64
RHEL	5 or higher, 6 or higher, 7	x86, x86-64
CentOS	5 or higher, 6 or higher, 7	x86, x86-64
Oracle Linux	5 or higher, 6 or higher, 7	x86, x86-64
Ubuntu	10.04LTS - 16.04LTS	x86, x86-64
Oracle Solaris	10, 11	x86-64, SPARC 64bit
HP-UX	11.0 or higher	64bit
IBM AIX	6.1 or later	64bit
Apple OS X	10.8 (Mountain Lion) or higher	x86, x86-64
Microsoft Windows	XP Professional (SP3)	32 bit
Microsoft Windows	Server 2003 / 2003 R2	32 bit
Microsoft Windows	Vista Enterprise / Ultimate	32 and 64bit
Microsoft Windows	Server 2008 / 2008 R2	32 and 64bit
Microsoft Windows	7 Professional / Enterprise / Ultimate	32 and 64bit
Microsoft Windows	Server 2012 / 2012 R2	32 and 64bit
Microsoft Windows	8 / 8.1 Pro / Enterprise	32 and 64bit
Microsoft Windows	10 Pro / Enterprise	32 and 64bit

Table 1: Supported Operating Systems, Versions and Architectures

*PLEASE NOTE:* Hosts running the Microsoft Windows operations system cannot be used as master or shadow hosts.

*PLEASE NOTE:* Univa Grid Engine qmaster is fully supported on Linux and Solaris. We provide binaries in Univa Grid Engine for running the qmaster on other operating systems but they are not supported and delivered as a courtesy. If you require qmaster support on other architectures please contact us at support@univa.com.

*PLEASE NOTE::* if you require Univa Grid Engine support for older versions of the above operating systems please contact our sales or support team.

## 3 Fixes and Enhancements

### 3.1 Summary

#### 3.1.1 Define behaviour if epilog script dies because of signal

The “`execd_params RESCHEDULE_ON_KILLED_EPILOG`” (see also man page `sgc_conf(5)`) was introduced to allow to prevent rescheduling of a job if the epilog dies because of a signal. The default value is “`TRUE`” or “`1`”, which leads to the old behaviour, i.e. the queue is set to error state and the job is re-enqueued in the pending job list. If set to “`FALSE`” or “`0`”, the job finishes, its data is written to accounting and the job is removed from the job list. The “`failed`” field of such a job is set to “`15 : in epilog`”, but the “`exit_status`” is the one of the job itself. This is useful in cases where the epilog does just unimportant cleanup work and the job workflow should not be interrupted because of problems doing this cleanup work. To detect if an epilog was signaled solely its exit status is taken into account, i.e. an epilog that exits with a status  $> 127$  is handled like an epilog that was signaled.

#### 3.1.2 Docker integration

##### Features of the Docker integration

Univa Grid Engine 8.4.4 provides an integration with Docker containers. This allows users to specify that a job has to be started in a Docker container that was created from a specific Docker image. Please see chapter 7.9 “Jobs using Docker Containers” in the UsersGuide and chapter 2.5.2 “Configuring Queues” in the AdminsGuide for details.

##### New in Univa Grid Engine 8.4.2

Univa Grid Engine allows now to use the Entrypoint that is defined in the Docker image the container is created from. This implies several limitations, but makes it easier to start “batch” containers. To let a job use the defined Entrypoint, submit it as binary job with job script “`NONE`”, e.g.:

```
$ qsub -l docker,docker_images="*myapp:latest*" -b y NONE
```

This is allowed for sequential jobs only. Interactive and parallel jobs are not supported. Please see chapter 7.9 “Jobs using Docker Containers” in the UsersGuide for details.

##### Known issues and limitations of the Docker integration

- Requirements:
  - Host architecture: `lx-amd64`. Other architectures are not supported yet.
  - Docker version: from 1.8.3 to 1.12.3. The Docker API is often changed and not in a backwards compatible way, so later versions of Docker are not supported. Because problems with the Docker API are hard to diagnose, if Univa Grid Engine detects a Docker daemon of a version that is not in the supported range on an execution host, it does not even report this host as an available Docker host, instead it reports it as a host that does not have Docker installed.
- Sometimes, the Docker daemon responds with a valid, but empty message to the “`docker images`” request. The execution daemon cannot distinguish this from a valid response of a Docker installation with no images available. If the execution daemon has a job to start

when it gets such an empty message, then this job fails, because the execution daemon assumes the image was deleted and the job cannot be started.

