

Qube 6.3-6 Release Notes

```
#####  
@RELEASE: 6.3.6  
#####
```

==== CL 10514 ====

@FIX: another patch for out-of-order issue. Fixed unexpected short-circuit evaluation that was happening in the startResources() routine

==== CL 10513 ====

@FIX: another patch for out-of-order issue. Fixed unexpected short-circuit evaluation that was happening in the startHost() routine

==== CL 10512 ====

@INTERNAL: QbJob object's _subjobswaiting data was not being initialized or copied correctly, causing some job comparisons based on subjobs waiting counts to unexpectedly fail.

==== CL 10504 ====

@INTERNAL: added more log output for debugging builds, added more comments while working on out-of-order issue.

ZD: 8198

==== CL 10477 ====

@FIX: Another out-of-order fix. Jobs at the same numerical and cluster priority should dispatch in the correct FIFO order now.

The FIFO enforcing should work most of the time, but there still will be occasional out-of-order behavior, due to the multi-threaded nature of the supervisor. ("qbshove"-ing the older job should correct it, when it's seen)

ZD: 8198

==== CL 10462 ====

@FIX: yet yet another fix for out-of-order dispatch behavior-- eliminate race-condition that would allow lower priority jobs that were just preempted to get workers before higher-priority jobs.

See also CL10440 10452

ZD: 8198

==== CL 10461 ====

@CHANGE: modified/compacted the multi-line "found a duty to replace" logging to be a single line.

==== CL 10452 ====

@FIX: yet another fix for out-of-order dispatch behavior-- eliminate race-condition that would allow lower priority jobs that were just preempted to get workers before higher-priority jobs.

See also CL10440

ZD: 8198

==== CL 10441 ====

@FIX: killing an already finished (complete, failed, killed) job leaves the job in the "dying" state.

==== CL 10440 ====

@FIX: another fix for out-of-order dispatch behavior-- eliminate race-condition that would allow lower priority jobs that were just preempted to get workers before higher-priority jobs.

ZD: 8198

==== CL 10429 ====

@FIX: out-of-order job dispatching issue with jobs using the "+" sign with the "host.processors" reservations.

ZD: 8198 8261 8229 8233 8228

==== CL 10189 ====

@FIX: timing issue where some worker resources (host.xyz) would disappear after the worker received a remote config.

@FIX: issue where supervisor tries to dispatch a subjob to a worker with insufficient resources (reduced the likeliness of that from happening)

@FIX: the above 2 fixes combined should now prevent some of the out-of-priority-order dispatch issues, especially in environments where worker resources are deployed.

ZD: 7885

==== CL 10118 ====

@FIX: fixed issue where agenda timeouts don't work properly on the first agenda item processed by a subjob, on Unix (Linux/OSX) workers

==== CL 10117 ====

@FIX: fixed issue where agenda items that fail because of timeout don't get automatically retried via retrywork

ZD: 7763

==== CL 10022 ====

@FIX: modified the worker to only report to the supe of its host status when subjobs are completely done and removed, and NOT when they are only marked/scheduled for removal.

This was causing jobs to sometimes run out-of-order, especially when there are many subjobs to each job (such as one subjob per frame), since that situation tends to increase the chance of the supervisor dispatching the same subjob to the same worker. The subjob will be dispatched to the same worker, but rejected since the worker thinks it's a duplicate assignment of a subjob that's being removed (and consequently a lower priority job will get the worker's slot, causing out-of-order job execution)

ZD: 7601

==== CL 9903 ====

@FIX: better message from worker when it rejects a dispatched subjob because it's a duplicate (being preempted or migrated on the same worker)

==== CL 9838 ====

@CHANGE: upped the default value for supervisor_max_threads to 100, and worker_max_threads to 32