

Job Management 101

By Yongjun Choi, iCER Research Consultant

Listing of your jobs

To see a list of your submitted jobs, type **qstat -u username**

```
dev-intel16-k80:~$ qstat -u choiyj

mgr-04.i:

Job ID          Username      Queue  Jobname      SessID  NDS   TSK   Req'd      Req'd      Elap
             Memory      Time   S           Time
-----
51293689[ ]mgr-04.i  choiyj      main   ucp_gptest   --      1     1     20gb    04:00:00 Q    --
51293730.mgr-04.i  choiyj      main   hello.qsub   --      1     1     200mb   01:00:00 Q    --
dev-intel16-k80:~$
```

If you need more detailed information, you can type **showq -u username**
 (It may take up to about 30 seconds for a job to appear in a **showq** output)

```
dev-intel16-k80:~$ showq -u choiyj

active jobs-----
JOBID          USERNAME      STATE  PROCS   REMAINING      STARTTIME
51293689(15)   choiyj      Running  15     3:59:05  Mon Jan 15 12:10:42

15 active jobs          15 of 16264 processors in use by local jobs (0.09%)
                        0 of 625 nodes active          (0.00%)

eligible jobs-----
JOBID          USERNAME      STATE  PROCS   WCLIMIT      QUEUE TIME
0 eligible jobs

blocked jobs-----
JOBID          USERNAME      STATE  PROCS   WCLIMIT      QUEUE TIME
51293689(179)  choiyj      Blocked  179     4:00:00  Mon Jan 15 12:07:30
51293730       choiyj      Idle     1       1:00:00  Mon Jan 15 12:10:14

180 blocked jobs

Total jobs: 195

dev-intel16-k80:~$
```

There are three lists provided by this command:

- 1) Active jobs: Jobs that are running now. The maximum is 1000 jobs.
- 2) Eligible jobs: Jobs are waiting for resources to run. The maximum is 15 jobs.

- 3) **Blocked jobs.** Jobs are blocked for a variety of reasons, the most common are as follows:
- Deferred: The job has been deferred to run at a later time (normally 1 hour). You can check the defer time using the checkjob command (see below). The system will defer a job when it does not run correctly. This can happen if the node the job is being submitted to is down. This can also happen if the job crashes on startup due to some resource or programming error. After the defer time the job will be submitted back into the queue. The defer time is used to prevent bad jobs from flooding the system while still letting good jobs run.
 - Idle: Jobs that are idle have reached the maximum 15 eligible jobs limit. Each of these jobs will move out of idle status when a running job has finished.
 - Hold: User has enforced some dependencies.

Detailed information about a job

To verify that your job has the right resources requested, use the command **qstat -f <jobid>**

```
dev-intel16-k80:~$ qstat -f 51293730
Job Id: 51293730.mgr-04.i
  Job_Name = hello.qsub
  Job_Owner = choiyj@dev-intel16-k80.i
  job_state = Q
  queue = main
  server = mgr-04.i
  Checkpoint = u
  ctime = Mon Jan 15 12:10:14 2018
  Error_Path = dev-intel16-k80.i:/mnt/home/choiyj/helloworld/hello.qsub.e512
              93730
  Hold_Types = n
  Join_Path = oe
  Keep_Files = n
  Mail_Points = a
  mtime = Mon Jan 15 12:10:14 2018
  Output_Path = dev-intel16-k80.i:/mnt/home/choiyj/helloworld/hello.qsub.o51
              293730
  Priority = 0
  qtime = Mon Jan 15 12:10:14 2018
  Rerunable = False
  Resource_List.nodes = 1:ppn=1
  Resource_List.walltime = 01:00:00
  Resource_List.mem = 200mb
  Resource_List.feature = gbe
  Resource_List.nodect = 1
  Variable_List = PBS_O_QUEUE=main,PBS_O_HOME=/mnt/home/choiyj,
```

```
PBS_0_LOGNAME=choiyj,
PBS_0_PATH=/mnt/home/choiyj/anaconda2/bin:/opt/software/fftw/3.3--GCC
-4.4.5/bin:/opt/software/cudatoolkit/6.0/include/CL:/opt/software/cuda
toolkit/6.0/include:/opt/software/cudatoolkit/6.0/computeprof/bin:/opt
/software/cudatoolkit/6.0/bin:/opt/software/allinea/5.0.1-42607/bin:/o
pt/software/powertools/bin:/mnt/home/choiyj/anaconda2/bin:/opt/softwar
e/R/2.15.1--GCC-4.4.5/bin:/opt/software/MATLAB/R2014a/bin:/opt/softwar
e/cmake/2.8.5--GCC-4.4.5/bin:/opt/software/OpenMPI/1.4.3--GCC-4.4.5/bi
n:/usr/lib64/qt-3.3/bin:/opt/software/lmod/bin:/opt/maob/bin:/usr/loca
l/bin:/bin:/usr/bin:/usr/local/sbin:/usr/sbin:/sbin:/usr/local/hpcc/bi
n:/opt/ibutils/bin:/mnt/home/choiyj/bin,
PBS_0_MAIL=/var/spool/mail/choiyj,PBS_0_SHELL=/bin/bash,
PBS_0_LANG=en_US.UTF-8,PBS_0_WORKDIR=/mnt/home/choiyj/helloworld,
PBS_0_HOST=dev-intel16-k80.i,PBS_0_SERVER=mgr-04

euser = choiyj
egroup = helpdesk
queue_type = E
etime = Mon Jan 15 12:10:14 2018
submit_args = hello.qsub
fault_tolerant = False
job_radix = 0
submit_host = dev-intel16-k80.i
init_work_dir = /mnt/home/choiyj/helloworld
request_version = 1

dev-intel16-k80:~$
```

Deleting a job

You can delete jobs by issuing the command **qdel <jobid>**

Deleting all submitted jobs

Sometimes, you might need to delete all the jobs you have submitted. In that case, use **qdel \$(qselect -u <username>)**

```
dev-intel16-k80:~$ qstat -u choiyj

mgr-04.i:
Job ID          Username      Queue      Jobname      SessID  NDS  TSK  Req'd  Req'd   Elap
                Memory      Time      S          Time
-----
51293689[ ] mgr-04.i  choiyj     main         ucp_gputest  --    1    1     20gb  04:00:00 Q  --
51293730 mgr-04.i  choiyj     main         hello.qsub   --    1    1     200mb  01:00:00 Q  --
dev-intel16-k80:~$ qdel $(qselect -u choiyj)
dev-intel16-k80:~$ qstat -u choiyj

mgr-04.i:
Job ID          Username      Queue      Jobname      SessID  NDS  TSK  Req'd  Req'd   Elap
                Memory      Time      S          Time
-----
51293689[ ] mgr-04.i  choiyj     main         ucp_gputest  --    1    1     20gb  04:00:00 C  --
51293730 mgr-04.i  choiyj     main         hello.qsub   --    1    1     200mb  01:00:00 C  --
dev-intel16-k80:~$
```

Jan. 2018



Deleting all jobs in a job array

Sometimes, you might need to delete all the jobs in a job array. In that case, add an open and close bracket (ex: []) to the end of the job id number:

```
dev-intel16-k80:~$ qdel 51293689[]
```

Reordering jobs in the queue

See <https://wiki.hpcc.msu.edu/display/hpccdocs/move2top>

module load powertools

move2top <jobid>x

There is also a companion command called move2bot which moves a job to the bottom of the queue, see <https://wiki.hpcc.msu.edu/display/hpccdocs/move2bot>.

If you have any further question, please feel free to contact us.