

# TORQUE Resource Manager 5.1.2 Release Notes

The release notes file contains the following sections:

- [New Features on page 2](#)
- [Differences on page 4](#)
- [Known Issues on page 7](#)
- [Resolved Issues on page 8](#)

## New Features

The following is a summary of key new features in TORQUE Resource Manager.

### 5.1.2

#### *Added Feature to Disable Editing pbsnodes File*

The qmgr command "dont\_write\_nodes\_file=true" disables all commands which edit nodes. If set to true, this setting locks all access to the nodes file, thereby ensuring the nodes file is never overwritten unintentionally.

#### *Able to Exclude Compute Nodes from allow\_node\_submit*

A new TORQUE server parameter "node\_submit\_exceptions" is available to specify nodes that are not allowed to submit jobs. This parameter is set in conjunction with "allow\_node\_submit".

#### *Ability to Trust Certain Users or Groups from Hosts without Allowing All Users from those Hosts to Submit Jobs*

Two new TORQUE server parameters "acl\_user\_hosts" and "acl\_group\_hosts" are available to support this new validation.

### 5.1.1

#### *Support for more OSs*

These additional OSs are now supported:

- CentOS 7.x
- RHEL 7.x
- Scientific Linux 7.x
- SUSE Linux Enterprise Server 12

**i** Support for Red Hat 7-based and SUSE 12-based systems requires the 5.1.1.2 maintenance release or later.

### 5.1.0

#### *Ability to Provide Condensed qstat Output*

A 'qstat -C' option, which specifies that TORQUE will provide only a condensed output (job name, resources used, queue, state, and job owner) for jobs that have not changed recently (as per the job\_full\_report\_time parameter), has been added. Jobs that have recently changed will continue to send a full output.

#### *Performance Enhancements to MOM Clean-up Time*

## New Features

Some minor performance enhancements were made to improve MOM clean-up time.

## Differences

This section contains differences in previously existing features that require a change in configuration or routine.

### 5.1.2

#### *down\_on\_error Server Parameter Now Defaults to TRUE*

By default, nodes that report an error from their node health check to pbs\_server will be marked down and unavailable to run jobs.

#### *Abnormal Job Exit States Combined Into Single Exit Code*

Added parameter "exit\_code\_canceled\_job" to force all canceled jobs to have the same exit code, regardless of exit state. Previously, exit codes would be inconsistent depending on the type of cancellation or abnormal exiting of each job. The "exit\_code\_canceled\_job" gets rid of extraneous exit codes and allows users to set their own exit code. For example, if a user sets "exit\_code\_canceled\_job = 300" in qmgr, the exit code for any job that exits with a non-zero status becomes 300.

#### *Qmgr Support Added for "loglevel" Attribute*

The "loglevel" attribute name was not recognized by qmgr, while it was recognized in momctl; the string "loglevel" is now equivalent to "log\_level" in qmgr. The user can now use either of these attribute names to check the current value of the loglevel.

#### *Node Health Checks Now Run on Sister Nodes*

Previously, node health checks would exclusively run on the mother superior. Node health checks now run on sister nodes as well.

#### *Submission Syntax Added to Prevent Mixing NCPUs and Nodes*

When mixing NCPUs and nodes, the checkjob command would show available procs as a negative; this mixing is now prohibited.

### 5.1.1

#### *\$prologalarm is Always Honored*

\$prologalarm was ignored on the prologue for a job. Also when the epilogue was run the \$prologalarm value was ignored if it was more than 300. Now the \$prologalarm value is always honored regardless of how large it is for both prologue and epilogue scripts. The default timeout is still 300 seconds.

#### *Disable the Automatic Requeuing of Jobs*

Added the ability to disable the automatic requeuing of jobs due to transient failures.

#### *pbs\_mom now sets environment variable for NVIDIA GPUs*

A new mom config parameter, `$cuda_visible_devices`, was added to specify whether `pbs_mom` sets the `CUDA_VISIBLE_DEVICES` environment variable when it starts a job. The default is `TRUE`.

#### *Log Milliseconds*

Added milliseconds in TORQUE's log files.

#### *pbs\_server Enhancement for Very Large Number of Jobs*

`pbs_server` has been enhanced to better handle a very large number of jobs (several hundred thousand or more) by enabling an alternate way for it to store job-related files in the directories `$PBS_HOME/server_priv/jobs` and `$PBS_HOME/server_priv/arrays`.