- Checkpointing a job that runs in an container is not supported.
- Only “builtin” interactive jobs are supported in containers.
- The `-xd` submission switch is added as an experimental feature. There might be issues with more sophisticated use cases, e.g. the combination of multiple switches for network settings or duplicate specification of data. The `-xdv` switch is deprecated and can be replaced by `-xd -v`.

### 3.1.3 Host resolving and `host_aliases` file

- The host name reported by a load sensor will get resolved at execution side now. Settings in the `host_aliases` file are now also used for load reports via external load sensors.
- At startup of `qmaster` it will verify hostname resolution of spooled objects in the database. If the resulting host names have changed for hosts referenced in the configuration list, the admin host list, the `execd` list or in the submit host list then all spooled data objects will get adjusted to match the new resulting hostnames.
- The resource hostname request (`-l h=<expr>`) now supports resolving plain hostnames in regular expression requests. This also includes using aliased hostnames.
- It is now supported to add new entries to the `host_aliases` file while `qmaster` daemon is running. Changing already referenced host names will need a `qmaster` restart. If a change to the current `host_aliases` needs a restart the `qmaster` process will log this information into the messages file. In order to change already active host aliases during runtime the corresponding hosts must be removed from Univa Grid Engine. A hostname change that affects the interface of a running `sge_qmaster` service or `sge_execd` service will always need a restart of the service (no change in behavior).
- The client tools `qping` and `gethostbyname` got a new option (`-all_rr`) to show the resulting host name after resolving on the service. This tools can be used to wait until a running `qmaster` has taken over changes to the `host_aliases` file.
- It is possible to set individual `qmaster` parameters for internal host name cache (via `qmaster_params`). The number of cached entries can be obtained via `qping` or `PROF_COMMLIB_TIME` `qmaster` param.

### 3.1.4 Scheduler specific changes

- The scheduler configuration parameter “`params`” can be used to enable profiling (`PROFILE=true`). In combination with the value of “`WARN_DISPATCHING_TIME`” it can be used to show additional information about the longest and shortest job scheduling time.
- The scheduler profiling is now thread based on architectures that supports thread specific user and system time measurement (linux kernel  $\geq 2.6.26$  and solaris operating systems). This results in correct system and user times for the scheduler thread. On other architectures the user and system times are measured for the entire process. This means that only the wallclock times are reflecting the overhead of the scheduler thread. System and user times will show the usage of all threads of the process in this case and is therefore not only scheduler thread specific.
- The profiling summary for the scheduler thread will contain information about time used for RQS (Resource Quota Sets) calculation. If some RQS Rule has an unexpected high

influence on scheduling time the profiling data for this rule will also occur in the profiling output.

- Defined RQS are now sorted alphanumerically on their names in order to have the possibility to define a clear RQS processing order while the scheduler is dispatching the jobs. The processing order might have an influence on the scheduling time and can be optimized now. The rule that limits the most should be the first one.
- The schedule info messages provided by qsub/qalter -w ... command might now provide different messages because filter rules which are used by scheduler are active now. The resulting scheduling information provided by the simulated scheduling run via qsub/qalter -w should produce better results now.

### 3.1.5 Execution daemon specific changes

- The implementation of GE-5849 made it necessary to change limit setting and observation behavior especially when the -masterl switch is used. For the parallel environment settings “master\_forks\_slave” the resulting limit will be increased by the limit used for the master task itself (see man page sge\_pe). The “master\_forks\_slave” PE settings are now also used for PE setting “control\_slaves” set to false (Loose Integration). If you now hit a limit for jobs using PEs with “control\_slaves=false” after installing Univa Grid Engine 8.4.2 you can use “master\_forks\_slave=TRUE” to increase the limit depending on the nr of slots for the job.

### 3.1.6 Intel® Xeon Phi™ x200 (Knights Landing) integration

Univa Grid Engine 8.4.4 provides an integration of Intel® Xeon Phi™ x200 (Knights Landing) Processors. The pre-compiled load-sensor automatically detects the current Cluster as also the current Memory Mode of the x200 machine. Furthermore the current MCDRAM distribution gets reported. Please see “Configure and Install Intel Xeon Phi x200 (Knights Landing) Processors support” in the AdminsGuide for details.

