

Qube 6.8 Complete Release Notes

#####

@RELEASE: 6.8-4a

#####

This is a cumulative patch release of the qube-core, supervisor, and worker packages, for all platforms, including several key fixes.

==== CL 17208 ====

@CHANGE: Populate the subjob (instance) objects with more data (like status), and not just the IDs, when subjob info is requested via "qbhostinfo" (qb.hostinfo(subjobs=True) for python API)

Previously, only jobid, subid, and host info (name, address, macaddress) were filled. Now, things like "status", "timestart", "allocations", etc. are properly filled in.

JIRA: QUBE-2073

ZD: 16541

==== CL 17206 ====

@FIX: When "migrate_on_frame_retry" job flag is set, prevent backend from doing further processing (especially another requestwork()) after a work failed

This was causing race-conditions that will get agenda items to be stuck in "retrying" state, while there are no instances processing them.

Now the reportwork() API routine is modified so that if it's invoked to report that a work "failed", and the "migrate_on_frame_retry" is set on the job, it will stop processing (does a long sleep), and let the worker/proxy do the process clean up.

JIRA: QUBE-2202

ZD: 16553

==== CL 17186 ====

@FIX: "VirtualBox Host-Only Ethernet Adapter" now when daemons (supe, worker) try to pick a primary mac address

JIRA: QUBE-2149

ZD: 16561

==== CL 17182 ====

@CHANGE: all classes that inherit from QbObject print as a regular dictionary, no longer have a __repr__ which prints the job data as a single flat string

@NEW: add qb.validatejob() function to python API, help find malformed jobs that crash the user interfaces

==== CL 17141 ====

@FIX: Any job submitted from within a running job picks up the pgrp of the submitting job

By design, if the submission environment has QBGRPID and QBJOBID set, the API's submission routine will set the job's pgrp and pid, respectively to the values specified in the environment variables.

One couldn't override this "inheritance" behavior even by explicitly specifying "pgrp" or "pid" in the job being submitted, for instance with the "-pgrp" command-line option of qbsub.

Fixed, so that setting "pgrp" to 0 on submission means that the job should generate its own pgrp instead of inheriting it from the environment.

JIRA: QUBE-2141

ZD: 16545

==== CL 17101 ====

@NEW: add "-dying" and "-registering" options to qbjobs.

@CHANGE: also add dying and registering jobs to the "-active" filter.

JIRA: QUBE-2091

ZD: 16469

==== CL 16804 ====

@TWEAK: added code to print what operation was requested, when printing out "permission granted to user..."

#####

@RELEASE: 6.8-4

#####

This is a cumulative patch release for all platforms.

==== CL 16628 ====

@FIX: "qb_default_string()" warning printed during linux qube-core installation

Corrected code so that warnings like the following won't print any more:

WARNING: qb_default_string() unknown value[1001]

WARNING: qb_default_string() unknown value[1002]

JIRA: QUBE-1894

==== CL 16602 ====

@FIX: misleading database name printed in error handler for MySQL stored procedures PFX_CALC_CPU_TIME() and PFX_CALC_AVG_WORK_TIME(); "ERROR: TABLE NOT FOUND IN DB pfx_dw.<actual_database_name>"

==== CL 16517 ====

@FIX: C4D appFinder jobs don't apply path translation properly on Windows, backslashes are converted too early

==== CL 16491 ====

@NOTES: Add support for AfterEffects point release scheme (2015.3)

#####

@RELEASE: 6.8-3c

#####

This is a patch release of core/supe/worker only, with some critical fixes to 6.8-3 for all platforms.

#####

==== CL 16389 ====

@FIX: calls to qb.reportwork that happen very close together can cause the supervisor to deadlock on a single frame's status

==== CL 16379 ====

@FIX: case-insensitive parsing of template names in qbwrk.conf when listed for template inheritance

The following now works (hostA will be in the "big" group):

[BigNode]

worker_groups = "big"

[hostA] : bignode

JIRA: QUBE-1809

==== CL 16369 ====

@FIX: don't mark the instance as failed if there is one more command to run, the child process has already exited, and the command is sys.exit(0); happens when maya is shut down with its native quit() function.

==== CL 16338 ====

@CHANGE: database checks script splits logging levels between stdout and stderr

==== CL 16286 ====

@FIX: checkDiskUsage fails when --mysql option is used and root can't authenticate

==== CL 16266 ====

@NEW: a new command-line utility for performing both database health checks and data integrity checks

==== CL 16247 ====

@FIX: fixed qb.workid() in callbacks to return the correct workid of the current callback context (it had been always returning None)

Also changed qb.jobstatus(), workstatus(), and subjobstatus() so that, if invoked in a callback giving no args (like a jobid and workid or subjobid),

they return the status of the respective thing (job, work, or subjob) of the current callback context.

