

Table of Contents

1 Release Notes Univa Grid Engine 8.1.....	1
2 License.....	2
3 Fixes and Enhancements.....	6
3.1 Summary.....	6
3.2 List of Fixes and Enhancements.....	7
3.2.1 Univa Grid Engine 8.1.0.....	7
3.2.2 Univa Grid Engine 8.1.1.....	12
3.2.3 Univa Grid Engine 8.1.2.....	12
3.2.4 Univa Grid Engine 8.1.3.....	13
3.2.5 Univa Grid Engine 8.1.4.....	13
3.2.6 Univa Grid Engine 8.1.5.....	14
3.2.7 Univa Grid Engine 8.1.6.....	15
3.2.8 Univa Grid Engine 8.1.7.....	16
4 Supported Platforms and Upgrade Notes.....	18
4.1 Supported Platforms.....	18
4.2 Upgrade Requirements.....	18
5 Known Issues and Limitations.....	20
5.1 cgroups Support after an Hot-Upgrade.....	20
5.2 Slotwise Preemption.....	20
5.3 Saving job related files.....	20

1 Release Notes Univa Grid Engine 8.1

2 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 means 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 with 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 right, 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 immediately 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

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) lost 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's 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)(l)(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 which Univa is bound to provide. We are hereby notifying you of these licenses.

Unicloud Kits

- Third Party Software means 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 means 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 means 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 means 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: March 2013

3 Fixes and Enhancements

3.1 Summary

Univa Grid Engine v8.1 support is available from <http://univa.com/products/grid-engine.php>

Here is a summary of things that have changed since version 8.0.1

- Introduced a new configuration object: Job Classes. They allow to
 - ◆ specify job templates that can be used to create new jobs.
 - ◆ reduce the learning curve for users submitting jobs.
 - ◆ avoid errors during the job submission or jobs which may not fit site requirements.
 - ◆ ease the cluster management for system administrators.
 - ◆ provide more control to the administrator for ensuring jobs are in line with the cluster set-up.
 - ◆ define defaults for all jobs that are submitted into a cluster.
 - ◆ improve the performance of the scheduler component and thereby the throughput in the cluster.
- Due to the Job Class enhancements qstat output has slightly been changed compared to 8.0.1. The qstat -ext|-urg|-pri shows an additional column with the name of the job class variant a job might have been derived from. Also qstat -j <jid> shows this information as well as the corresponding XML output when the switches are used in combination with the -xml switch.
- The processors attribute of the queue configuration has been removed.
- Moved decision about core binding from execution host to scheduler in order to guarantee a binding and make a better host selection.
- Core binding request on command line for parallel job is now a per host request (instead of a per qsh -inherit request before). Core binding is now better supported for parallel jobs, i.e. when submitting with "-binding pe ..." the pe_hostfile contains now different core binding decisions for different hosts. Before only the core selection of the master host is used for all slave nodes.
- Added a new complex type RSMAP, which allows to set a set of strings as resources. The job is mapped to the selected strings. The selected strings are available for the job through an environment variable (SGE_HGR_<rsmap>). qstat -j output shows the selected strings as well. RSMAP do support all kinds of jobs: resource reservation, parallel jobs, and array jobs.
- Added NUMA scheduling capabilities, which respects memory per NUMA node (using -mbind with -l m_mem_free) and sets memory allocation modes on lx-amd64 hosts.
- Added new resource complexes:

m_cache_l1	mcache1	MEMORY	<=	YES	NO	0	0
m_cache_l2	mcache2	MEMORY	<=	YES	NO	0	0
m_cache_l3	mcache3	MEMORY	<=	YES	NO	0	0
m_mem_free	mfree	MEMORY	<=	YES	YES	0	0
m_mem_free_n0	mfree0	MEMORY	<=	YES	YES	0	0
m_mem_free_n1	mfree1	MEMORY	<=	YES	YES	0	0
m_mem_free_n2	mfree2	MEMORY	<=	YES	YES	0	0
m_mem_free_n3	mfree3	MEMORY	<=	YES	YES	0	0
m_mem_total	mtotal	MEMORY	<=	YES	YES	0	0

m_mem_total_n0	mmem0	MEMORY	<=	YES	YES	0	0
m_mem_total_n1	mmem1	MEMORY	<=	YES	YES	0	0
m_mem_total_n2	mmem2	MEMORY	<=	YES	YES	0	0
m_mem_total_n3	mmem3	MEMORY	<=	YES	YES	0	0
m_mem_used	mused	MEMORY	>=	YES	NO	0	0
m_mem_used_n0	mused0	MEMORY	>=	YES	NO	0	0
m_mem_used_n1	mused1	MEMORY	>=	YES	NO	0	0
m_mem_used_n2	mused2	MEMORY	>=	YES	NO	0	0
m_mem_used_n3	mused3	MEMORY	>=	YES	NO	0	0
m_numa_nodes	nodes	INT	<=	YES	NO	0	0
m_topology_inuse	utopo	RESTRING	==	YES	NO	NONE	0
m_topology_numa	unuma	RESTRING	==	YES	NO	NONE	0