## 3.2 Full List of Fixes and Enhancements

Univa Grid Engine 8.4.0alpha (also fixed for a 8.3 patch release)

- GE-2716 interactive jobs (qlogin, qrsh without command) don't set the TZ environment variable correctly
- GE-3392 Job reservation with wildcards in PE names doesn't work correctly
- GE-3858 pe job does not start: cannot run in PE "my\_pe" because it only offers 2147483648 slots
- GE-4229 Reduce executable sizes by removing extra symbols
- GE-4288 Confusing message on h\_rt or s\_rt limit
- GE-4293 qsub -w e -l exclusive=true rejects job, even if the request is valid (THIS FIX WAS WITHDRAWN FOR 8.4.0)
- GE-4296 the unit of the io usage value is missing in qstat -j <job\_id> output and not explained in the man page
- GE-4297 report io wait time and number of io operations
- GE-4384 User lists do not handle space separated user names correctly
- GE-4404 Rounding error, when memory values are reported by execd
- GE-4641 jobs with high job\_id may starve when job\_id roll over happens when wait\_time is not recognized in job priority calculation
- GE-4739 print unique thread names in messages file
- GE-4943 shepherd closes FDs, needed by AD authentication
- GE-4983 port DRMAA2 C API to other architectures
- GE-5033 setting ENABLE\_SUBMIT\_LIB\_PATH in qmaster\_params has no effect for LD\_PRELOAD env variable
- GE-5045 qlogin and qrsh without command does not inherit expected variables (e.g. TERM)
- GE-5074 sessionusers ACL not present after installation
- GE-5081 wrong reference to "MONITOR\_TIME" in admin guide
- GE-5156 Non-existing paths for input and error files should be implicitly created
- GE-5268 event client id of DRMAA2 event clients displayed as "unknown"
- GE-5289 add a note about the msvc redist dll to the installation guide
- GE-5340 hard coded timeout for PE \*\_proc\_args, prolog, epilog of 120 s not documented and changeable
- GE-5401 h\_vmem kill done by execd even when cgroups is setup to handle this limit
- GE-5486 introduce per job profiling
- GE-5487 introduce performance improvement for -masterq switch
- GE-5536 requesting more than one tmpdir per job
- GE-5543 drmaa2 functions drmaa2\_get\_drmaa\_name() and drmaa2\_get\_drmaa\_version() are missing
- GE-5557 Add functionality to search primary and secondary groups when '@' used to specify group in Grid Engine
- GE-5587 allow the Cray XC load sensor to update the slots counter in the queue
- GE-5588 communication errors at first startup not logged into /tmp/execd\_messages.<pid> file
- GE-5595 GetAdminUser() fails and is setting ADMINUSER to 'default'

- GE-5597 with `accounting_summary=true`, "wallclock" usage of PE jobs is wrong
- GE-5605 test and release `qping.exe` for win-x86
- GE-5619 `drmaa2_*session_create()` should ignore contact string instead of expecting null
- GE-5620 `drmaa2_open_msession()` should return a `msession` handle even `msession` is opened already
- GE-5624 Unix group entries in predefined `userlist` as well as `manager` or `operator` list are ignored
- GE-5625 CUDA and XEON PHI complex attribute installation fails
- GE-5637 `jsv` task job related params are not transferred for 1 task arrays
- GE-5638 Windows (win-x86) does not forward or collect the job exit code
- GE-5639 `gid` range observation not always un-blocking additional group ids
- GE-5641 user list man page should mention all predefined lists or list with a special meaning
- GE-5643 `qalter -when` now does not work for PE jobs with exclusive consumables
- GE-5647 `qconf -mu, -au, -du` triggers crashes when RQS'es are configured
- GE-5649 add an automated TS test for the error scenario
- GE-5653 `ulx-amd64` packages seem not to be built with HWLOC library
- GE-5654 `execd` crashes on win-x86 when sending a job related admin mail
- GE-5655 `Qmaster` get unresponsive after error "invalid task number 0"
- GE-5656 `qconf -ke` does not completely cleanup `execd` information
- GE-5663 setting host to unheard might block `qmaster` under certain conditions
- GE-5664 array jobs can oversubscribe consumables with `qalter -when` now
- GE-5672 develop library to communicate with the Docker Remote API
- GE-5673 job lost detection is logging strange error regarding granted resource list
- GE-5674 `qmaster` crash can be trigger with `qconf -mattr` on an `execd` object.
- GE-5678 implement load sensor that reports docker version and available images
- GE-5679 add a `"-xdv"` switch to the submit clients to allow the user to specify directories to mount into a Docker container
- GE-5680 forward information about the selected Docker image and the paths to mount to the shepherd
- GE-5681 make sure Docker jobs are not registered in PDC/PTF
- GE-5682 use Docker API to get online usage of a job
- GE-5684 cleanup finished Docker containers after job ended
- GE-5685 implement a `coshepherd` that is started in a Docker container to keep it alive and to run methods and the job
- GE-5687 use Docker Remote API to run methods and job and signal container
- GE-5689 fix support for foreign filedescriptors in `commlib`
- GE-5690 `qquota "limit"` value rendered as `-2^31` for large limits
- GE-5693 fix container stats acquisition via docker communication library
- GE-5694 Designation of events in logs - many events are labeled as 'Errors', where perhaps they should be 'Warnings'.
- GE-5703 non-admin user cannot trigger preemption of own jobs
- GE-5718 `lothread` needs to send reservation information to

