



UNIVA CORPORATION

GRID ENGINE DOCUMENTATION

Grid Engine Release Notes

Author:
Univa Engineering

Version:
8.5.3

July 19, 2017

Contents

1	License	1
2	Supported Operating Systems, Versions and Architectures	6
3	Fixes and Enhancements	7
3.1	Summary	7
3.1.1	8.5.1: Changed limit calculation	7
3.1.2	8.5.1: Improved rescheduling behaviour	7
3.1.3	8.5.1: Possibility to reduce qhost data request sizes at sge_qmaster	7
3.1.4	8.5.1: New environment variables in the job environment	7
3.1.5	8.5.1: New example script for jsv and core-binding	8
3.1.6	8.5.1: sgepasswd renewal	8
3.1.7	Performance Improvements and Memory Requirements	9
3.1.8	Standing Reservations	9
3.1.9	Policy Scheme: Consider Slots Instead of Jobs	9
3.1.10	RSMAP Enhancements	10
3.1.11	Improved Scheduler Profiling	11
3.1.12	Improved Logging	11
3.1.13	Encryption in CSP mode / sgepasswd	11
3.1.14	Online usage of running Windows jobs	11
3.1.15	Docker Related Enhancements	13
3.1.16	Host Aliasing and Resolving	13
3.1.17	Intel® Xeon Phi™ x200 (Knights Landing) integration	14
3.2	Full List of Fixes and Enhancements	15
3.2.1	Univa Grid Engine 8.5.0alpha1 (also fixed for a 8.3 or 8.4 patch release)	15
3.2.2	Univa Grid Engine 8.5.0alpha1 (also fixed for a 8.4 patch release)	18
3.2.3	Univa Grid Engine 8.5.0alpha1	21
3.2.4	Univa Grid Engine 8.5.0alpha2	23
3.2.5	Univa Grid Engine 8.5.0beta1	24
3.2.6	Univa Grid Engine 8.5.0 FCS	25
3.2.7	Univa Grid Engine 8.5.1	25
3.2.8	Univa Grid Engine 8.5.2	26
3.2.9	Univa Grid Engine 8.5.3	27

4 Upgrade Notes	29
4.1 Upgrade Requirements	29
5 Compatibility Notes	30
5.1 Changes in Output Format of Commands	30
5.2 Changes in Windows execution host sgepasswd file	31
5.3 Deprecated Functionality	32
6 Known Issues and Limitations	33
6.1 setting halftime to -1 (GE-6497) not supported in qmon	33

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

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 8.5.1: Changed limit calculation

The limit calculation for jobs was improved for Univa Grid Engine 8.5.1 and changed compared to previous versions of Univa Grid Engine. The most important changes are:

- The configured consumable type (NO, YES, JOB, HOST) will not influence any resulting limit for tight integrated parallel jobs
- Previous Univa Grid Engine version calculated limits that were too high (depending on pe and consumable settings)
- This also affects the cgroups h_vmem observation

Univa Grid Engine versions prior to 8.5.1 had limit values that were set too high for limits like “h_vmem”. As a result jobs were not terminated although they exceeded one of their limits. The limit calculation is now corrected. A detailed overview of the limit observation and how it works is described in the `sge_diagnostics(1)` man page (JOB LIMITS).

ATTENTION: If you are updating to 8.5.1 from a previous version please verify the used limit requests of your jobs. It may be necessary to change the requested limit value. If the limit is set too low or adjusted to fit the old limit calculation, jobs which were running fine may fail after installing this version.

3.1.2 8.5.1: Improved rescheduling behaviour

The new `execd_params` parameter `RESCHEDULE_ON_MISSING_EPILOG` is introduced. The default value is `true`, which causes the old behaviour. If set to `false`, the job is not rescheduled and the queue not set to error state if the configured epilog script cannot be found. Instead, Univa Grid Engine behaves as if no epilog script was configured. This parameter also applies to the `stop_proc_args` scripts of a parallel environment (also called `pe_stop` script).

3.1.3 8.5.1: Possibility to reduce qhost data request sizes at sge_qmaster

By setting the environment variable `SGE_GDI_REQUEST_REDUCE_LEVEL` it is possible to reduce the amount of data transferred from `sge_qmaster` to `qhost` clients. A detailed description can be found in the `qhost(1)` man page (ENVIRONMENTAL VARIABLES).

3.1.4 8.5.1: New environment variables in the job environment

Univa Grid Engine sets two new environment variables in the environment of the job, the prolog, `pe_start`, `pe_stop` and epilog scripts:

```
SGE_RERUN_REQUESTED=<0|1|2>
```

A value of 0 means there was no `-r <y|n>` request on the submit command line of the job, 1 means `-r y` was requested and 2 means `-r n` was requested.

```
SGE_RERUN_JOB=<0|1>
```

A value of 1 means the job will be rescheduled on error. The value is determined from the `SGE_RERUN_REQUESTED` value and the configuration value `rerun` of the queue the job runs in.

Additionally, Univa Grid Engine sets this new environment variable in the environment of the `pe_stop` and `epilog` scripts:

```
SGE_JOB_EXIT_STATUS
```

This variable is set to the exit status of the job. This is the same value that is written to the accounting to the `exit_status` field.