- The resource complex "m_mem_free" is a consumable as well as a reported load value. The scheduler takes the minimum for scheduling decisions into account. It is initialized in the complex_values field of the host configuration. When using with core binding it automatically turns into "m_mem_free_nX" requests (memory per socket (NUMA node) requests) depending which cores the scheduler had chosen. These additional requests are called implicit requests.
- The qstat -j output contains now implicit requests, which are displayed per host.
- The qstat -j output for "binding" changed from a topology based output (e.g. "SCCcc") to a numerical output (e.g. 0,2:0,3). It shows all bindings for all hosts in cases of parallel jobs.
- New scheduler parameter: PE_SORT_ORDER which could be ASCENDING or DESCENDING. This determines the order in which PE's are traversed.
- New scheduler parameter: PREFER_SOFT_REQUESTS which can be true or false (default). If true fulfilling soft requests is more important than scheduling time in case of reservation.
- New fair urgency policy used to achieve an even distribution of jobs on resources (scheduler configuration attribute fair_urgency_list).
- New sge_qmaster spooling method to a PostgreSQL database as an alternative to Berkeley DB spooling on NFS4.
- Improved behaviour of the -masterq switch. The -masterq request is now always fulfilled. If it contradicts the allocation_rule of the parallel environment, the allocation_rule is obeyed and a further task is added automatically to fulfill the -masterq request.
- Added templates for out of the box tight integration of the most common MPI implementations.
- The execd-parameter KEEP_ACTIVE was extended by the options "ERROR" and "ALWAYS". If set to ERROR, all job relevant log files will be sent to \$SGE_ROOT/\$SGE_CELL/faulty_jobs/\$job_id if the jobs fails. If set to ALWAYS, the log files of all jobs will be sent. KEEP_ACTIVE=ERROR is set on every default installation.

3.2 List of Fixes and Enhancements

3.2.1 Univa Grid Engine 8.1.0

GE-1412 additional pseudo variable \$sge_root for pe definition
 GE-1926 no info messages in execd messages file on aix
 GE-2418 qrsh fails with 'connection refused' error message

GE-2601 multiple occurrence of same compile parameters in aimk
 GE-2603 qsub option -q breaks -masterq
 GE-2643 accounting and online usage of jobs are wrong on aix
 GE-2841 submit(1) man page reports that qrsh does not support -display option
 GE-3132 Job validation behaviour changed since 6.0 / 6.1
 GE-3214 manpage queue_conf does not fully describe 'slots' notation
 GE-3265 array jobs with PE and dependencies killing qmaster
 GE-3299 On Windows Vista Enterprise, sgeexecd can fail to start up at boot time
 GE-3302 net continue SGE_Helper_Service.exe STOPS the service
 GE-3304 no accounting information for Windows GUI jobs
 GE-3354 Cache sizes and cache topology should be reported by GE execution hosts per default
 GE-3363 new spooling method writing data into a relational database
 GE-3364 evaluate / fix / improve the spooling performance tests
 GE-3365 create a prototype for database spooling using postgres
 GE-3373 Create a default parallel environment for OpenMP jobs (pe_slots), which is available right after installation
 GE-3386 Multi-core NUMA awareness and binding.
 GE-3390 qrsh does not forward necessary environment variables
 GE-3403 Add support to Grid Engine for GPUs
 GE-3409 Out of Box support for MPI Libraries - likely OpenMPI
 GE-3414 sge_execd sometimes hangs during daemonization
 GE-3440 file descriptor -1 passed to system calls in interactive job support
 GE-3441 change default shell_start_mode from posix_compliant to unix_behavior in global config and in default queues (all.q)
 GE-3456 adding a new complex type "resource map" RSMAP
 GE-3474 PDC_INTERVAL=NEVER does not work
 GE-3480 qstat -xml -j output changes after a job is altered with qalter -l
 GE-3483 Within JSV it is not possible to distinguish if -v or -V was used
 GE-3484 submit client and host is not available in JSV
 GE-3498 remove reporting_param log_consumables
 GE-3503 on Windows, the loadcheck.exe binary output misses a line break
 GE-3505 fstype binary doesn't detect NFS4 on Linux
 GE-3509 qacct segfaults when bootstrap file was not found
 GE-3512 gdi_retries option shall also have effect on sending gdi requests
 GE-3514 Pass data as part of GDI return value
 GE-3546 adjust JC from JSV
 GE-3575 qmaster can't read spooled jobs after a hot upgrade to version 8.0.0
 GE-3583 Cleanup: Move sgeijs-lib in clients/common
 GE-3591 sge_shepherd might not deliver a signal because "remaining_alarm" might be 0
 GE-3595 output np_load_avg instead of load_avg in qstat -f and qhost
 GE-3603 RSMAP complex must be compatible with a per JOB consumable
 GE-3604 Qconf man pages have wrong object_spec for Resource Quotas
 GE-3607 62u5 clients causes a segmentation fault of a 8.0.0 qmaster
 GE-3609 CMDNAME needs to be documented in the wiki
 GE-3611 adding qstat resource map output to qstat -xml output
 GE-3612 add new version 8.1.0alpha1
 GE-3620 adding memory per NUMA node reporting
 GE-3623 Add all inherited environment variables of qrsh to documentation and man-page
 GE-3627 Limit number of multi GDI get requests in qmaster
 GE-3628 Make it possible to disable sending of environment variables in combination with qstat -j requests
 GE-3629 qstat -j * should show only own jobs per default
 GE-3630 documentation shall contain a matrix listing the rights of user/operator/manager for the different commands
 GE-3637 wrong header file might be included when GE is compiled with UGE extensions