License Orchestrator

- GE-5722 preempted LO job stays in "dr" state after qdel
- GE-5723 qalter -p not transferred to LO
- GE-5728 suspend/unsuspend endless loop for Preemption of jobs with the same priority
- GE-5731 qmod -p is only allowed on admin-hosts
- GE-5732 User should get a warning when preempting a higher prio job
- GE-5734 execd dumps core when a large tightly integrated parallel job is submitted to that host
- GE-5737 for Docker jobs, adjust paths in environment variables set by the container\_coshepherd
- GE-5740 fix libnuma dependency of shepherd to allow starting the coshepherd in the container
- GE-5741 rework container start mechanism in shepherd
- GE-5745 Add a man page containing all kind of error codes
- GE-5746 give the docker containers meaningful names
- GE-5750 cleanup container creation struct
- GE-5756 make the 'docker' and 'docker\_images' complexes builtins
- GE-5757 qmaster/execd/job protocol lacks check for active\_jobs cleanup
- GE-5759 show statistics about request types in worker and reader request queues
- GE-5760 add information about start and end (duration) of requests to the DEBUG log\_level
- GE-5763 improve per thread profiling
- GE-5764 create a man page sge\_diagnostics that summarizes and explains output of profiling/monitoring/logging/debug functionalities of UGE and LO
- GE-5767 Add a means to limit the job script size
- GE-5768 keep\_active sends all files of a job regardless of the file size
- GE-5770 add a means to switch on and off debug logging (DPRINTF) of sge\_qmaster during runtime
- GE-5775 improve logging in all Docker related components
- GE-5779 improve error handling in communication with docker daemon
- GE-5796 qmaster crashes with MONITOR\_REQUEST\_QUEUES=1
- GE-5798 performance regression with RQS rules
- GE-5803 enhance error logging of Windows (win-x86) qloadsensor.exe
- GE-5807 enhance scheduler profiling to show information for RQS calculation
- GE-5813 On Windows (win-x86), the execd cannot send the first CR to the qloadsensor.exe, causing it to never send load
- GE-5814 On Windows (win-x86), the execd logs a misleading warning about load sensors at startup time
- GE-5816 commlib external file descriptor support not thread save
- GE-5817 Docker jobs fail if the mount points of the binds are not unique
- GE-5818 resource quota cleanup for profiling
- GE-5825 improve qdel performance for bulk job deletions
- GE-5829 Docker jobs fail to start on some Linuxes because MemorySwappiness cannot be set
- GE-5845 possible race condition in event master at event client registration or total update