3.1.5 8.5.1: New example script for jsv and core-binding

A new example script that demonstrates core-binding using JSV can be found at “`$SGE_ROOT/util/resources/jsv/core_binding_jsv.sh`”

3.1.6 8.5.1: sgepasswd renewal

The upgrade and installation scripts have been updated to ensure that CSP/sgepasswd key store is backed up and is restored correctly in a clone upgrade. If you are currently using CSP or sgepasswd you need to save your configuration as user root with:

```
# $SGE_ROOT/util/upgrade_modules/save_sge_config.sh <backupdir>
```

then replace the original script in your existing installation with the new Univa Grid Engine 8.5.1 one to also ensure backup of existing sgeCA infrastructure. Now an upgrade with `inst_sge-upd -csp` will restore your backed up key store. If you create new key store by creating a new sgeCA infrastructure you will have to reencrypt an existing sgepasswd file manually with the following command as root:

```
# $SGE_ROOT/bin/<sge_arch>/sgepasswd -n \  
    /var/sgeCA/<old port number>/<old sge_cell>.backup/private/key.pem
```

The original sgepasswd file is stored as

```
# ls $SGE_ROOT/$SGE_CELL/common/sgepasswd.oldcert_backup
```

and the reencrypted file is available as

```
# ls $SGE_ROOT/$SGE_CELL/common/sgepasswd
```

Please do not repeat this process without first saving the original files otherwise you might lose your original information and need to recreate the sgepasswd file from scratch.

Please be aware that the encryption algorithm has changed in version 8.5.0. If you are upgrading from an older installation than 8.5.0 you have to first follow the steps under section: [Changes in Windows execution host sgepasswd file](#)

3.1.7 Performance Improvements and Memory Requirements

For Univa Grid Engine 8.5.3 we invested quite some time to improve the performance of various Univa Grid Engine components and libraries.

As a consequence the following metrics of the cluster have been improved compared to Univa Grid Engine 8.4.4 versions of Univa Grid Engine:

- Submit rate (increased by 5-15 % depending of the jobs types and requested functionality)
- Scheduling times (reduced by 5-30 % depending on the used policies)
- Faster delivery of dispatched jobs to `sge_execd` esp. for interactive jobs
- Memory requirements for request handling (reduced by 5-10%) especially for read-only requests like `qstat`, `qhost`, ... (reduced by 5-30 %)
- Processing and response time of requests send by execution hosts (certain requests will now be handled in parallel within `qmaster`)
- Processing of clients requests like `qstat/qhost` (which results in about 30% more requests that can be handled in the same amount of time with the same memory requirements)
- Job turnaround times

This improves the overall cluster throughput as well as interactions with the Univa Grid Engine cluster.

The speedup in your cluster depends on the details of the cluster setup and on the features of Univa Grid Engine that are enabled or disabled.

3.1.8 Standing Reservations

In Univa Grid Engine 8.5.3 the Advance Reservation feature has been extended to allow for Standing Reservations.

A Standing Reservation is a recurring Advance Reservation. Start and end times of the individual Advance Reservations are specified via a calendar, additional command line options allow for the specification of the number of reservations at a time and the behaviour in case a reservation cannot be granted.

All options available for Advance Reservations such as resource requests are also available for Standing Reservations.

See [User Guide -> Reservations](#) for details.

3.1.9 Policy Scheme: Consider Slots Instead of Jobs

Univa Grid Engine 8.5.3 provides a configuration option where the scheduler will consider the number of slots used by running jobs and by pending jobs when calculating users and projects contribution toward their sharing goals as defined by the share tree. That is, a parallel job using 4 slots will be considered equal in terms of resource usage to 4 serial jobs. The previous share tree algorithm did not take into account slot use which meant that if a mix of parallel and serial jobs were running or queued, the number of tickets granted to pending jobs did not result in the correct run-time sharing ratios and the share tree targets were not met. For example, if two

projects “a” and “b” are configured at the same level in the share tree with equal shares, the scheduler should try to schedule jobs so that the projects get equal usage. However, if project “a” has mostly parallel jobs, it will tend to get more usage because the previous share tree algorithm treats all jobs equally. In fact, with the old algorithm, if we look at the prioritized order of pending 4-slot jobs for project “a” and pending 1-slot jobs for project “b” for a share tree with no usage, we would see the pending jobs interleaved (a b a b a b a b . . .). With the new algorithm, we would see the pending jobs ordered based on their slot usage (a b b b b a b b b b . . .), which is more likely to lead to the proper runtime sharing ratios.

The urgency_slots PE attribute will be used to determine the assumed number of slots used by a pending job with a slot range. See urgency_slots in the sge_pe(5) man page for additional information.

The old behavior (sharing based on jobs) can be configured by setting SHARE_BASED_ON_SLOTS=false (default) in the sched_conf(5) params attribute. The new behavior (sharing based on slots) can be configured by setting SHARE_BASED_ON_SLOTS=true in the sched_conf(5) params attribute. See the sched_conf(5) man page for more details.

```
$ qconf -msconf
...
params SHARE_BASED_ON_SLOTS=true
...
```

Please note that beginning with version 8.6.0 of Univa Grid Engine the default for SHARE_BASED_ON_SLOTS will be changed from false to true.