GE-3638 ship the CUDA load sensor with Grid Engine
 GE-3640 GUI- and text installer both don't show the trial license when they are compiled in trial mode
 GE-3642 Make it possible to extract source code documentation with aimk
 GE-3643 queue/job error states should be explained in more detail
 GE-3648 About-Dialog of qmon is not readable
 GE-3651 Better description of master and slave queues in Parallel Environments.
 GE-3656 add description of UGE_Starter_Service.exe to documentation
 GE-3659 add a way to control error state behavior of DRMAA jobs
 GE-3673 fix Insure errors reported for IJS in 8.0.1alpha2
 GE-3674 make qrsh without command work with Insure instrumented sge_shepherd, qrsh_starter, qsh
 GE-3682 add support for kernel 3.0
 GE-3686 scheduler monitoring enable/disable is not clear
 GE-3695 sharetree man page doesn't explain internal nodes
 GE-3708 bundle a script which adds the complexes reported by the CUDA load sensor to the complex configuration
 GE-3713 qrsh (without command) fails if something is entered in qrsh client during job start
 GE-3715 online job usage is not reported on AIX
 GE-3718 on Windows, if run as the local Administrator, qstat prints 'invalid user name "Administrator"'
 GE-3722 qsub -sync y and drmaa clients on AIX cannot connect to sge_qmaster
 GE-3724 FD are close to fast which leads to problems in combination with nsswitch module from BeyondTrust
 GE-3731 Java DRMAA Error : can't send response for this message id - protocol error
 GE-3743 prevent sge_execd to crash when "/" is not a directory and in out of memory scenarios
 GE-3746 pe array jobs put queues in error state with a file not found error for the job script
 GE-3748 scheduler has to make the decision about core selection in case of core binding
 GE-3755 tightly integrated parallel jobs are not correctly handled when suspending/unsuspending queues
 GE-3756 qalter -verify is not shown in qalter -help output
 GE-3757 reduce impact of qstat -j "*" on qmaster in clusters with many jobs
 GE-3758 pe tasks of tightly integrated parallel jobs are never suspended
 GE-3760 "qsub -pty yes -b y tty" exits with an exit code 3
 GE-3761 security hole in UGE when setting LD_PRELOAD or LD_LIBRARY path
 GE-3763 enhance loadcheck -cb with memory binding capability checks and memory binding testing
 GE-3764 shepherd consumes 100% CPU if IJS does not use builtin as starter method
 GE-3765 documentation of XFILESEARCHPATH in qmon man page and Qmon config file missing
 GE-3767 the output of qsub -pty yes jobs is not written to the jobs output file
 GE-3768 a pty is created for the pre- and post-commands of any -pty yes job
 GE-3773 qrsh -pty yes fails if invoked within a qsub-job
 GE-3775 UGE crash with error in qmaster message: got NULL element for JB_type
 GE-3781 calling JSV "jsv_set_param binding_exp_n 0" twice segfaults qsub or qmaster
 GE-3783 Missing field in mpich template
 GE-3785 bash shell functions are not properly transferred in the environment
 GE-3790 after job end sge_shepherd processes stay running
 GE-3792 auto installation fails if EXEC_HOST_LIST points to a file containing host names
 GE-3793 the scheduler thread can be stopped by a normal user
 GE-3810 resource reservation does not work correctly with serial jobs
 GE-3812 user can be added to multiple departments but it should be denied