- GE-5851 jemalloc 3.6.0 can cause qmaster core dump
- GE-5853 qalter -tc prints incorrect success message
- GE-5858 job\_load\_adjustments may prevent any job dispatching in scheduler run after parallel job was scheduled
- GE-5860 communication specific enhancements for profiling and startup behavior
- GE-5865 keep\_active option does not copy all job related files into faulty job directory
- GE-5870 confirm Windows 10 Pro/Enterprise support and add it to list of supported OS
- GE-5876 where and what data structures are not used to prepare data for event clients
- GE-5878 inplace upgrade with postgres spooling breaks the upgrade script (inst\_sge -upd)
- GE-5880 profiling shows zero value for utilization in some scheduler profiling lines
- GE-5894 cluster queues are rejected due to missing project even if job has a project request
- GE-5895 hosts or qinstances are skipped by dispatch algorithm in scheduler but no valid reason is shown.
- GE-5902 a pe job requesting a per slot memory resource is not scheduled despite sufficient resources available
- GE-5907 shepherd aborts after a tightly integrated job was killed
- GE-5915 locale of qmaster process gets distorted by JVM\_thread
- GE-5917 jobs are not dispatched with open ended PE requests
- GE-5925 wrong qdel message when a job is already in deletion
- GE-5929 sge\_qmaster crashes when submitting a job to advance reservation using d\_rt.
- GE-5932 qghost and qmon round NLOAD to integer

#### Univa Grid Engine 8.4.0alpha

- GE-4497 PE job is not scheduled when a non-requestable consumable is setup in global host
- GE-4603 Job 124205 cannot run in PE "OpenMP" because it only offers 0 slots
- GE-4908 native Windows (win-x86) binaries can't find the SGE\_ROOT directory if it is the root directory of a share
- GE-5123 qdel array syntax from manpage fails
- GE-5129 regular "ckpt\_command" in CKPT interface not executed
- GE-5135 user has to login at least one time on each native Windows (win-x86) exec host to get the PROFILE created
- GE-5258 Using qconf for creating GDI sessions always returns exit status 1
- GE-5345 UGE to auto resolve host\_aliases
- GE-5509 host\_aliases not working for resource hostname OR request
- GE-5510 host\_aliases not working for qconf -purge request
- GE-5528 hostname resolving changes should trigger update of all affected data objects at qmaster/execd daemon startup



- GE-5559 Grid Engine upgrade procedure is running into an issue with `chmod` call
- GE-5667 describe in win-x86 installer and documentation that the UGE Starter Service doesn't work with mounted network directories
- GE-5671 integration of Docker into UGE
- GE-5692 enhance `sge_container_shepherd` to handle `stdin/stdout/stderr` stream to allow interactive and parallel jobs
- GE-5710 changed host aliases can trigger `qmaster abort()` at startup
- GE-5712 need concurrent array jobs where either all tasks run concurrently or no task at all
- GE-5714 unbelievable high `ru_wallclock` values in accounting
- GE-5739 `qmaster` installation script does not add admin host if its hostname cannot be resolved (error message unclear)
- GE-5747 `sharetree` might be wrong if host clock changes
- GE-5758 possible segmentation fault in `commlib` when static clients try do reconnect
- GE-5776 bind lists of UGE directories properly into the Docker container
- GE-5780 update `openssl` library to current version 1.0.2
- GE-5781 `exec` host goes into unknown (u) state when the system time is set to an earlier time
- GE-5787 reported `wallclock` time is too low when the system time is set to an earlier time
- GE-5792 docker jobs are reported as failed on `execd` restart
- GE-5793 `Qmaster` uninstall script tries to remove `startup-script` even if it was not installed
- GE-5795 exit dispatching loop immediately when shutdown of scheduler thread is triggered
- GE-5820 update Cray XC documentation
- GE-5821 remove basic Docker integration that uses `load sensor` and `starter_method` when real Docker integration works
- GE-5824 fix spelling mistakes in Grid Engine output messages
- GE-5834 bad performance of RQS rules on host groups
- GE-5836 during scheduling messages explaining why a job is not scheduled are generated but not used
- GE-5844 ship SGI MPT integration in Univa Grid Engine `mpi` folder
- GE-5877 `max_aj_instances` and `-tc` submit option are not respected with enrolled tasks
- GE-5879 when `max_aj_instances` is set to 0 (unlimited) the submit option `-tc` does no longer have effect
- GE-5913 On native Windows (win-x86), improve error logging of `wl_connect_to_service()` function and subfunctions
- GE-5920 not all functions retrieving `passwd` information do resize used buffer if it turns out that it is too small
- GE-5935 header of HTTP response from Docker daemon > 1.9 not handled properly
- GE-5937 Typo in UGE Admin Guide
- GE-5955 `host_aliases` not working for `sge_shadowd`
- GE-5962 the Docker daemon doesn't download an image when its not

available locally

#### Univa Grid Engine 8.4.0beta1

- GE-5956 update jemalloc to version 4.1.0
- GE-5975 upgrade reports: The built-in complex "docker\_images" cannot be deleted
- GE-5977 update openssl library to version 1.0.2h
  - fixed several memory leaks in sge\_qmaster