3.1.10 RSMAP Enhancements

In Univa Grid Engine 8.5.3 it is possible to request specific Ids of a Resource Map Complex (RSMAP) via command line with the syntax used for RESTRING (see complex(5) man page for details). The following example submits a job that requests four Ids of the complex “GPU”, three Ids with the name “gpu1” or “gpu2” and one Id with the name “gpu3”.

```
qsub -l GPU=3(gpu1|gpu2)&1(gpu3) $SGE_ROOT/examples/jobs/sleeper.sh 3600
```

Depending on the host configuration and the available Ids, one possible combination of assigned Ids for this job is gpu1 gpu1 gpu2 gpu3.

A job cannot be scheduled if the scheduler cannot find enough free Ids with the requested names, even if there are enough free Ids with different names available.

It is still possible to use the RSMAP complex without the syntax enhancements introduced in Univa Grid Engine 8.5.3. The scheduler will then behave like in previous versions and use any free id.

Please be aware that very complicated requests may slow down the scheduler.

To make the configuration of RSMAPs easier, a shortcut has been added.

The syntax is:

```
complex_values    complex_name=amount(complex_id:amount)
```

The following example defines a complex named “GPU” with ten available Ids, five with the name “gpu1” and five with the name “gpu2”:

```
qconf -me exechost1
...
complex_values    GPU=10(gpu1:5 gpu2:5)
...
```

3.1.11 Improved Scheduler Profiling

In previous versions of Univa Grid Engine The scheduler profiling did not completely cover the scheduling main loop. This lead into some wrong or missing profiling data. Univa Grid Engine has additional diagnostics to cover the main loop. Please read the updated `sge_diagnostics(1)` man page where the scheduler profiling is described in detail.

3.1.12 Improved Logging

The `sge_diagnostics(1)` man page was introduced to provide an overview on available logging and diagnostic options. The most important changes and new options are:

- show statistics about request types in worker and reader request queues (see “MONITOR_REQUEST_QUEUES”, man page “`sge_conf(5)`”)
- log spooling exceeding a certain threshold (see “LOG_SPOOLING_TIME”, man page “`sge_conf(5)`”)
- communication errors at first startup not logged into `/tmp/execd_messages.` file
- communication specific enhancements for profiling and startup behavior (see “PROF_COMMLIB_TIME”, man page “`sge_conf(5)`”)
- log job verification time exceeding a certain threshold (see “LOG_JOB_VERIFICATION_TIME”, man page “`sge_conf(5)`”)
- log request processing exceeding a certain threshold (see “LOG_REQUEST_PROCESSING_TIME”, man page “`sge_conf(5)`”)

3.1.13 Encryption in CSP mode / `sgepasswd`

The encryption algorithm has been changed from RC4 to AES256_CBC. This affects CSP encryption and the encryption of the Windows `execd sgepasswd` file. There are no additional upgrade steps necessary for CSP mode and the steps for Windows are described below in section: [Changes in Windows execution host `sgepasswd` file](#)

3.1.14 Online usage of running Windows jobs

For Univa Grid Engine jobs running on Microsoft Windows, some usage values are now reported while the jobs are running. The usage values that are reported are:

wallclock, cpu, mem, io, ioops, vmem, maxvmem

The online usage value `iow` is not reported on Microsoft Windows.

Not all Windows usage values exactly match the corresponding UNIX usage values; this list shows what Windows system value is retrieved to report the corresponding usage value:

Online usage	Corresponding Windows system value
wallclock	(current wallclock time - start wallclock time) of the job.
cpu	(kernel time + user time) of all processes in the job.
mem	integral of vmem over time.
io	(read bytes + write bytes + other bytes) of all processes in the job.
ioops	(read operations + write operations + other operations) of all processes in the job.
vmem	private usage of all processes in the job. The private usage is documented to be the “commit charge for this process”, which is nearly the same as the virtual memory of UNIX.
maxvmem	maximum of all measured vmem values of the whole job.

These online usage values are printed e.g. using the command:

```
> qstat -j <job_id>
```

Additionally, for finished jobs that ran on Microsoft Windows, these usage values are accounted: `ru_wallclock`, `ru_utime`, `ru_stime`, `wallclock`, `cpu`, `mem`, `io`, `ioops`, `maxvmem`

Accounting usage	Corresponding Windows system value
ru_wallclock	the wallclock of the whole job measured by Windows.
ru_utime	user time of the whole Windows Job object that contains all processes of the job.
ru_stime	kernel time of the whole Windows Job object that contains all processes of the job.

wallclock, cpu, mem, io, ioops, maxvmem are identical to the online usage values.

The accounting values for finished jobs are printed e.g. using the command:

```
> qacct -j <job_id>
```


3.1.15 Docker Related Enhancements

With Univa Grid Engine 8.5.3, variable placeholders are allowed in sub-options of the “-xd” option on the submit command line, sge_request files, job scripts, job classes and job submission verifier. These variable placeholders are resolved by corresponding elements of specific RSMAP complexes the Scheduler selects for the tasks of a job.

The format of these placeholders is:

```
placeholder := '${ complex_name '(' index ')' '}' .
```

where `complex_name` is the name of the corresponding RSMAP complex and `index` is the index of the element the scheduler selects from the RSMAP for this job, starting with 0.

