

Torque

Why are workload queries timing out and job completions not being detected?

SYMPTOMS

When using Torque, the workload query times out repeatedly, and after that Moab no longer sees job completions or jobs submitted through qsub. Raising the timeout seems to make Moab stop responding. Additionally Moab may see jobs still running that have actually run well beyond their walltimes.

CAUSE

There was too much information to complete in one timeout period. This can be due to a lot of job-specific activity (for example, a large job array failed or was cancelled). After the first query times out, the situation will only get worse, as each successive query still has that original, plus additional information, so they just continue timing out. Since Moab no longer has job updates, it can't detect job completions, nor new "qsub"-submitted jobs.

In addition, bumping up the timeout may make Moab appear to have stopped responding. What's likely happening, though, is Moab is simply busy trying to process all of the information. Customers often will restart Moab, which does not actually resolve anything.

SOLUTION

The solution for this is to bump up Torque's timeout until the logs show the query completed, and then just wait for Moab to process all of the back information. The timeout is adjusted by the "TIMEOUT=<value>" option on the RMCFG line for Torque (eg., assuming RM's name is torque: RMCFG[torque] TIMEOUT=60).

Having the longer timeout will not impact normal operations, so you can leave it at a higher value if desired. If it used to be 30 seconds (the default), for example, you might want to go to 45 or 60.

Then monitor the "moab.log" file (tail -F \$MOABHOMEDIR/log/moab.log) and make sure Moab is really doing something and not actually hung. Note that by using the upper-case "-F", "tail" will follow the file by name, not by inode, so if the logs rotate you should still see the latest. If no activity is being logged, then Moab is likely hung and you need to look for other problems.

Unique solution ID: #1230

Author: Rob Greenbank

Last update: 2020-12-11 19:37