GE-3829 importance of soft requests in resource reservations should be higher than earlier start time (make it configurable)
 GE-3831 need a "fair urgency" policy
 GE-3832 Execd init script does not stop sge_execd daemon on non MacOS hosts
 GE-3839 getJobProgramStatus call in drmaa is throwing DRMCommunicationException
 GE-3841 possible buffer overflow in command line parsing of sgepasswd
 GE-3846 builtin qrsh <command> fails if data still has to be transferred after job end
 GE-3847 the 48 core limit in the test binaries should apply only to running execution hosts, not to all that are registered at the QMaster
 GE-3848 qstat -fjc
 GE-3849 document required mount options for the spooling shared file system
 GE-3857 People allowed to create AR's is not documented.
 GE-3860 parallel environment selection order in case of wildcards pes should be configurable
 GE-3863 host_alias file is not covered in documentation
 GE-3866 qsub -w p \$SGE_ROOT/examples/jobs/sleeper.sh is crashing
 GE-3887 not helpfull error message in qmaster messages file, regarding sharetree
 GE-3900 both man page and documentation must explain how the new -masterq switch works
 GE-3903 sge_aliases file is not backed up and restored
 GE-3904 man page sge_conf is unclear about the meaning of projects/xprojects
 GE-3905 Reconcile the Archimedes Hadoop docs with the current Grid Engine Hadoop docs
 GE-3907 slave tasks are wrongly scheduled to the master queue
 GE-3918 qalter man page: wrong tag is used for italic text
 GE-3921 qmod man page should mention that only the master task of a parallel job gets suspended
 GE-3923 GUI Installer must support Postgres spooling method
 GE-3924 if "--soft -q <queue>" is specified, "-masterq" doesn't prevent slave tasks from being scheduled to the master queue
 GE-3925 Support chkconfig Tool for RHEL Variants
 GE-3926 qstat -j "*" -xml is crashing
 GE-3932 can delete a complex attribute which is referenced in the load_formula by shortcut
 GE-3943 'long_term_usage' in the 'users' spool object is not spooled when QMaster is shut down
 GE-3944 sgeexecd script returns 0 even if it was not able to stop the execution daemon because of missing permissions
 GE-3949 Extend KEEP_ACTIVE execd param
 GE-3950 the fstype binary doesn't report the right file system types anymore
 GE-3952 test_sge_lock_fifo test binary throws a floating point exception from time to time
 GE-3954 'job duration is longer than duration of AR' message is no longer printed if requested job runtime exceeds AR duration
 GE-3955 the caller of qalter, not the job owner, is checked against the ACL of a project
 GE-3956 the argument of the '-r' option of qsub is ignored, if '-r' is provided, it is assumed it is always set to 'yes'.
 GE-3957 remove the warning "Job Done" in the execd messages file
 GE-3958 qresub of a job in hold state doesn't clear the hold state for the new job
 GE-3959 make Berkeley db spooling platform independent
 GE-3960 update openssl used for CSP mode to openssl-1.0.0
 GE-3961 update Berkeley DB to current version 5.3.21
 GE-3962 job hold doesn't ignore jobs that are not currently in the system
 GE-3963 the halflife_decay_list parser doesn't reject invalid separators
 GE-3966 qmon works only partly on lx-amd64
 GE-3971 qsub option "--shell <y[es]|n[o]>" is ignored
 GE-3972 Releasing jobs with qrls -u <username> aborts all held jobs of this user

GE-3973 failing DRMAA jobs are always set to error state, even if
 SGE_DRMAA_ENABLE_ERROR_STATE is not set
 GE-3974 Remove 'Deprecated' Message on SHARETREE_RESERVED_USAGE option in sge_conf
 GE-3975 Remove Hadoop integration from default install
 GE-3978 qalter -cwd does not change the working directory of a job
 GE-3979 qalter -c does not change the time when a job should be checkpointed
 GE-3980 the sge_conf(5) man page doesn't mention the key word "infinity" for
 resource limits
 GE-3982 adding support for different core binding decisions (selected processors)
 to PE hostfile
 GE-3984 for PE jobs different core bindings on different hosts should be displayed
 in the qstat output
 GE-3988 qalter prints error message but alters the job
 GE-3991 SUBMIT_HOST parameter in JSV is missing
 GE-3992 on AIX, win and 32 bit linux, qstat prints out error message:
 PE_RANGE_ALG=bin is not a valid parameter, qconf is not working
 GE-3993 Interactive jobs have no name
 GE-3994 cannot start a 32bit drmaa application on Solaris when CSP mode is installed
 GE-3998 it is not possible to change mbind with qalter
 GE-4000 abort of qmaster when job derived from a JC changes PE request
 GE-4001 qalter -pe shows incorrect message in case of success
 GE-4002 classic spooling of job fails sometimes when pe name or range of a parallel
 job is changed
 GE-4004 csp mode is broken on multiple platforms like aix51, hp11-ia64
 GE-4006 DISPLAY variable is send as part of the 'full environment' although -V is
 not used
 GE-4014 Set KEEP_ACTIVE=ERROR as default
 GE-4015 adding sections about RSMAP and mbind in man pages
 GE-4017 qconf -mc always shows message about m_topology_numa
 GE-4018 fair urgency policy is broken after some time
 GE-4025 qsub/qalter -mods is not able to change CMDARG elements of a job
 GE-4029 qstat -xml -j <not_existing_jid> prints invalid XML
 GE-4031 when a tightly integrated job is rescheduled with qmod -rj the slave tasks
 are not signalled
 GE-4032 qmaster crashes when a parallel job requests soft resources
 GE-4036 adding a execution daemon parameter in order to disable reporting of
 m_mem_free
 GE-4038 backup/restore mechanism of inst_sge does not restore all files
 GE-4039 sge_execd silently ignores duplicate job delivery
 GE-4041 restore with spooling method postgres should clear the database
 GE-4044 make the cuda loadsensor Makefile available in the load sensor directory
 GE-4051 wrong message when resubmitting a job with hold
 GE-4052 test_enumeration binary crashes because it calls lFreeDescr() with NULL
 pointer
 GE-4053 setting SGE_BINDING variable in any case of binding
 GE-4056 epilog with exit status 1 doesn't set queue in error state
 GE-4057 adding SGE 6.5u5 compatibility wrapper script for qhost which transforms
 non existent "-cb" switch
 GE-4060 qalter man page is lacking description of the -clearp option
 GE-4061 'qrsh -inherit -v ENVVAR=value' doesn't transfer the value to the execution
 daemon
 GE-4062 It is possible to add not existing queue references to a job via
 -soft -q of qalter/qsub.
 GE-4063 qalter -mods q_soft/q_hard does not change the specified queue name
 GE-4064 qresub as non deadline user fails
 GE-4065 Deadline and start time of a job cannot be cleared with qsub/qalter -clearp
 GE-4070 restore (inst_sge -rst) is failing with Postgres spooling to a remote
 database server
 GE-4072 qstat -xml output contains invalid tag <context list> when job context is

changed with qalter.