E.g.:

If a resource map defines these values on a host: `gpu_map=4(0 1 2 3)`
this `qsub` command line is used:

```
# qsub -l docker,docker_images="*some_image*",gpu_map=2
  -xd "--device=/dev/gpu${gpu_map(0)}:/dev/gpu0,
      --device=/dev/gpu${gpu_map(1)}:/dev/gpu1" ...
```

and the scheduler selects the elements “1” and “3” from the resource map, the command line is resolved to

```
# qsub -l docker,docker_images="*some_image*",gpu_map=2
  -xd "--device=/dev/gpu1:/dev/gpu0,
      --device=/dev/gpu3:/dev/gpu1"...
```

which means the physical GPUs “gpu1” and “gpu3” are mapped to the virtual GPUs “gpu0” and “gpu1” inside the container and at the same time are exclusively reserved for the current job among all Univa Grid Engine jobs.

3.1.16 Host Aliasing and Resolving

Univa Grid Engine now better supports changes to the `host_aliases` file while Univa Grid Engine is running.

Periodically naming services like DNS or NIS may be updated which may result in changed hostnames, additionally administrators may update the `host_aliases` file. Both of these situations result in changes to the host name resolution in Univa Grid Engine. Univa Grid Engine has been enhanced to handle the following situations:

Adding `host_aliases` while Univa Grid Engine is running: Adding new entries to the `host_aliases` file is supported while Univa Grid Engine is running if the resulting name and none of the mapped hostnames is referenced in any Univa Grid Engine configuration. Host names that are changed or added which are referenced in any Univa Grid Engine configuration object will be ignored and a message will be logged in the `qmaster` messages file.

Update of internal name resolution database on daemon startup: At startup of the qmaster daemon any change of a hostname in the configuration will be detected and the name resolution database will be adjusted to reflect this change. If the name resolution change affects the execution daemon nodes they must be restarted by the Administrator.

Additional Improvements: Improved several places (e.g. plain hostnames used in regular expressions, hostnames reported by loadsensors) where hostnames are entering the system. This resulted in not scheduleable jobs and other problems in the past. The final Univa Grid Engine version will get an updated `host_aliases` man page and an updated admin guide (GE-6013).

3.1.17 Intel® Xeon Phi™ x200 (Knights Landing) integration

Univa Grid Engine 8.5.3 provides an integration for Intel® Xeon Phi™ x200 (Knights Landing) Processors. The pre-compiled load-sensor automatically detects the current Cluster and 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 AdminsGuideGE for details.

3.2 Full List of Fixes and Enhancements

3.2.1 Univa Grid Engine 8.5.0alpha1 (also fixed for a 8.3 or 8.4 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-4229 Reduce executable sizes by removing extra symbols
- GE-4288 Confusing message on h_rt or s_rt limit
- 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-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-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-5289 add a note about the msvc redistrib dll to the installation guide
- GE-5332 DRMAA2 job template needs to support a native specification replacement
- 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-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-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-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 -2³¹ 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 L0 job stays in "dr" state after qdel
- GE-5723 qalter -p not transfered to L0
- GE-5726 licence_constraints in UGE clusters are not updated
- 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-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-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-5892 RQS limits with Job Classes do not work when max_reservation > 0
- 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-5925 wrong qdel message when a job is already in deletion

GE-5933 PE jobs with start/stop procedure or jobs with prolog/epilog requesting a pty change ownership of /dev/null to job user

GE-5978 performance regression when using DRMAA2 monitoring session

GE-5981 add memory usage values as extensions in DRMAA2 job info

GE-5993 qmaster segfaults when deleting jobs with non-existing LO licenses

GE-6030 possibility to switch off 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-6097 log request processing exceeding a certain threshold

GE-6111 sge_qmaster crashes when deleting an advance reservation with a partially finished array job

GE-6152 log spooling exceeding a certain threshold

GE-6167 add new scheduler profiling line for wait times

GE-6172 jdrmaa2 openJobSession implementation is missing

GE-6174 qmaster crashes with: ->|C|!!!!!!!!!!!! QU_qname not found in element !!!!!!!!!!!!!, host is removed from queue

GE-6176 qdel -f prints unexpected messages, e.g. debug information

GE-6221 sge_qmaster might crash with specific logging settings

GE-6230 Null DRMAA2 Native Specifications Fields Can Cause Segfaults

UWSA-164 advance reservations does not contain the start time

UWSA-165 creation of an AR results in a wrong duration error

UWSA-175 support CUDA load_sensor values in resourceNumericValues

3.2.2 Univa Grid Engine 8.5.0alpha1 (also fixed for a 8.4 patch release)

GE-3146 resource reservation is broken with SGE calendar

GE-3227 AR shouldn't be scheduled to already disabled queues at time of submission

GE-4158 Some of the job class attributes are incorrect in man page and users guide

GE-4293 qsub -w e -l exclusive=true rejects job, even if the request is valid

GE-4425 SGE_LONG_QNAMES=-1 lead to qstat segfault

GE-4497 PE job is not scheduled when a non-requestable consumable is setup in global host

