

Torque

How to migrate queued jobs to a new Torque server with a different host name

Issue:

We need to migrate our queued Torque jobs from one server to another one with a new host name.

Solution:

Disclaimer: the steps provided here have not been tested, so you will first want to try this on test jobs.

The process for migrating Torque essentially entails doing a fresh installation on the new server, copying over the files from the old server, and then updating `server_name`, `serverdb`, and the `.JB` and `.AR` files. Adaptive does not have a supported tool to perform this task. This article suggests one possible method for doing so. This method is not officially supported, although Adaptive Support will make every reasonable effort to assure your success.

There are a number of ways you can accomplish this. Obviously, you can simply offline nodes (with `"pbsnodes -o <node ID>"`), a few at a time, and then remove them from the old cluster and add them to the new one, once all jobs drain from the offlined nodes.

As far as retaining queue jobs:

The following method differs slightly, and is mostly straightforward (but still involved):

1. Install the new "vanilla" Moab and `pbs_server`, and make sure they're both running and talking to each other.

– You can temporarily run `pbs_mom` on the server and add it to the nodes file, just to make sure a node shows up in `pbsnodes` and `checknode/mddiag -n`.

– Once you've confirmed they're running properly, shut them both down.

2. To begin migrating Torque, it's probably easiest to copy the entire `TORQUE_HOME` directory over to the new server and overwrite `TORQUE_HOME`, but first back up the new `serverdb` file (and/or the entire new install directory).

3. Reconcile the old `serverdb` file copy with the new backup copy, making any appropriate corrections.

4. Delete any files in `server_priv/jobs` with running, exiting, or completed states (as indicated by the `"<state>#</state>"` tags). Here's the list of codes from `pbs_job.h`:

```
/* sync w/PJobState[] */
```

Torque

```
#define JOB_STATE_TRANSIT 0
#define JOB_STATE_QUEUED 1
#define JOB_STATE_HELD 2
#define JOB_STATE_WAITING 3
#define JOB_STATE_RUNNING 4
#define JOB_STATE_EXITING 5
#define JOB_STATE_COMPLETE 6
```

Example:

```
# # Delete all files for running, exiting and completed jobs:
# cd /var/spool/torque/server_priv/jobs/
# find . -type f -exec grep -q '<state>[4-6]</state>' {} \; -delete
```

5. Rename the remaining .JB files, replacing the old server name with the new one (eg., "1234.oldservers" to "1234.newserver.JB").

6. Replace all occurrences of the old server name inside the .JB files with the new server name (for instance, with a tool like sed).

7. Repeat step #5 for all .AR file names under server_priv/arrays/.

8 Repeat step #6 for the contents of the .AR files.

9. Delete all .AR files that don't have queued jobs. One way to do this:

- get a list of .AR file names with "ls -l arrays/*.AR"
- look for matching "jobs/<JOBID>*.JB" files. For the ones with no hits, delete the .AR file.

Example:

```
[root@headnode server_priv]# ls -l arrays/*.AR | sed -e "s%arrays/%% ;
s%\\([1-9][0-9]*\\)\\.\\.%.%1-%" | xargs -l{} bash -c "ls -l jobs/{ }"
```

10. In the old Moab, create an (inaccessible) administrative or standing reservation covering a portion of the cluster. Just to be safe, I'd also pause scheduling in Moab by running "mschedctl -p" (and adding "DISABLESCHEDULING TRUE" to moab.cfg, in case Moab restarts at any point before the migration finishes).

11. On each node (as jobs drain from the compute hosts in the admin reservation), shut down the pbs_mom, update the server_name file, remove it from the old Torque, and add it to the new Torque (via the nodes file, or with "qmgr -c "create/delete node...""). (Obviously, this is an ongoing process until every node has been migrated over to the new server.) Also check for the old name in TORQUE_HOME/mom_priv/config, specified with a "\$pbsserver" line.

Torque

12. Back up TORQUE_HOME on the new server.
13. Start up pbs_server and make sure the jobs queue up. If they don't, you'll get a bunch of "JB.BD" files in the jobs directory.
14. For the MOMs, read the "Rolling Upgrade" section of Appendix E, "Considerations Before Upgrading"
15. Check the job submission filter on the server for any pertinent updates.
16. Also check torque.cfg anywhere you have client commands installed, if necessary.

Unique solution ID: #1107

Author: Rick McKay

Last update: 2018-01-19 19:41