GE-4073 qstat -xml output might contain a tag without a name if the list of mail recipients is changed

GE-4078 add scripts which create a OpenMPI rankfile out of the pe_hostfile

GE-4081 low level spooling info or error messages may not be propagated to clients (e.g. qsub)

GE-4086 the qmaster generates wrong m_mem_used complex value if it doesn't exist

GE-4088 parallel job which could run is not started when using job exclusive complex

GE-4092 job classes bug when submitting binary job

GE-4094 when qmaster is started on command line in debug mode unnecessary information is printed out at the beginning

GE-4096 provide a script for reformatting an accounting file to reporting file

GE-4098 misleading error message when initializing the Postgres database during installation

GE-4104 support non standard port for postgres spooling

GE-4107 user_lists in job classes does not check existence of ACL if this ACL is preceded by user entry

GE-4108 strange error message when showing job class variant with qconf -sjc

GE-4117 free and used memory per NUMA node is wrong for multi-socket hosts

3.2.2 Univa Grid Engine 8.1.1

TS-531 create testsuite test for GE-4122

TS-501 jsv_ge_mod test is broken in rulelevel 1

GE-4162 automatically create per socket memory consumables (m_mem_free_nX) for more than 4 sockets as soon as bigger exec hosts connect to master

GE-4156 -mbind nlocal with parallel jobs does not decrement per socket memory correctly

GE-4152 GUI installer has problems with installation in case of BDB spooling

GE-4148 qalter using job names instead of JID's to identify jobs cause job names to be reset to (null)

GE-4144 very short qrsh job sometimes seems to fail

GE-4133 the qrsh client hangs for parallel jobs if consumables are requested

GE-4129 past usage information of tightly integrated jobs can get lost at qmaster restart with postgres spooling

GE-4121 qrsh sometimes crashes at job end

GE-4113 qsub -help output shows -pty not in alphabetic order

GE-3919 GE-3386 adding a possibility to express that amount of cores can also be amount of threads depending on the selected execution host

GE-3769 GE-3386 provide a technical paper which discusses all memory requests UGE offer for sequential and parallel jobs

Fixed with the first patch (UGE 8.1.1p1)

GE-4199 qstat -j does not return with an error code if no job can be selected according to a specified job name or pattern

3.2.3 Univa Grid Engine 8.1.2

GE-4198 qmon crash if user tries to submit an array job

GE-4196 qmon opens a error message dialog at job submission -> job submission is not possible

GE-4195 The JB_tgt job attribute is cleared when using the AFS security model

GE-4172 qmaster crash if qsub job class is specified with a binary submission but without -b y switch

GE-4151 add option for triggering a scheduling run without writing the schedd_runlog file

GE-4150 too slow reservation scheduling

GE-3776 jobs submitted from an old qsub binary are accepted and might crash qmaster

3.2.4 Univa Grid Engine 8.1.3

GE-4272 wrong DRMAA errno after session reconnect and status check of an array job which is in hold
GE-4266 tightly integrated jobs do not set SGE_HGR_ for slave tasks correctly
GE-4244 KEEP_ACTIVE=ERROR does not transfer all files and doesn't transfer files to the right location
GE-4243 qmaster crashes when KEEP_ACTIVE=ERROR and the path to \$SGE_ROOT/\$SGE_CELL/faulty_jobs is very short
GE-4231 Non-LSB RedHat Installations Fail Init Service Install
GE-4226 RSMAP id is not freed in certain exceptional error conditions
GE-4211 qmaster becomes unresponsive when submitting a large job net
GE-4201 event client id's are not always reused
GE-4199 qstat -j does not return with an error code if no job can be selected according to a specified job name or pattern
GE-4194 improve scheduling performance with exclusive complexes
GE-4147 GE-4135 enhance RMSAP consumable so that it can be a per HOST consumable
GE-4143 qstat -u @unix_group does not work
GE-4137 GE-4135 create an installation script for Intel Phi load sensor complexes
GE-4136 GE-4135 create a load sensor which reports metrics about Intel Phi (MIC) co-processor cards
GE-4101 resource reservation for a (job) exclusive consumable does not work
GE-4058 The KEEP_ACTIVE=ERROR functionality does not always transfer the generated_job_script
GE-4023 GE-4135 adding support for binding a job near a specific (by the scheduler granted) PCIe device for better I/O rate
GE-3433 Add time stamp to debug output
GE-3313 on Windows hosts reporting their full qualified hostname, certificates got wrong names
GE-2894 SGE 6.2u2 Beta execd install failed on Windows 2008 server
GE-2842 listener threads get stuck in cl_complib_receive_message