GE-4603 Job <jid> cannot run in PE <pe_name> because it only offers 0 slots

GE-4672 default_jc and enforce_jc are not documented in the man page

GE-4908 native Windows (win-x86) UGE binaries can't find the SGE_ROOT directory if it is the root directory of a share

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-5345 UGE to auto resolve host_aliases
- GE-5492 cuda loadsensor bash script compatibility issue
- GE-5509 host_aliases not working for resource hostname OR request
- GE-5510 host_aliases not working for qconf -purge request
- GE-5524 newline in job submission breaks reporting/accounting lines and qstat -j
- GE-5528 hostname resolving changes should trigger update of all affected data objects at qmaster/execd daemon startup
- 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-5589 changes to host_aliases file should be updated when hosts are re-resolved
- GE-5604 install_cuda_complexes.sh produces invalid complex
- GE-5635 multiple occurrences of same resource in RSMAP is not working
- 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-5675 Fully integrate Univa Grid Engine with systemd
- 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-5717 basic environment check during startup of Grid Engine daemons
- 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-5774 enable execution of tasks of parallel jobs in Docker containers
- 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-5795 exit dispatching loop immediately when shutdown of scheduler thread is triggered
- GE-5821 remove basic Docker integration that uses load sensor and starter_method when real Docker integration works
- GE-5834 bad performance of RQS rules on host groups
- GE-5849 new masterl switch s_rss limit setting is enforced for slave

- tasks
- 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-5901 allow Docker containers to automatically start the application configured in the Docker image
 - 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-5955 host_aliases not working for sge_shadowd
 - GE-5962 the Docker daemon doesn't download an image when its not available locally
 - GE-5989 jobs are not started in Docker containers if the job user has insufficient permissions to write docker specific files
 - 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-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-6028 User/group management done via Windows Active Domain might break UGE
 - 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-6045 on native Windows (win-x86), the win_getpwnam_r() always tries to load the user profile, even if called with insufficient permissions
 - GE-6046 infinite loop writing to trace file in docker container
 - GE-6049 default job class is not used correctly
 - GE-6056 Cray XC integration needs to set unlimited timeout for epilog for newer UGE version
 - GE-6059 job gets rescheduled when epilog gets a SIGABRT signal
 - GE-6064 on native Windows (win-x86), error handling of (Un)LoadUserProfile() overwrites real error message
 - GE-6067 adding a queue with calendar generates error messages and no jobs will be scheduled into this queue
 - GE-6068 Interactive Docker jobs fail because sge_container_shepherd doesn't have the permission to "write the shepherd_about_to_exit" file
 - GE-6073 qsub does not support multiple use of "-binding" or "-t" parameters
 - GE-6076 on native Windows (win-x86), all job spool directories are deleted when the first job on an execution host finishes
 - 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-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-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-6126 submitting a docker autostart job fails if no job name is given
 - 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)
 - GE-6156 execd crashes when a docker job is submitted, requesting memory resources
 - GE-6163 errors for sge_mirror_process_events() result in scheduler inactivity until it times out after 600 seconds
 - GE-6173 Description of certain states for qstat -s is missing in help output and man page
 - 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-6211 the provided Docker container name is overwritten by automatically generated one
 - GE-6219 missing l10n for INFO message plus wrong newline character at the end of message
 - GE-6262 broken sge_usleep() could cause very long sleeps
 - GE-6168 sgepasswd issues

3.2.3 Univa Grid Engine 8.5.0alpha1

- GE-575 Parallel jobs exceeding wall clock time are not killed
- GE-3909 job submitted with user not on every execd node is stuck in zombie state
- GE-4991 loading dynamic libraries fails if uid != euid
- GE-5317 remove "verify_suitable_queues" from qstat -j output

- GE-5505 Job that is restarted and that was in certain error states previously stays in t-state
- GE-5540 GE-5949 drmaa does not provide the submit_cmd line
- GE-5560 Java Binding for DRMAA2
- GE-5583 enhance RSMAP so that necessity of OR requests in command line can be avoided
- GE-5600 Windows (win-x86) queue instance is set to error state if job can't be executed
- GE-5608 Atomically Delete Jobs by State
- GE-5648 cleanup and speed-up execd job start when using additional group ids
- GE-5651 Unify behavior of UGE calls that trigger library functions depending on directory services (NIS, LDAP, Active Directory, ...)
- GE-5657 Enhance RSMAP so that only elements / instances with the same name within a RSMAP are chosen.
- GE-5658 Enhance RSMAP so that instances can be requested via REGEX
- GE-5659 Allowing to configure RSMAP that the above OR is restricted only to one type.
- GE-5660 Allow shortcut for RSMAP definition
- GE-5661 Allow load sensor to restrict RSMAP
- GE-5668 Posix priority 0 not anymore scaled to 0.5
- GE-5696 add xd switch to job classes
- GE-5721 improve and reduce qhold/qrls messages and logging
- GE-5855 skip pending read-only requests where clients are already gone
- GE-5856 replace 'old' execd reports in the list of pending worker requests
- GE-5875 functionality missing that allows to filter thread specific debug outputs of components
- GE-5899 AR accepts jobs after cleanup of running jobs before end of AR
- GE-5957 jobs submitted with qsub -sync are not killed when Ctrl+C is pressed
- GE-5985 sharetree algorithm should consider slots, not jobs
- GE-6015 It takes 2 load_reports_intervals to get load_values when a new load_sensor is setup
- GE-6054 docker jobs started with qrsh have incorrect environment (SGE_ROOT, ...)
- GE-6075 Improved scheduler performance due to CULL performance improvement of all lGet/lSet functions
- GE-6088 make sure AR and SR are correctly written to the reporting file
- GE-6096 improve and reduce qdel messages/logging
- GE-6110 Reduce communication overhead by optimizing communication layer
- GE-6120 sge_execd reports "can't resolve hostname "_sge_pseudo_host" in messages file
- GE-6128 update openssl to 1.0.2j
- GE-6134 qmaster receives empty report lists that qmaster tries to process.
- GE-6136 Lock handling done in qmaster for reports send by execd might cause slowdown of qmaster
- GE-6138 Reports from execd's are always send in groups. Qmaster uses the same lock for all.
- GE-6139 Introduce possibility to set job report flush time @ execd

