

Moab

Why do I have incorrect job charges in Moab Accounting Manager?

Problem: Some of the charges in MAM are not reflecting the actual job run times, as reported by Torque. The "tracejob" output is showing the correct CPU utilization, but MAM has recorded different values.

Solution:

MAM will record exactly what Moab tells it to record for its job charges. When the charges are incorrect, it's due to Moab providing incorrect information. There are configuration settings that can cause Moab to record incorrect charges, and these are easily corrected.

There are four places in the Moab configuration where Moab is told what to do when it cannot contact MAM. These four are CONTINUEFAILUREACTION, CREATEFAILUREACTION, RESUMEFAILUREACTION, and STARTFAILUREACTION. Each of these configuration parameters has three options defining what Moab should do in each case:

- ConnectionFailureAction - there is a communication problem with MAM
- FundsFailureAction - there are insufficient funds to run the job
- GeneralFailureAction - MAM rejects the job for any other reason

The defaults for all of these are IGNORE,IGNORE,IGNORE. This is definitely not the best configuration to use, so these options should be changed.

Many sites do not allow job suspension, and for those sites the CONTINUE and RESUME options do not matter. It's still a good idea to configure them as shown below, however. In the case of the CREATE option, it's actually acceptable to ignore MAM connection errors, as the job is created and ready to be queued, but no job charges will have accrued. However, the CREATEFAILUREACTION policy will only be applied if the VALIDATEJOBSUBMISSION is set to True.

The biggest potential problem lies in the START option. With the flags set to IGNORE, Moab will go ahead and start a job if it cannot contact MAM or if MAM returns an insufficient funds error. With strict accounting, this means a job without any funds will be started anyway.

The defaults are all set to IGNORE, which tells Moab to proceed anyway. In a perfect cluster, this might never be an issue as Moab should always be able to connect to MAM. In the real world, however, there are situations where a connection to MAM might fail. This could be a network issue, or it could be MAM is unresponsive due to database access. The result could be jobs starting that have no funds.

For this reason, we recommend the following settings:

```
AMCFG[mam] CONTINUEFAILUREACTION=DEFER,CANCEL,DEFER
AMCFG[mam] CREATEFAILUREACTION=DEFER,CANCEL,DEFER
AMCFG[mam] RESUMEFAILUREACTION=DEFER,CANCEL,DEFER
```

Moab

AMCFG[mam] STARTFAILUREACTION=DEFER,CANCEL,DEFER

These settings will not cancel jobs when connections are a problem, but will instead defer them, allowing Moab to retry the connection during another iteration. Some sites prefer to place jobs with insufficient funds on hold, instead of canceling them, and this is also perfectly acceptable. Placing them on hold allows an administrator to add to the job's fund and un-hold the job.

Finally, there is one other MAM option that deserves close scrutiny. For very large sites, or sites that have frequent connection issues, the "TIMEOUT" value may need to be higher than the default 30 seconds. The timeout value for the Torque resource manager also needs to be set high enough to allow iterations to complete. If either of these timeouts are exceeded, what might happen is Moab fails to get the job information from Torque, and eventually it may process the job based on what information it has. If Moab is just starting after an outage, and it does not receive information about a job, it's possible it may cancel the job and use the time that happens instead of the resources the job actually used. The symptoms of this will be Torque reporting a job as completed, Moab showing it as canceled, and the charges will be based on the time of the cancellation as shown in the event logs (in Moab's "stats" sub-directory).

Unique solution ID: #1237

Author: Rob Greenbank

Last update: 2022-06-24 16:27