3.2.5 Univa Grid Engine 8.1.4

GE-4155 adding a facility which transforms between different CPU core representations
GE-4255 add a tool (mic_load_check) which allows to access load values of Intel Xeon Phi cards within jobs
GE-336 h_vmem and shared memory segments
GE-2704 on Windows, the installation should check if the installing user is really the Administrator
GE-2907 execd intallation on Windows doesn't provide global spool dir as default, even if it's local
GE-3487 missing index for sge_job table on j_job_number, j_task_number, j_pe_task_id
GE-3709 SGE_BINDING variable contains just one ID in case of hyperthreading
GE-3836 on Windows 7, the installer tries to use the mapadmin.exe binary which doesn't exist there
GE-3844 execd crash immediately after startup when reading existing active_jobs dirs
GE-3985 You cannot qalter the h_rt limit of a job that has run then exited with 100 to enter the error state with the expected outcome that the new h_rt will be used when the job is re-run
GE-4273 m_topology_inuse is not adapted in case the execution topology is changed
GE-4278 accounting data is copied when doing an upgrade from 8.0 to 8.1 but qacct is not able to parse it
GE-4282 allow qrsh to use new -mbind parameter
GE-4289 Submitting user and host are not automatically used as mail address (-M) for

email notifications.

GE-4290 Core of sge_qmaster if m_topology is not reported by a sge_execd anymore

GE-4291 Cluster is not filled completely because scheduler is skipping remaining orders

GE-4299 execd profiling is broken and (if working) needs improvement

GE-4311 logchecker.sh generates a error message if execd_spool_dir is set to none

GE-4318 qstat output order is changed after qmaster restart

GE-4320 qrsh client waiting for job scheduled can unexpectedly abort while waiting for job execution

GE-4321 sge_resource_quota man page shows strange characters in section jcs

GE-4330 jsv_add_env doesn't work any more in server-side jsv in version 8.1

GE-4332 QMON does not support RSMAP

GE-4333 distinst fails due to phisource (though exit code might still be 0 with "-noexit" flag)

GE-4335 SGE_HGR_ environment variable must show consumable values in an appropriate format

GE-4347 Job sequence number roll-over mechanism not working under special conditions

GE-4351 wrong error message and exit code for qconf -dconf <non existing config>

GE-4366 qstat -j error reason output is truncated

GE-4369 AFS token problem after qmaster failover

GE-4380 sge_execd exits when files cannot get written at job start

GE-4386 sge_qmaster does not recognize set filedescriptor limit

GE-4388 JSV script does reject core binding settings for linear and striding core allocation

GE-4396 creating a wrong job class leads to qmaster crash

GE-4398 Temporary name server problems can produce resolve errors on client commands

GE-4423 qmaster abort with "Removing element from other list !!!"

GE-4426 client command fails with reresolve hostname error

GE-4435 qmon jobs showing an empty line in qstat -j output

GE-4306 document format of execd spool directory and explain the files that can be found there

GE-3486 on Solaris the reported amount of hardware supported threads must be calculated

GE-4066 enhance qstat -j output to show current wallclock and cpu time when SHARETREE_RESERVED_USAGE is set

GE-4310 m_mem_free should not exclude cached memory

GE-4360 optionally break schedulers dispatching loop after a configurable time interval

GE-4370 add an option for limiting the number of jobs getting dispatched in one scheduling run

GE-4395 adding JSV support for the mbind submission parameter

GE-4434 RSMAP value can't be initialized in the host configuration of qmon

GE-4308 allow for the active "decreasing" of a m_mem_free memory request with qalter

3.2.6 Univa Grid Engine 8.1.5

GE-3074 end job with SIGINT instead of SIGKILL if qrsh <command> was quit with Ctrl-C

GE-4233 Berkeley DB utility html pages are missing from \$SGE_ROOT/doc

GE-4268 Allow deletion of Complexes that have load_values present

GE-4350 qconf -aconf called on a non cluster host reports misleading error message

GE-4390 incorrect scheduling decision with pe allocation rule \$fill_up

GE-4445 possible hostname resolving problems for installations at sides having large amount of hosts

GE-4459 Invalid parsing of AR durations longer than 1 day

GE-4461 re-resolving hosts after reading in the host_aliases file should be done with standard host name for the current machine

GE-4465 Client side jsv script gets wrong job priority if -p switch is not used

GE-4471 monitor.sh script (dtrace) fails with "probe description :::BEGIN does not match any probes"