- GE-6142 All read-only-requests from execd's will be executed by worker threads
- GE-6162 qmaster crashes at startup when a SR with a multi time calendar is spooled
- GE-6169 Improve performance of packing operations and reduce size of objects within packbuffers
- GE-6232 reporting of job end is delayed up to one second in sge_execd
- GE-6241 update jemalloc to version 4.3.1
- GE-6242 reduce the number of events subscribed by scheduler
- GE-6250 speed up delivery of scheduler job start orders
- GE-6266 -clearp switch not available for xd-attribute of job
- GE-6267 -mods does not work correctly with rou-attribute
- GE-6268 rou-attribute of JC is not tagged as list attribute in man page
- GE-6269 XML output might be incorrect for JB_mail_list depending on the API or submit client that was used to submit the jobs
- GE-6153 use stronger cipher for CSP mode communication and sgepasswd file encryption

3.2.4 Univa Grid Engine 8.5.0alpha2

- GE-3060 in order to connect through a firewall, qrsh/qlogin should use only specified port numbers
- GE-3341 pe_min value is not correct in JSV if multiple slots ranges are requested
- GE-3928 Document and/or change default shell for qsub '-b y'
- GE-5631 -mods/-clears does not work for masterl
- GE-5632 -adds/-mods/-clears for "masterl" not described in qsub man page
- GE-5906 possible endless loop in test_drmaa2 -> test_job_wait_all
- GE-5980 getInfo() does always return a null
- GE-5997 allow to specify placeholders for Scheduler selected values in the argument list to the -xd option
- GE-6000 Have UGE documentation available in HTML format.
- GE-6026 Documentation for job name starting with a digit needs to be corrected
- GE-6027 scheduler thread repeats logging of cluster configuration
- GE-6131 update postgres library to version 9.6
- GE-6145 qsub client crashes with specific sge_request file
- GE-6153 use stronger cipher for CSP mode communication and sgepasswd file encryption
- GE-6154 Configuration API needs to be part of the regular package build
- GE-6164 qmaster, execd and scheduler params should be in upper case letters in man page and documentation
- GE-6165 log a warning when qmaster reads unknown cluster config parameters
- GE-6166 log a warning when qmaster and execd read unknown *_params
- GE-6168 sgepasswd issues
- GE-6170 qsub_time is set to a time at year 1970, in case of job error
- GE-6192 reduce spooling operations for a job lifecycle
- GE-6251 Rest and Config API should also support 'port_range' parameter
- GE-6283 switch and argument in xd attribute of job classes cannot be delimited by space although this is possible at command line in qsub
- GE-6289 load sensor script triggered endlessly (independent from the configured load report interval)

GE-6291 scheduler profiling does not show time spend for updating event master settings
 GE-6314 JC variant output for xd attribute missing
 GE-6317 resource bookkeeping is broken, scheduler stops dispatching large parallel jobs, qmaster getting overloaded
 GE-6323 Scheduler tries to find a master_queue for not schedulable job even if there is no master queue request
 GE-6329 changing a resource with qalter -mods fails if mem_free is part of the requested resources in the qsub command
 GE-6331 docker autostart jobs run through but exit_status is 1 and start/end time missing
 GE-6336 incorrect hostname resolving for qsub -q parameter
 GE-6337 transient builtin load values are spooled
 GE-6358 sge_execd might crash on AIX
 GE-6341 Config API should also support new job class 'xd' parameter
 GE-6345 build sol-sparc64 packages on Solaris 10 instead of Solaris 9
 GE-6352 qalter displays success message even if modification is not allowed due to job class
 GE-6353 shepherd crashes on MacOS and possibly other OS after job finishes
 GE-6357 FQDN Hostnames are limited to 63 Characters
 GE-6359 remove obsolete load sensor scripts
 UWSA-177 automatic installation is broken -> missing SGE_CELL setting
 UWSA-179 Document should mention exact file name for adjusting logging
 UWSA-183 UGERest Api crashes when masterl switch is used