A new boolean server attribute, `use_jobs_subdirs`, lets an administrator direct the way `pbs_server` will store its job-related files. When `use_jobs_subdirs` is unset (or set to `false`), job and job array files will be stored directly under `$PBS_HOME/server_priv/jobs` and `$PBS_HOME/server_priv/arrays`. This is the default behavior and the way the server has stored these files in the past. When `use_job_subdirs` is set to `true`, job and job array files will be distributed over 10 subdirectories under their respective parent directories. This method helps to keep a smaller number of files in a given directory.

If an administrator wishes to change the `use_jobs_subdirs` attribute from its previous value (or when setting it to `true` when it has not previously been set), it is highly recommended that TORQUE be drained of all jobs. Failing to take this action may result in the loss of existing jobs.

#### 5.1.0

#### *CLIENTRETRY Configuration Option Support for qdel*

The `qdel` parameter now includes the `-b` option and `CLIENTRETRY` configuration option support. This feature functions similar to `-b` option for `qsub`. Specifically:

- `-b <num>` command line argument support
- `CLIENTRETRY` configuration option support
- `PBS_CLIENTRETRY` environment variable support

#### *Jobs Deleted When the Dependency Can No Longer Be Satisfied*

When a job is deleted because its dependency can no longer be satisfied, that job will follow the `keep_completed` parameter set (if any) for `pbs_server`. Previously, jobs were purged immediately.

#### *pbs\_server -t No Longer Supports hot|warm|cold Options*

The `pbs_server -t` option no longer supports the `hot|warm|cold` options. Other options are still supported.

## Known Issues

The following are known issues in TORQUE Resource Manager. Following each issue description is an associated issue number in parentheses. Known issues are aggregated and grouped by the release version for which they first occurred or were reported.

### 5.1.2

No known issues.

### 5.1.1

- Jobs do not exit the routing queue if the job is submitted with a hold (e.g. `qsub -h`). (TRQ-2788)

### 5.1.0

- ~~DNS caches are not purged of removed nodes when Elastic Computing is enabled. SLES 11 SP1 has an issue with giving the old IP address to TORQUE even after updating `/etc/hosts`. Do not use SLES SP1 if you are using this method to manage the IP addresses for `pbs_server`. The mom's OSs are irrelevant. (TRQ-2765)~~ This issue is resolved in 5.1.1.
- Some limitations exist in the way that `pbsdsh` can be used. Please note the following situations are not currently supported:
  - Running multiple instances of `pbsdsh` concurrently within a single job. (TRQ-2851)
  - Using the `-o` and `-s` options concurrently; although requesting these options together is permitted, only the output from the first node is displayed rather than output from every node in the chain. (TRQ-2690)

## Resolved Issues

The following is a list of some key bugs fixed in TORQUE Resource Manager. Following each issue description is an associated issue number in parentheses.

Resolved issues are aggregated and grouped by the release version in which they were resolved.

### 5.1.2

- **Server timed out to MOM when running a lot of small jobs.** Added load balancing to login nodes when login nodes begin to get busy. (TRQ-3367)
- **Numbered directories in server\_priv/jobs (and arrays) were missing when the server attribute "use\_jobs\_subdirs" was set to TRUE.** This issue is fixed. (TRQ-3185)
- **Random group names displayed for users that did not belong in the group.** A race condition was fixed by changing to thread safe calls to get group and user ids. (TRQ-3190)
- **Read timeouts were retried indefinitely by pbs\_server.** This issue is fixed. (TRQ-3306)
- **End-of-job times were not consistent.** End-of-job times are now recorded closer to when nodes are free for use. (TRQ-2840)
- **Torque init scripts were not LSB compatible.** The pbs\_mom service was not returning exit code 3 when exiting, which is an LSB requirement. As a result, the builder was failing to properly initialize pbs\_mom, causing nose tests to fail. Init script return codes are now altered to correct values for LSB compatibility. (TRQ-3254)
- **Reporter MOM was unable to handle ALPS UNKNOWN role.** This issue is fixed. (TRQ-3245)
- **Jobs were stuck in EXITING state on mother superior.** This issue is fixed. (TRQ-3165)
- **Jobs with square brackets "[]" in their name were aborted on restart if they were not array subjobs.** Torque now allows all jobs that have "[]" in their name. (TRQ-3214)
- **Improperly escaped xml caused job files and dependencies to fail upon reload attempt.** This issue is fixed. (TRQ-3284)
- **Jobs were not created properly if the path had one or more spaces in the name.** Added multiple spacing options for job file path name. (TRQ-3235)