JIRA: QUBE-1763
ZD: 16105

==== CL 16235 ====
@FIX: a problem with the filtering added to avoid jobs with an ID of 0, in CL15821

This was causing preemption to not function in many cases.

ZD: 16006

==== CL 16229 ====
@FIX: On Windows, daemons (supe, worker) now ignore VMWare Virtual Ethernet Adapters when trying to pick a primary mac address (QbConnection.cpp) for the host, which is used to uniquely identify hosts

ZD: 14481

==== CL 16214 ====
@FIX: aerender AppFinder mangling first path conversion on Windows when using UNC

==== CL 16064 ====
@FIX: when job 'dev' attribute True, printing the job package with regex_errors causes the logParser to generate a false positive for the regex_errors match

==== CL 16049 ====
@NEW: add 'outputPath match required' to python-based jobs, frame/work is failed if no match is found

#####

@RELEASE: 6.8-3

#####

==== CL 15964 ====

@NEW: changes to code that generates/modifies my.cnf

@CHANGE: some refactoring of the configure_mysql script (run on linux on (un)installation of the supervisor to modify my.cnf.

@NEW: make sure "default-storage-engine=MyISAM" is set on Linux too

@NEW: add "query_cache=0" to my.cnf on all platforms

JIRA: QUBE-1663

==== CL 15960 ====
@FIX: jobs submitted with pgrp set to a (null) string end up having a pgrp of 0

JIRA: QUBE-1668

==== CL 15957 ====
@FIX: use of single-quotes in job dependency "info-*" syntax results in hung job instances

JIRA: QUBE-1571

==== CL 15947 ====
@CHANGE: adding "default-storage-engine=MYISAM" to the my.cnf generated for Linux/OSX supe installations

JIRA: QUBE-1663

==== CL 15936 ====
@CHANGE: add InnoDB to MyISAM conversion code in upgrade_supervisor program for all "qube" tables

JIRA: QUBE-1664

==== CL 15909 ====
@CHANGE: change flaw in auto-wrangling logic in which it sometimes won't detect a bad worker, and allows it to fail many job agendas.

When a single job instance/worker has failed all of its assigned frames (at least aw_activation_work_count frames) for a job, while other workers are still processing their first frame (i.e., no other worker/instance has finished a frame), the system deems this worker "bad", locks it, and migrates the failed frames and instance, and notify the admin.

JIRA: QUBE-1475
ZD: 15219

==== CL 15865 ====
@CHANGE: Made section headers (such as "[default]" or "[node[001-199]]") case-insensitive in config files such as qbwrk.conf

JIRA: QUBE-1356

==== CL 15848 ====
@NEW: add Ubuntu 16.04 LTS support

==== CL 15821 ====
@FIX: add code to the DB routines and doPreemption() routine to silently ignore job records with job ID of 0 (likely due to corrupt DB records), which was spewing out many warning messages into the supelog

ZD:15739

==== CL 15809 ====
@FIX: backslashed characters in VRED jobs get treated as escape characters

==== CL 15761 ====
@NEW: add CentOS 7.2 support

JIRA: QUBE-1482

==== CL 15700 ====
@NEW: add "--conf filename" option to supervisor to specify an alternate location and name for the qb.conf file

JIRA: QUBE-253

#####

@RELEASE: 6.8-2

#####

==== CL 15673 ====
@FIX: orphaned job processes left behind on Windows workers, especially when the proxy.exe program dies unexpectedly

ZD: 15518

==== CL 15653 ====
@FIX: setting jobss "pgrp" value prior to submission is ignored for all but the first job when submitting a list of jobs via a single call to the qbsubmit() API routine

JIRA: QUBE-1536
ZD: 15528

==== CL 15650 ====
@FIX: Explicitly setting "host.memory" in worker_resources broken on Linux

ZD: 15505
JIRA: QUBE-1531

==== CL 15642 ====
@FIX: Unix (Linux/OSX) workers, when running a cleanup process for a terminating job instance (via removeJob()), would sometimes inadvertently kill processes belonging to other job instances, due to process IDs once owned by the terminating job being reused by the system.

ZD: 15548

==== CL 15587 ====
@FIX: cmdline and cmdrange jobtypes don't report the jobtype version in the job logs

==== CL 15567 ====
@FIX: supervisor_default_max_cpus value was not being applied properly

ZD: 15503
JIRA: QUBE-1528

==== CL 15560 ====
@CHANGE: "modify" operation will print, into the supelog and the job's .hst file, the values of the newly modified parameters

JIRA: QUBE-1318
ZD: 14979

==== CL 15555 ====

@FIX: prevent "upgrade_worker --reset" from printing out "table does not exist" error message.

JIRA: QUBE-817

==== CL 15531 ====

@NEW: add [run_program_and_convert_encoding.pl](#) script, which is a wrapper to run any given program and convert its stdout from and to specified encodings (like UTF-16le to UTF-8).