#### Univa Grid Engine 8.4.0beta2

- GE-5983 qmake crashes on sol-sparc64
- GE-5589 changes to host\_aliases file should be updated when hosts are re-resolved

#### Univa Grid Engine 8.4.0

- GE-5991 host names used for host\_aliases should be handled case insensitive
- GE-5994 sge\_qmaster startup fails with critical "setup failed" logging message
- GE-5995 job is executed even if prolog fails
- GE-5999 Preempted jobs stay in P-state
- GE-6006 A negative posix priority at submission results in a very high priority value in report
- GE-6008 dbwriter installation installs incorrect database version 14 instead of 15
- GE-6009 dbwriter cannot parse accounting record with io operations
- GE-6010 upgrade script does not recognize patch releases, e.g. 8.3.1p9
- GE-6012 Syntax error in installer script causes wrong error message during install

#### Univa Grid Engine 8.4.1

- GE-4293 qsub -w e -l exclusive=true rejects job, even if the request is valid
- GE-5850 allow to specify more Docker properties when submitting a job
- GE-6016 submitting job dependencies and deleting them again can trigger huge amount of qmaster mem usage
- GE-6028 User/group management done via Windows Active Domain might break UGE
- GE-6030 Introduce possibility to switch of commlib's internal hostname cache
- GE-6031 on native Windows (win-x86), the shepherd of running jobs can produce huge trace files if the UGE job starter service ends the connection unexpectedly
- GE-6036 job loss on exed restart after host\_aliases changes
- GE-6039 clients report "failed to extract authentication information" error
- GE-6042 Scheduling run might take longer than with previous versions of UGE
- GE-6052 cwd is not mapped into docker container bind
- UWSA-80 UWSA-135 wrong headline in installer
- UWSA-82 Symmetric Data Types
- UWSA-83 UGERest installation is broken, the way how it is documented
- UWSA-84 installer should check for dependencies when it starts
- UWSA-85 installer can not create certificates
- UWSA-87 UWSA-135 UGE REST API exports complexentrys instead of complexentries

- UWSA-91 job submission template uses unimplemented DRMAA1 job template attributes
- UWSA-96 documentation contains a broken JSON description of a job object
- UWSA-99 UWSA-135 allow to modify (qalter) the POSIX priority of a job
- UWSA-100 Create combined rest request to "get all host summary"
- UWSA-101 Create combined rest request to "get all clusterqueue summary"
- UWSA-102 CPU usage returned from hostsummary REST call always 0
- UWSA-103 JOBLIST attribute in hostsummary always empty even if jobs are running
- UWSA-104 JOB COUNT attribute in hostsummary call always 0
- UWSA-105 usedSlots in hostsummary always 0
- UWSA-106 Misspelled attribute parrallelEnv
- UWSA-110 inst\_ugerest fails in non CSP mode since the keystore cannot be created correctly
- UWSA-112 run / debug rest service from IDE
- UWSA-113 UGERestService.pdf - misleading usage for range urls
- UWSA-114 UGERestService.pdf - java binding section missing format for challenge.txt file
- UWSA-115 UGEWSClientApi deleteExecHost(id) results in SEGV and crash
- UWSA-116 UGEWSClientApi getExecHostById(id) when id doesn't exists returns confusing message
- UWSA-117 UGEWSClientApi modifyExecHost(object) results in SEGV and crash
- UWSA-120 Adjust refresh rate for hostsummary to make job updates available in time
- UWSA-121 verify if hostsummary and clusterqueuesummary are supported in client api bindings
- UWSA-133 HostSummary Job start/submit times are incorrect
- UWSA-136 HostSummary should show ResourceValues as numeric values where possible
- UWSA-138 Missing support for 'expand' type RQS rules
- UWSA-139 Unable to get (and perhaps set) Scheduler Configuration
- UWSA-141 Sharetree Puts/Posts/Deletes Don't seem to work
- UWSA-142 ugerest 8.3.1p10 crashed
- UWSA-143 Add block on job manager cache during cache update
- UWSA-144 RuntimeException for UserResourceRange call
- UWSA-145 The project resource doesn't handle xacl properly
- UWSA-146 Expansion of Hosts property is broken on RQS
- UWSA-147 In 8.3.1p10 still unable to operate on schedconf
- UWSA-151 Fetch additional fields for JobManager in getQinstanceSummary
- UWSA-152 ugerest script fails in line 102 due to invalid comparison
- UWSA-154 setfileperm.sh does no longer setuid bit for authuser
- UWSA-156 Job Context Variables no longer returned
- UWSA-158 ugerest service is not starting if the key and keystore passwords are different
- UWSA-159 No support for complex resources with '/' in their name
- UWSA-160 inst\_ugerest fails on root->nobody file systems
- UWSA-161 ugerest startup script fails when a stale pidfile exists