3.2.5 Univa Grid Engine 8.5.0beta1

GE-4876 make the native Windows PDC report IO and memory values
 GE-4956 qmon about box contains incorrect copyright information
 GE-4967 general documentation and man pages typos and issues
 GE-5721 improve and reduce qhold/qrls messages and logging
 GE-5999 Preempted jobs stay in P-state
 GE-6107 add information about not properly started Docker daemon to AdminGuide
 GE-6316 rework parts of the UGE documentation
 GE-6332 docker containers are not always removed after job end
 GE-6354 Improve qalter message logging
 GE-6372 possible execd crash in Docker job handling
 GE-6374 job start fails with Docker 1.13.0
 GE-6375 job lost detection might cause severe problems in qmaster
 GE-6376 for Docker jobs, a wrong mem online and accounting usage value is reported
 GE-6381 possible handshake problems with external load sensors
 GE-6382 AAPRE cplx column is not set correctly in qmon complex dialog
 GE-6383 qmon crash when job is suspended
 GE-6387 qtcsh does not provide an error message when command should be executed on remote host but SGE_ROOT is not set
 GE-6388 unexpected CRITICAL ERROR logging in qmaster messages file

- GE-6394 wrongly, wallclock is summed up over all PE tasks if
accounting_summary=true
- GE-6395 qalter -w p does not display all schedd_job_info messages

3.2.6 Univa Grid Engine 8.5.0 FCS

- GE-4170 -help for ./inst_sge -upd
- GE-4514 Spooledit can create strange 'CONFIG:CONFIG:global' entry
- GE-4699 one cannot create CSP credentials for a user 'a', if a user
'ab' already exists
- GE-5502 modularize install and upgrade scripts
- GE-5804 far too much usage data spooled in user objects
- GE-6013 Enhance host_aliases man page and admin guide
- GE-6026 Documentation for job name starting with a digit needs to
be corrected
- GE_6216 fix memory leaks in native Windows (win-x86) execution daemon
- GE-6289 load sensor script triggered endlessly (independent from
the configured load report interval)
- GE-6291 scheduler profiling does not show time spend for updating
event master settings
- GE-6336 incorrect hostname resolving for qsub -q parameter
- GE-6341 Config API should also support new job class 'xd' parameter
- GE-6357 FQDN Hostnames are limited to 63 Characters
- GE-6362 fix memory leaks and access issues in 8.5.0
- GE-6384 ARs oversubscribe queue slots
- GE-6397 shutdown of execd in early startup phase might take long
- GE-6398 jobs running in Docker containers write files with wrong
group ownership
- GE-6399 if the version comparison in the upgrade script doesn't find
the given version, upgrading should be stopped
- GE-6400 invalid load and memory usage reported by execd during install
- GE-6401 qmon Modify queue dialog fails with 'no default value for slots'
- GE-6405 qmon Clone queue dialog fails with "'Multiple values for one
queue domain/host..."
- GE-6409 autostarting Docker job cause the shepherd to use 100% CPU for
many seconds at job end
- GE-6411 Jobs sometimes don't get a reservation when using RQS
- UWSA-184 Expose the 'qmod -p' preemption functionality in the REST API

3.2.7 Univa Grid Engine 8.5.1

- GE-4305 Better documentation for basic share tree use case
- GE-4389 enhance qsub man page with JSV modification examples for core
binding
- GE-5650 with allocation rule \$fill_up, slave resource requests are not
obeyed when master resource and queue requests are provided
- GE-5806 parallel jobs might not startup due to wrong RQS calculations
- GE-5941 renewing certificates makes sgepasswd file unreadable

- GE-6020 allow native Windows (win-x86) functions to retry to logon users several times
- GE-6103 jobs are bound to cores even if no binding is requested
- GE-6183 sge_shepherd sets limits too high for master task
- GE-6237 jobs are being restarted even if the '-r no' was specified during submission
- GE-6402 Scheduler might not respect RQS limits during the time RQS rules are changed
- GE-6407 add documentation for config-api
- GE-6413 shepherd does not handle all error responses to a pull Docker image request
- GE-6434 incomplete binding requests shown in qstat for long lists of binding requests
- GE-6462 on native Windows (win-x86), environment variable values containing an equal sign are truncated
- GE-6467 installer fails to restart qmaster when upgrading and using SYSTEMD
- GE-6478 Very long load value of loadsensor causes segfault of execd
- GE-6480 wrong/missing error messages and wrong exit status when initializing invalid RSMAP ranges
- GE-6489 using of external loadsensors might deadlock execution daemon
- GE-6490 load sensor specific errors are not logged into execd messages file
- GE-6493 RSMAP map entry selection request does not work
- GE-6510 Core binding: striding-strategy counts needed cores wrong and rejects viable hosts
- GE-6511 gdi_request_limits does not behave as documented
- GE-6521 Qmaster crashes with LO enabled and job dependencies
- GE-6534 sge_execd crash with core dump with GPGPU jobs
- GE-6537 Submitting RSMAP-range jobs via -adds is broken
- GE-6553 RQS limits incorrectly applied when PE job submitted with "-l h=<host>"
- GE-6573 reduce resulting qhost client requested data sizes transferred via network
- GE-6575 stree-edit utility broken
- GE-6583 scheduler is wrongly skipping hosts or queue instances for parallel jobs that are using master task specific requests
- GE-6587 abort of sge_qmaster if a JC where V is set to yes is used to create a job
- UWSA-81 Add support for listening only on localhost
- UWSA-186 requested jobEnvironment is not shown in jobs
- UWSA-188 jobEnvironment ugerestsdk contains error in json converter
- UWSA-189 ugerest is showing same scheduler conf twice
- UWSA-190 upgrade of restlet-jse-2.3.6 to restlet-jse-2.3.10