GE-4473 builtin interactive job may not terminate on pty loss
 GE-4483 prepared statements in postgres spooling are lost after server restart
 GE-4484 job context variables and environment variables passed via -v variable=value cannot contain commas
 GE-4486 MAX_DISPATCHED_JOBS can disable a cluster in some cases
 GE-4492 wrong error message is logged in spooling if hostname can't be resolved
 GE-4499 When using enable_mem_details the execd process can exit with STDERR message "input in flex scanner failed"
 GE-4502 prefer launchd for UGE daemon installation on darwin-x64
 GE-4505 scheduler is sending orders for changing user or project objects even if no change was done -> triggered by spooling outage
 GE-4516 qstat -j without job id shows wrong output in case a job was rejected through insufficient amount of cores/sockets
 GE-4523 PE slot debiting broken
 GE-4527 Advanced Reservations stop reserving global host consumables after global exec host updates
 GE-4528 Postgres library is missing in darwin-x64 package
 GE-4533 when a tightly integrated job is rescheduled with qmod -rj the old master task is not killed
 GE-4537 tightly integrated job fails as sge_shepherd loses connection to the qrsh -inherit
 GE-4553 qstat -xml -j <job_id> creates for parsers unknown output
 GE-4556 commlib debug information is written to stdout
 GE-4558 qmaster does not read pe tasks from spooling database with Berkeley DB spooling
 GE-4559 potential qrsh and sge_shepherd crash under high load
 GE-4566 builtin qrsh affected by race condition of communication system
 GE-4568 builtin qrsh -suspend_remote y affected by race condition
 GE-4569 qmaster crash when a invalid array size is entered at submission
 GE-4570 jobs with core binding are rejected when load sensor reports a load value into global host
 GE-4573 submitting a job with array task number ULONG_MAX causes scheduler endless loop
 GE-4263 Add environment variable to allow load sensor to find architecture-specific binaries
 GE-4387 use scheduler pe sort algorithm "bin" (binary search) as default value
 GE-4455 add qmaster param ENABLE_REDUCE_MEM_FREE to allow qalter to reduce mem_free consumable during job run-time
 GE-4506 Display qconf complex_values and load_values in alphabetic order
 GE-4576 improve logging of SGE Helper Service on Windows Vista and later
 GE-3727 add MacOS X launchd integration in order to get GE running on newer MacOS X
 GE-4068 allow dynamic job runtimes that are used when s_rt and h_rt cannot be specified for AR jobs and that can be adjusted later
 GE-4481 improve performance and reliability of builtin interactive job support
 GE-4589 Each slave PE task report transfers trace files of all PE tasks on the same host if KEEP_ACTIVE=ERROR

NOTE: Performing a Hot upgrade on 8.1.5 can only be done if the 8.1.X cluster is drained of all jobs before performing the hot upgrade. Please make sure you do a backup of your cluster before performing any upgrade.

3.2.7 Univa Grid Engine 8.1.6

GE-4662 Command line options from script file can break qrsh execution
 GE-4661 Command line options from sge_request can break client commands
 GE-4660 Resource quota for forced requestable complex not working with "--soft" job request
 GE-4657 Forced requestable complex not working with "--soft" request
 GE-4654 Added missing warning message if an execd cannot get notified to send full load report
 GE-4653 potential qstat and scheduler (sge_qmaster) crash with inconsistent spooling data

GE-4610 Wrong statement regarding "qstat -j" job visibility of jobs from other users
GE-4602 qmon crashes when doing qalter of job
GE-4601 Change in s_rt functionality
GE-4598 Modification of job context with qalter silently ignores variables starting with "-" or "+" letter
GE-4597 New sorted adding of elements into cull list with same value may show unexpected behaviour
GE-4596 qmaster becomes irresponsible with large dependency nets
GE-4593 qalter -ac core dump for invalid variable parameter if SGE_OLD_VARIABLE_PARSING env is set
GE-4588 scheduler thread may run out of qmaster event synchronisation
GE-3868 qalter -ac/-dc will crash qmaster if both options are used at the same time
GE-2914 classic spooling method (flatfile spooling) should be more error tolerant

NOTE: Performing a Hot upgrade on 8.1.6 can only be done if the 8.1.X cluster is drained of all jobs before performing the hot upgrade. Please make sure you do a backup of your cluster before performing any upgrade. If you are unsure about how to perform your upgrade please feel free and contact our support department and send a mail to support@univa.com or open a ticket in our support portal.

3.2.8 Univa Grid Engine 8.1.7

GE-2675 After installing the xterm is not found because the wrong path was set
GE-3429 use linux control groups (cgroups) for resource management on execd
GE-4188 qstat man page discusses invalid "-cp" option
GE-4507 Linux Cgroups support for Univa Grid Engine.
Add functionality to support Linux cgroups subsystems in Univa Grid Engine.
GE-4609 Bad paths printed in installer when installing UGE
GE-4615 Extend sync option of qsub/losub
GE-4618 Qmaster terminates when KEEP_ACTIVE is enabled and the faulty_jobs directory cannot be created
GE-4628 Installing the qmaster with postgres spooling fails if install-user == admin-user
GE-4630 builtin qrsh creates tty which is owned by root instead of job user
GE-4642 Installer prints error messages when using -start-all switch
GE-4646 client side JSV scripts are waiting too long (min. 1 sec) when they finish
GE-4651 The scripts 'install_qmaster' and 'install_execd' should only support command line options for qmaster or execd host installation
GE-4663 Install script (inst_sge) does not backup sge_request file in cell directory with -bup
GE-4665 adding Makefile to Intel Xeon Phi load sensor
GE-4667 make Phi load sensor more robust
GE-4669 SGE_HGR_ environment variable is not set under some circumstances for array tasks
GE-4674 Wrong online sharetree usage calculated for task array jobs
GE-4676 Check for duplicate queue instances when reading in configuration at sge_qmaster startup
GE-4677 Wrong booked sharetree usage for deleted task array jobs
GE-4681 wrong XML output in escape sequence for "<" character
GE-4682 lothread is failing when running a tight pe job, which leads to a qmaster crash
GE-4683 after qmaster restart m_topology_inuse is not reset and re-accounted correctly
GE-4684 add support for Mac OS Maverick
GE-4685 in rare conditions core accounting in m_topology_inuse is wrong
GE-4687 accounting is not complete after rescheduling a job
GE-4706 dbwriter installation prints "infotext: too few arguments"

