

Torque Resource Manager 6.0 Release Notes

Revised December 2015

The release notes file contains the following sections:

- [New Features](#)
- [Differences](#)
- [Known Issues](#)
- [Resolved Issues](#)

New Features

This section contains a summary of key new features.

6.0.0

cgroup Support

Torque is enhanced to create one Linux control group (cgroup) per task based on the new NUMA-aware, task-based job submission option (-L) and to create one cgroup for all tasks of a job on the same compute node for the older job-based option (-l). Torque uses cgroups to manage CPU and memory accounting, enforce memory usage limits, set up Cpuset management, and bind cores/threads, memory, and accelerators, such as GPUs and MICs, to jobs.

When binding resources that include an accelerator to a task, Torque will make a best-effort attempt to place a task on the cores/threads and memory of the socket/NUMA node to which the accelerator attaches.

Ability to Prevent Nodes Being Dynamically Edited

A new qmgr parameter is available. When 'dont_write_notes_file' is set to true, the nodes file cannot be overwritten for any reason; qmgr commands to edit nodes will be rejected. The default is FALSE.

Execute the Job Starter Script with Elevated Privileges

The '\$job_starter_run_privileged' MOM configuration parameter is added and lets you specify whether Torque executes the job starter script with elevated privileges. The default is FALSE.

Differences

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

6.0.0

Default RPM Installation Path Is Changed

The Torque default path for an RPM installation has been changed to match the path used during a tarball (Manual) installation. The default path for both install methods is `/usr/local`.

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.

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`.

\$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 the 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.

Known Issues

This section lists known issues. Known issues are aggregated and grouped by the release version for which they were first reported. Following each issue description has the associated issue number in parentheses.

6.0.0

Running multiple instances of pbsdsh concurrently within a single job is not supported.

pbsdsh will fail to return under certain conditions (not-passing high-stress tests). *Resolved 6.0.0.1*

Kernel crashes may occur when using cgroups on CentOS or RHEL prior to 6.6. See https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/6/html/6.6_Technical_Notes/kernel.html; especially RHEL6.6 fix BZ#1204626. If cgroups are part of your configuration, Adaptive Computing recommends running a more recent version of CentOS or RHEL.

Resolved Issues

This section lists resolved issues. Resolved issues are aggregated and grouped by the release version in which they were resolved. When applicable, each resolved issue has the associated issue number in parentheses.

6.0.0.1

- A hang in pbsdsh occurred if the pbs_mom daemon was started with a -q or -r option. (TRQ-3308)
- Array templates were being reported as jobs. (TRQ-3405)
- Typo found in the error message reported when the swap memory limit could not be set.

6.0.0

- 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. (TRQ-2993)
- Array slot limits were not getting decremented when a job is preempted or rerun. (TRQ-3110)
- Jobs were getting stuck in a running state when an asynchronous run failed. (TRQ-3114)
- Interactive jobs not staying on the node from which they were submitted. (TRQ-3122)
- Occasionally a random group name would show up for a user who 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)
- When \$thread_unlink_calls is set to true in /var/spool/torque/mom_priv/config, job files were not being deleted at job end in the mom; threadpool in pbs_mom was not being started. (TRQ-3232)
- Reporter mom did not correctly handle UNKNOWN role. (TRQ-3245)
- Read timeouts were being retried indefinitely by pbs_server. (TRQ-3306)