3.2.8 Univa Grid Engine 8.5.2

- GE-5569 native Windows (win-x86) sge_execd exits if it cannot access

- the act_qmaster file
- GE-6454 improve documentation of Docker integration
- GE-6236 accounting file broken on Solaris if accounting line has 1023 characters
- GE-6615 enhance error logging if load sensor cannot be started
- GE-6618 native Windows (win-x86) execution daemon crashes if qloadsensor does not work
- GE-6620 error messages from qloadsensor.exe (win-x86) are not forwarded to the execd messages file
- GE-6623 add the UGE admin user to the Performance Monitor Users group on native Windows 10 (win-x86)
- GE-6629 qhost NSOC and NCOR incorrect on lx-arm64
- GE-6637 using hostgroups in rqs limit definition can trigger short qmaster hang at startup or rqs modify request
- GE-6640 qloadsensor.exe report "no error" if an error occurs while initializing the PDH service
- GE-6644 any epilog SIGSEGVs and sets queue in error state with execed_params INHERIT_ENV=false
- GE-6660 Requesting a RSMAP without ID can crash qmaster

3.2.9 Univa Grid Engine 8.5.3

- GE-3721 qstat -j "*" -u "user1" is not working
- GE-5290 qstat (-xml) does not accept filter switches when -j "*" is specified
- GE-6432 qdel -u "*" is only allowed to managers, not to operators
- GE-6483 document placeholders in Docker requests in the UserGuide and man pages
- GE-6497 support halftime -1 setting in scheduler config to disable past usage for sharetree
- GE-6568 Allow for forced job deletion through UGE REST
- GE-6578 Duplicate calendar entries associated with host_aliases
- GE-6595 Docker interactive job can't be deleted by qdel
- GE-6617 provide workaround in native Windows (win-x86) sgeexecd.bat script for start /b bug on Windows 10, version 10.0.15063
- GE-6664 When host aliases are configured qsub -sync y reports "commlib info: successfully updated host aliases (add: 0, del: 0)"
- GE-6670 The global configuration parameter "gdi_request_limits" not working for aliased hostnames
- GE-6671 sudo requests for same user as ugerest service user are rejected
- GE-6678 Improve accept() handling in commlib
- GE-6682 exec host cannot startup if no admin or submit host
- GE-6687 RSMAP-topology-masks jobs and -binding jobs result in wrong scheduling decisions
- GE-6695 qrstat does not output cal_depth and cal_jmp information for standing reservations
- GE-6696 save_sge_config.sh needs to dump advance/standing reservations for upgrade to newer versions

- GE-6697 `qstat -j "*" does not support "$user" placeholder set in sge_qstat request file`
- GE-6706 `unexpected logging and possibly incorrect accounting if multiple array tasks of a job are running on a host`

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.X.X to Univa Grid Engine 8.5 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.X.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 4: Upgrading from SGE, OGE, UGE 8.X.X to Univa Grid Engine 8.5.X

5 Compatibility Notes

5.1 Changes in Output Format of Commands

The output format of some commands has been changed:

- ‘qstat -xml’ output for jobs did sometimes show the mail recipient in the XML node “JB_mail_list/element/MR_user” or “JB_mail_list/mail_list/MR_user” of jobs depending on the used submit interface or submit client. Beginning with Univa Grid Engine 8.5.3 it will always be: “JB_mail_list/element/MR_user”. (see GE-6269)
- ‘qstat -j’ output showed an Univa Grid Engine internally used parameter with the name ‘verify_suitable_queues’. This line of output is not shown anymore beginning with Univa Grid Engine 8.5.3 (see GE-5317)
- ‘qstat -j’ might show additional online usage values compared to previous versions of UGE. Depending on the architecture of the underlying execution node of a job following values might be reported additionally in the ‘usage’ line: ioops (number of io operations), iow (io waiting time). (see GE-4296)
- If resource requests of jobs contained newline-characters then they were also shown in the ‘qstat -j’ output. This issue has been resolved. (see GE-5524)

5.2 Changes in Windows execution host sgepasswd file

The encryption algorithm for the “\$SGE_ROOT/\$SGE_CELL/common/sgepasswd” file passwords has been changed from RC4 to AES-256-CBC:

If you upgrade to Univa Grid Engine 8.5.3 you need to convert your existing sgepasswd file during the upgrade procedure.

Become superuser and execute the following command on the master machine:

```
# sgepasswd -c
```

This will create a backup of your original ‘sgepasswd’ file as ‘sgepasswd.old_algorithm_backup’ and create the new compatible sgepasswd file. Otherwise encryption related error messages might show up.

If you create a new sgepasswd file from scratch no additional steps compared to previous versions are required.

5.3 Deprecated Functionality

Following components/features are deprecated and will be removed with version 8.6.0 of Univa Grid Engine:

- Graphical Installer
- qtsh

6 Known Issues and Limitations

6.1 setting halftime to -1 (GE-6497) not supported in qmon

It is recommended to use the qconf command line client.