

Torque

My MPI jobs are failing to run when I run them through TORQUE however they run externally

Issue: My MPI jobs are failing to run when I run them through TORQUE however they run externally.

Affected Versions: All

Symptom:

MPI_Init: ibv_create_cq() failed
MPI_Init: Can't initialize RDMA device
MPI_Init: Internal Error: Cannot initialize RDMA protocol

MPI_Init: ibv_create_cq() failed
fMPI_Init: Can't initialize RDMA device
MPI_Init: Internal Error: Cannot initialize RDMA protocol
MPI Application rank <NN> exited before MPI_Init() with status 1
mpirun: Broken pipe

Solution:

When a batch job is created through TORQUE sometimes the jobs environments ulimits are inherited from the pbs_mom. Normally a jobs ulimits inherit the ulimits from the OS environment. In this case the "max memory size" (ulimit -m) is not sufficient for the mpi job to run correctly.

You can edit the init script for the pbs_mom to set the ulimit -m before pbs_mom is executed. This will allow the mom to have a mem limit greater then 64. Alternatively if you start your pbs_mom by hand you will set your current shells ulimit with the command (ulimit -m unlimited) before running "pbs_mom".

Unique solution ID: #1022

Author: Jason Booth

Last update: 2015-06-11 00:00