Added to support 3dsmax batch (i.e., "cmdrange") submissions.

JIRA: QUBE-1210

#####

@RELEASE: 6.8-1a

#####

==== CL 15462 ====

@FIX: removed submission-time check for jobtype existence on the farm, as it was causing false negatives in certain cases and disallowing submissions

ZD: 15328, 15831

#####

@RELEASE: 6.8-1

#####

==== CL 15384 ====

@NEW: add Mac OS X 10.11, aka "El Capitan" support

==== CL 15380 ====

@CHANGE: modification now allowed on "done" jobs

ZD: 15281

==== CL 15347 ====

@FIX: Windows issue where wireless network interfaces are ignored when licenses are verified, causing license keys bound to such interfaces to not work.

#####

@RELEASE: 6.8-0

#####

==== CL 15324 ====

@CHANGE: supervisor on Win32 to build against Perl 5.8 (upgraded from 5.6) to avoid build issues on new build platform.

==== CL 15154 ====

@CHANGE: supervisor now rejects workers that have newer major/minor version than itself.

Such workers will essentially stay in "down" state, or never appear in the host list.

JIRA: QUBE-1341

==== CL 15137 ====

@FIX: Windows qbsevice tool to back up existing my.cnf file before writing a new one when invoked with the "--mysqlprepare" option (i.e., via the supervisor installer)

For consistency with the Mac OS X supe installer, the back up file is named "mysql.qubebak.\$\$" where \$\$ is the current process ID (pid).

JIRA: QUBE-1229

==== CL 15077 ====

@NEW: add bin/qbdeleteworkerresources and qbdeleteworkerproperties programs

==== CL 15053 ====

@NEW: Basic admin UI for central prefs

==== CL 15052 ====

@CHANGE: automatically adjust host.processors of all jobs on farms with Designer licensing to 1.

==== CL 15048 ====

@FIX: "ERROR: unable to contact worker." - checkDiskUsage.py throws error when run on a machine which is not running as a worker.

==== CL 15014 ====

@FIX: fixed Python API docstring for deleteworkerresources and deleteworkerproperties

JIRA: QUBE-1322

==== CL 15011 ====

@CHANGE: allow "retry" of "badlogin" jobs (attempts to change their status to "pending")

JIRA: QUBE-642

==== CL 14948 ====

@FIX: "scoped" global resources aren't being tracked in the data warehouse

==== CL 14923 ====

@FIX: decrease the frequency of reporting progress and errors

@CHANGE: only do a file size check on the first 5 frames in a chunk

@FIX: setting fileSizeMin validation size to 0 disables the size checking.

==== CL 14919 ====

@FIX: log parsing not finding any matches in stderr, only stdout

==== CL 14751 ====

@CHANGE: decrease sampling and polling intervals to allow for consecutive fast-running commands to complete quicker, cuts down on application startup time for some apps

==== CL 14750 ====

@CHANGE: python job classes can take option 'prototype' arg in the constructor

==== CL 14749 ====

@CHANGE: child_bootstrapper for python loadOnce jobs is passed in as an argument, allows for application-specific bootstrappers

==== CL 14702 ====

@FIX: add code so that python27.zip is also added to 64-bit supe MSI builds

JIRA: QUBE-1228

==== CL 14698 ====

@NEW: adding python27.zip to be shipped with supervisor's MSI package

JIRA: QUBE-1228

==== CL 14691 ====

@FIX: add code to properly load python 2.7 modules shipped with the supervisor, in python27.zip (which contains files from Python 2.7.10 distribution)

==== CL 14657 ====

@FIX: add missing python27.dll file to supervisor MSI package

JIRA: QUBE-1228

==== CL 14581 ====

@CHANGE: changed ("new") worker behavior when auto-mount drives are unmountable due to duplicate drives.

Now, failed attempts to auto-mount a drive due to the drive letter already being in use will only generate a WARNING message in the workerlog, instead of rejecting the job and sending it back to the supe as "pending".

==== CL 14579 ====

@CHANGE: add more useful info to print to the workerlog when a job is rejected due to duplicate drive mounting (attempt to mount to a drive letter that's already mounting something else)

==== CL 14574 ====

@FIX: Secondary jobs were being dispatched even when supervisor_smart_share_mode is set to NONE

ZD: 14613

==== CL 14528 ====

@FIX: issue when modifying job's "env": "cwd", "umask", and "drivemap" are wiped-- additional fix to allow "env" modification of multiple jobs with a single call to qbmodify()

See also CL14516.

JIRA: QUBE-1161
ZD: 14549

==== CL 14523 ====
@CHANGE: upgraded supervisor's embedded Python to version 2.7.2 on Windows

JIRA: QUBE-1164

==== CL 14518 ====
@CHANGE: worker_boot_delay defaults to 10 seconds on workers running in service mode, on ALL platforms.