#### Univa Grid Engine 8.4.1p1 (windows package patch)

- GE-5135 user has to login at least one time on each native Windows (win-x86) exec host to get the PROFILE created
- GE-6022 native Windows (win-x86) execution daemon installer unnecessarily

- needs sgepasswd file
- GE-6023 native Windows (win-x86) submit clients need private keys of sgepasswd file
- GE-6031 on native Windows (win-x86), the shepherd of running jobs can produce huge trace files if the UGE job starter service ends the connection unexpectedly
- GE-6045 on native Windows (win-x86), the win\_getpwnam\_r() always tries to load the user profile, even if called with insufficient permissions
- GE-6076 on native Windows (win-x86), all job spool directories are deleted when the first job on an execution host finishes

#### Univa Grid Engine 8.4.2

- GE-3146 resource reservation is broken with SGE calendar
- GE-4158 Some of the job class attributes are incorrect in man page and users guide
- GE-4672 default\_jc and enforce\_jc are not documented in the man page
- GE-5492 cuda loadsensor bash script compatibility issue
- GE-5547 The install\_cuda\_complexes.sh doesn't handle '\n' correctly
- GE-5577 Server side JSV parameters l\_hard, l\_soft and masterl contains job class access specifier
- GE-5579 Cray XC integration needs to support multiple Crays in a cluster out-of-the box
- GE-5604 install\_cuda\_complexes.sh produces invalid complex
- GE-5675 Fully integrate Univa Grid Engine with systemd
- GE-5726 licence\_constraints in UGE clusters are not updated
- GE-5849 new masterl switch s\_rss limit setting is enforced for slave tasks
- GE-6025 create hash tables on CULL lists only on demand
- GE-6046 infinite loop writing to trace file in docker container
- GE-6049 default job class is not used correctly
- GE-6051 do docker run option parsing and checking in sge\_qmaster instead of on client side
- GE-6056 Cray XC integration needs to set unlimited timeout for epilog for newer UGE version
- GE-6058 Support for Ubuntu 16.04
- GE-6059 job gets rescheduled when epilog gets a SIGABRT signal
- GE-6063 a soft consumable extending a hard request might not get granted despite enough resources available
- GE-6064 on native Windows (win-x86), error handling of (Un)LoadUserProfile() overwrites real error message
- GE-6068 Interactive Docker jobs fail because sge\_container\_shepherd doesn't have the permission to "write the shepherd\_about\_to\_exit" file
- GE-6070 DRMAA2 submission crashes qmaster
- GE-6073 qsub does not support multiple use of "--binding" or "-t" parameters
- GE-6087 error message complains about missing Qmaster port when Execd port is missing
- GE-6095 log job verification time exceeding a certain threshold
- GE-6097 log request processing exceeding a certain threshold
- GE-6098 RC script fails to install when the OS is using LSB

