

Moab

What to check after a crash.

Question: Our machine running Moab and Torque recently crashed. Do I need to take any action before bringing moab and pbs_server back up to verify that things are ok?

Answer: Torque keeps data about jobs that are on the system in /var/spool/torque/server_priv/jobs/ and /var/spool/torque/server_priv/arrays for array jobs. After a crash you should make sure these directories still have job data in them. For non-array jobs you should see a .JB and .SC file for each job. Array jobs will have an .AR extension. If there are problems with a job there will be a file with a .BD extension.

If you submit jobs with msub then the directory you would look in for job files is /opt/moab/spool/. The files would have .cp extensions instead. In this case there could still be files for the same jobs in the torque directory because moab defaults to sending the job data to the resource manager as soon as it knows about them. You can change this behavior too by setting the migrate policy to "JUSTINTIME", in which case the idle jobs would have files in the moab directory and the active jobs would be in torque. You probably won't have a list of jobs that were there before the crash to make sure they're all accounted for, but it's worth verifying that these directories contain roughly the number of files you'd expect for the number of jobs on your system.

If moab or pbs_server crashed rather than the machine itself then you should also look for core files in /var/spool/torque/ and /opt/moab/. These should help support quite a bit in figuring out the cause of the crash.

When bringing things up after a crash you should start out with pbs_server and then start trqauthd. Then run qstat to make sure Torque sees a list of jobs that looks reasonable for the amount of work you expect to see. Then run 'pbsnodes' to make sure the nodes are being recognized. If that looks good then start moab and it should resume scheduling.

If moab and torque are down because of a crash it shouldn't have too much impact on the running jobs. The moms will continue to run the jobs whether or not the server is available. One problem would be that the moms can't report to the server that a job has completed. But when the server comes back up it will get an update from all the nodes about what is currently running on them and it should get correct data about which jobs are currently active. The other problem caused by the server being down is that users can't submit new jobs or have existing ones start.

Unique solution ID: #1139

Author: Ben Roberts

Last update: 2016-10-25 17:08