- **Completed jobs were not cleaned up properly.** Fixed various issues relating to restarting dependency jobs. (TRQ-3175)
- **Job files were never deleted when "\$thread\_unlink\_calls" value was set to TRUE.** Fixed issue where the threadpool in pbs\_mom was not started properly causing job files to be left behind. (TRQ-3232)
- **Some threads would cause Torque to hang and become completely unresponsive.** Count\_proc() now returns 0 on a NULL argument. (TRQ-3196)
- **Jobs were run in the wrong order and canceled prematurely.** Running jobs are no longer deleted due to dependencies. (TRQ-3189)
- **Server entered deadlock around completed jobs map.** This issue is fixed. (TRQ-3226)
- **Flooding server with client commands disabled some functionality.** Resolved issue where server gets stuck due to large volume of client requests. (TRQ-3296)
- **Kill\_delay sent multiple signals.** Signals sent to job pids are now tracked to eliminate redundant sending. (TRQ-3239)
- **Qsub -W stage-in displayed incorrect user grouping.** The log was showing users belonging to groups that the users were not a part of. This issue is fixed. (TRQ-3312)
- **pbs\_submit() method changed for stable job submissions.** The qsub command currently utilizes the pbs\_submit\_hash() method to submit jobs and is stable. Submitting jobs via the pbs\_submit() method would occasionally render pbs\_server unresponsive. Now pbs\_submit() is functioning properly. (TRQ-3314)
- **Moab crashes in Torque API.** This issue is fixed. (TRQ-3368)

#### 5.1.1

- **Queue default resources were being modified.** Corrected an issue where the all memory values were converted to bytes after a restart. (TRQ-3139)
- **Some completed jobs were taking quite a while to clear out.** Hardened the removal of completed jobs from pbs\_server. (TRQ-3044)
- **Ability to have the procct values displayed on the queues where it was set was lost.** This issue is fixed. (TRQ-3135)
- **max\_user\_queueable issues.** Fixed some bugs causing incorrect max\_user\_queueable values. (TRQ-2841)
- **PBS\_NUM\_NODES on multi-req job wrong when using specific hostnames.** This issue is fixed. (TRQ-1949)

- **Dependency jobs fail to start during HA fail over.** Dependencies now display in the format of `depend=type:jobid` instead of `depend=type:jobid@server`. (TRQ-2332)

```
# Was
depend=afterok:12.napali@napali

# Now
depend=afterok:12.napali
```

- **Bug reported that caused jobs to not start when ALPS incorrectly returns a permanent confirmation failure.** This issue is fixed. (TRQ-3023)
- **Problem reported that caused MOM restarts to intermittently fail.** This issue is fixed. (TRQ-2307)
- **TORQUE will not compile with --enable-debug on configure.** This issue is fixed. (TRQ-2969)
- **pbs\_server would count completed jobs against queue limits when pbs\_server was restarted.** This issue is fixed. (TRQ-3087)
- **Compute node offline state doesn't persist across server restarts.** This issue is fixed. (TRQ-2790)
- **With kill\_delay and \$exec\_with\_exec set, a job would be set to a completed state after running qrerun instead of getting set back to queued.** This issue is fixed. (TRQ-2993)
- **qmgr refuses numerical node ids.** This issue is fixed. (TRQ-2946)
- **X11 very slow with qsub -I -X.** The port forwarding buffer size has been increased to improve performance when enabling X11 forwarding from qsub (-X switch). (TRQ-2740)
- **Problem reported with interactive jobs not staying on the node from which they were submitted.** This issue is fixed. (TRQ-3122)
- **Jobs were getting stuck in a running state when an asynchronous run failed.** This issue is fixed. (TRQ-3114)
- **Array slot limits were not getting decremented when a job is preempted or rerun.** This issue is fixed. (TRQ-3110)

#### 5.1.0

- **HOST\_NAME\_SUFFIX was no longer adding suffix to job names.** This issue is fixed. (TRQ-2956)
- **pbs\_mom filling up the logs in a HA environment.** Reduced verbosity in error logging in HA environments. (TRQ-2863)
- **Delays reported with Preemption.** Speed up process data collection to help diminish time to rerun a job, including correctly sum job memory and cpu time resources for processes created under `setsid()` calls within a job,

and avoid an infinite loop in pbs\_mom when running under a Linux Container (lxc). (TRQ-3026)