JIRA: QUBE-989

==== CL 14516 ====
@FIX: issue when modifying job's "env": "cwd", "umask", and "drivemap" are wiped

JIRA: QUBE-1161
ZD: 14549

==== CL 14514 ====
@FIX: add agenda item (aka "work") status to print properly to the job's history log when it's recalled, because of the instance that's processing being migrated, interrupted, failed, killed, or blocked.

There will be a line like the following in the .hst history log:

```
[Sep 15, 2015 17:09:05] 495670145 work 45765 1 __QUBE_SYSTEM__@supervisor recalled in supervisor by user[] from host[supervisor] on host[shinyambp] (127.0.0.1)
```

Note that this will also show, as expected, when a job instance reaches timeout (if specified) and "failed" by the system.

JIRA: QUBE-829
ZD: 13521

==== CL 14507 ====
@FIX: issue where subst mounted local drives will disappear from Explorer after a job finishes on DU mode workers.

@FIX: also fixed a bug where an already-mounted network/subst drives weren't being detected properly

ZD: 14009
JIRA: QUBE-1030

==== CL 14500 ====
@FIX: issue where cmd* jobtype jobs fail when paths given to QB_CONVERT_PATH() include parentheses

Note: problem was with the command-line tokenizer, QbExpressions::commandtokenize() routine, commonly used by all cmd* jobtypes, not respecting double-quoted and single-quoted strings.

JIRA: QUBE-1139

==== CL 14479 ====
@FIX: QB_CONVERT_PATH() runtime path conversion fails when the path to be converted contains parentheses

==== CL 14473 ====
@FIX: Allow custom algorithms to decide how to preempt SmartShare secondary instances, or just default to using value set in supervisor_smart_share_preempt_policy.

Custom algorithms may define a qb_preemptcmp_secondary() routine to control how secondary jobs are preempted.

ZD: 14472
JIRA: QUBE-1145

==== CL 14406 ====
@FIX: fixed missing job parameters in the job object returned by "qbjobobj()" (qb.jobobj() in python) in jobtype backends.

The following parameters were added:

- queue
- max_cpus
- omithosts
- omitgroups
- notes
- cpustally
- tototally
- automigratecount
- retrysubjob
- retrywork
- retrywork_delay

dependency
mailaddress
sourcehost
prod_show
prod_shot
prod_seq
prod_client
prod_dept
prod_custom1
prod_custom2
prod_custom3
prod_custom4
prod_custom5

==== CL 14397 ====

@FIX: performance tweak, cut down on the number of times backends and automated scripts fetch the supervisor config

==== CL 14360 ====

@CHANGE: agenda-based job instance is immediately interrupted, even if the global preemption policy is set to passive, if it hasn't started processing an agenda item

JIRA: QUBE-1077

ZD: 14109

==== CL 14352 ====

@FIX: added QB_FRAME_NUMBER, QB_FRAME_START, QB_FRAME_END, QB_FRAME_STEP, and QB_FRAME_RANGE to be defined in the environment just before a frame is processed

ZD: 14203

==== CL 14326 ====

@FIX: make appropriate invocation of approvemodify(qb_approvemodify()) perl routine) for Custom Policy

ZD: 14173

JIRA: QUBE-1082

==== CL 14320 ====

@FIX: catch case in checkUserPermission where traceback error "e" is not defined and an attempt is made to report the error message - occurs when user running the script is not a qube admin

==== CL 14305 ====

@TWEAK: print queuing policy (Internal or custom/Perl) message to supelog

==== CL 14273 ====

@FIX: properly report back failing status when an regex_error is matched early on, but then not found in the last pass through the logs.

==== CL 14207 ====

@FIX: log sections that match an error regex from before an auto-retry are being scanned and matching for errors; now either "qube! - retry/queue" or "auto-retry" messages trigger a reset

==== CL 14204 ====

@NEW: a script and modules to sync external 3rd-party license server counts with Qube's global resources

@NEW: first external license server modules are for FLEXlm and sesinetd servers

==== CL 14191 ====

@NEW: add path translation to all python-based loadOnce jobtypes

==== CL 14162 ====

@FIX: issue where the supervisor, when starting secondary instances for a job, can preempt more instances than necessary-- i.e., preempt more instances than there are agenda items for the job.

ZD: 13969

JIRA: QUBE-1007

==== CL 14064 ====

@FIX: issue where global time-based callbacks (i.e., "dummy-time-self" callbacks) sometimes not triggering

ZD 13366

JIRA: QUBE-807

==== CL 13971 ====

@CHANGE: add job "name" and "lastupdate" columns to be added at time of job ID creation (available while job is still in "registering" state).