GE-4716 SGE_BINDING environment variable does not appear on 32 bit Linux
GE-4718 The built-in PLPA core binding is not working on linux hosts
GE-4720 Fixed memory leak when reading in sgepasswd file
GE-4724 unexpected sge_qmaster daemon shutdown with exit state 100

NOTE: Performing a Hot upgrade on 8.1.7 can only be done if the 8.1.X cluster is drained of all jobs before performing the hot upgrade. Please make sure you do a backup of your cluster before performing any upgrade. If you are unsure about how to perform your upgrade please feel free and contact our support department and send a mail to support@univa.com or open a ticket in our support portal.

4 Supported Platforms and Upgrade Notes

4.1 Supported Platforms

Univa Grid Engine 8.1 supports various hardware architectures and versions of operating systems.

Supported Platforms, Operating Systems and Architectures

Operating System	Version	Architecture
SLES	10,11	x86, x86-64
RHEL	4-5.6, 6-6.3	x86, x86-64
CentOS	4-5.6, 6-6.3	x86, x86-64
Oracle Linux	4-5.6, 6-6.3	x86, x86-64
Ubuntu Server	10.04LTS-10.10	x86, x86-64
Microsoft Windows ¹	Vista, 7, Server 2008 R1, HPC Server 2003, XP SP3 (x86 only)	x86, x86-64
Oracle Solaris	9,10	x86_64
HP-UX	11.0 or higher	32 and 64bit
IBM AIX	5.3, 6.1 or later	64 bit
Apple OS X	10.8 (Mountain Lion) or higher	x86, x86-64
Ubuntu	12.10	ARMv7

¹ Hosts running the Microsoft Windows operations system cannot be used as master or shadow hosts.

4.2 Upgrade Requirements

This is a summary of the Upgrade Matrix that describes how you can make the transition from Sun or Oracle Grid Engine 6.2uX or Univa Grid Engine 8.0.X to Univa Grid Engine 8.1 when you currently use classic or BDB local spooling. If your current version of Grid Engine you are using is missing in the overview then please have a look into the full Upgrade Matrix located in the section [Updating Univa Grid Engine](#) of the Installation Guide.

Upgrading from SGE, OGE and UGE 8.0.X to UGE 8.1.X

Version	Upgrade Method
Sun Grid Engine 6.2u5	Backup/Restore
Sun Grid Engine 6.2u4	Upgrade to SGE 6.2u5 and then Backup/Restore
Sun Grid Engine 6.2u3	Upgrade to SGE 6.2u5 and then Backup/Restore
Sun Grid Engine 6.2u2	Upgrade to SGE 6.2u5 and then Backup/Restore

Sun Grid Engine 6.2u1	Upgrade to SGE 6.2u5 and then Backup/Restore
Sun Grid Engine 6.2 FCS	Upgrade to SGE 6.2u5 and then Backup/Restore
Oracle Grid Engine 6.2u6	Backup/Restore
Oracle Grid Engine 6.2u7	Backup/Restore
Univa Grid Engine 8.0.X	Backup/Restore

For any upgrade the backup/restore mechanism as outlined in the section Clone Configuration of the Installation Guide must be used to upgrade to Univa Grid Engine 8.1.

5 Known Issues and Limitations

5.1 cgroups Support after an Hot-Upgrade

After a fresh installation of UGE 8.1.7 a new field is added to the global configuration (qconf -mconf). This field ("cgroups_params") contains the configuration of cgroups feature. It is missing after an hot-upgrade (unpacking new packages in the installation folder), hence it must be added manually when cgroups feature is going to be used.

This is done by simply copying the configuration into new lines when having the global configuration open (qconf -mconf).

These are the default entries:

```
cgroups_params          cgroup_path=none cpuset=false mount=false \  
                        freezer=false freeze_pe_tasks=false killing=false \  
                        forced_numa=false h_vmem_limit=false \  
                        m_mem_free_hard=false m_mem_free_soft=false \  
                        min_memory_limit=0
```

5.2 Slotwise Preemption

If OGE 6.2u6 or OGE 6.2u7 was used with parallel jobs in slotwise preemption, be aware that Univa Grid Engine does not support this.

5.3 Saving job related files

If the "execd_params" parameter "KEEP_ACTIVE" is set to "ERROR" or "ALWAYS", the job finish of a larger parallel job could make the QMaster response slowly for several seconds or minutes. This depends highly on the number of tasks in the parallel job and on the performance of the storage the \$SGE_CELL directory is located on. If such a slow-down is observed and causes trouble, setting "KEEP_ACTIVE=FALSE" is recommended.

The setting "KEEP_ACTIVE=TRUE" does not affect the performance of the QMaster, but should be used in debugging sessions only.