- GE-6099 RC script not being uninstalled when update-rc.d is used as RCFILE
- GE-6100 the name of the named pipe between UGE Job Starter Service and SGE\_Starter.exe is not always unique
- GE-6109 dbwriter log is showing invalid integer value for a\_ioops field
- GE-6111 sge\_qmaster crashes when deleting an advance reservation with a partially finished array job
- GE-6112 ar\_attr and ar\_log records in reporting have incorrect record time
- GE-6113 the SGE\_Starter on win-x86 logs always, even if /log was not specified
- GE-6117 dbwriter / arcodb - error parsing the reporting file using MySQL backend
- GE-6118 the UGE Job Starter service on native Windows (win-x86) sometimes crashes when starting qrsh with command jobs
- GE-6148 uge\_js\_service.exe crashes on native Windows (win-x86) if the slave task of a tightly integrated parallel job finishes
- GE-6149 it is not possible to submit the slave task of a tightly parallel job from native Windows (win-x86)
- UWSA-140 Job Classes do not work as expected
- UWSA-150 Jobs QueueReq field not set
- UWSA-152 ugerest script fails in line 102 due to invalid comparison
- UWSA-154 setfileperm.sh does no longer setuid bit for authuser
- UWSA-156 Job Context Variables no longer returned
- UWSA-158 ugerest service is not starting if the key and keystore passwords are different

#### Univa Grid Engine 8.4.3

- GE-6152 log spooling exceeding a certain threshold
- GE-6156 execd crashes when a docker job is submitted, requesting memory resources

#### Univa Grid Engine 8.4.4

- GE-4425 SGE\_LONG\_QNAMES=-1 lead to qstat seqfault
- GE-5524 newline in job submission breaks reporting/accounting lines and qstat -j
- GE-6067 adding a queue with calendar generates error messages and no jobs will be scheduled into this queue
- GE-6126 submitting a docker autostart job fails if no job name is given
- GE-6163 erros for sge\_mirror\_process\_events() result in scheduler inactivity until it times out after 600 seconds
- GE-6167 add new scheduler profiling line for wait times
- GE-6173 Description of certain states for qstat -s is missing in help output and man page
- GE-6176 qdel -f prints unexpected messages, e.g. debug information
- GE-6181 automated install still using systemd initscript even when it is turned off in the config file
- GE-6185 job start fails when Docker daemon is busy
- GE-6186 the sge\_container\_shepherd fails to distinguish between yet unhandled IJS control messages and container related messages
- GE-6188 Windows (win-x86) interactive uninstaller should default to remove the services

- GE-6193 mirror interface reports unexpected error  
"callback function ("callback\_default") for event MOD QUEUE INSTANCE failed
- GE-6195 startup scripts are being installed during the automated install even if ADD\_TO\_RC=false
- GE-6200 child shepherd of builtin interactive job waits infinitely for mutex in shepherd\_trace()
- GE-6203 communication to Docker daemons built with Go (golang) > 1.6 does not work
- GE-6210 Add Intel KNL support
- GE-6211 the provided Docker container name is overwritten by automatically generated one
- GE-6212 automatic installation fails on OpenIndiana if the -auto switch is not the first in the command line

## 4 Upgrade Notes

### 4.1 Upgrade Requirements

This is a summary of the Upgrade Matrix that describes how you can carry out the transition from Sun or Oracle Grid Engine 6.2uX, Univa Grid Engine 8.0.X, 8.1.X, 8.2.X, 8.3.X to Univa Grid Engine 8.4 when you are currently using classic, BDB local spooling or PostgreSQL spooling. If the current version of Grid Engine you are using is missing in the overview, then please look at the full Upgrade Matrix located in the section Updating Univa Grid Engine in the Installation Guide.

Version	Upgrade Method
Univa Grid Engine 8.3.X	Backup/Restore
Univa Grid Engine 8.2.X	Backup/Restore
Univa Grid Engine 8.1.X	Backup/Restore
Univa Grid Engine 8.0.X	Backup/Restore
Oracle Grid Engine 6.2u6-6.2u8	Backup/Restore
Sun Grid Engine 6.2u5	Backup/Restore
Sun Grid Engine 6.2u1-6.2u4	Upgrade to SGE 6.2u5 and then Backup/Restore
Sun Grid Engine 6.2 FCS	Upgrade to SGE 6.2u5 and then Backup/Restore

Table 2: Upgrading from SGE, OGE, UGE 8.1.X, UGE 8.2.X to Univa Grid Engine 8.3.X

Upgrading to Univa Grid Engine requires a drained cluster, which means: No pending, running, ... jobs are allowed.

## 5 **Compatibility Notes**

## 6 **Known Issues and Limitations**