

Peer scheduling between Cheyenne and Casper

Cheyenne and Casper use the PBS scheduler, and each system has its own dedicated PBS server to manage job submission and execution. These "peer" servers can share data between each other, and *jobs can be submitted from one system to another*. It is also possible to create dependencies between jobs on each server, enabling simulation-analysis workflows that target the appropriate system for each task. (*Effective October 20, 2021*)

Page contents

- [Submitting a job](#)
- [Querying jobs](#)
- [Creating dependencies between peer-scheduled jobs](#)

Submitting a job

To submit a job to a queue on a peer server, you need to append the name of the server to the queue directive in your job script. The names of the PBS servers are:

System	PBS server name
Cheyenne	chadmin1.ib0.cheyenne.ucar.edu
Casper	casper-pbs

Example 1: You want to submit to the Cheyenne "regular" queue from a Casper login node or compute node. Append the Cheyenne server name as follows in your job script when specifying the queue:

```
#PBS -q regular@chadmin1.ib0.cheyenne.ucar.edu
```

Example 2: You want to submit a job to Casper from Cheyenne. Append the Casper server name as follows in your job script when specifying the destination:

```
#PBS -q casper@casper-pbs
```

The server-specific queue names will be understood by both PBS servers, so if you will want to submit the same script at times from either Cheyenne or Casper, *always append the server name to your queue*.

The **qinteractive** and **execcasper** scripts, which start interactive jobs on Cheyenne and Casper respectively, will adjust the queue name for you to include the server, so *you do not need to append the server name manually for interactive jobs*.

Querying jobs

You can use **qstat** to query jobs from peer servers by including the server name in your field of interest. You can also use the system names noted above when running qstat.

Note that the separator character differs for jobs (.) and queues (@).

```
qstat 123456.chadmin1.ib0.cheyenne.ucar.edu
qstat regular@chevyenne
qstat 654321.casper
```

Creating dependencies between peer-scheduled jobs

Creating job dependencies between submissions on peer servers is straightforward; there is nothing unique about this workflow in PBS. As with all jobs, pay close attention to specifying the destination server in your queue designations. The job IDs returned by PBS include the server name, so you do not need to append a server to the job ID you specify in your dependency argument.

Here is an example of a workflow that runs a simulation on Cheyenne and, if successful, then runs a post-processing job on Casper. Thanks to peer scheduling, these jobs can be submitted from either Cheyenne or Casper login nodes.

tcsh example

```
set JID=`qsub -q economy@chadmin1.ib0.cheyenne.ucar.edu run_model.pbs`  
qsub -q casper@casper-pbs -W depend=afterok:$JID run_postprocess.pbs
```

bash example

```
JID=$(qsub -q economy@chadmin1.ib0.cheyenne.ucar.edu run_model.pbs)  
qsub -q casper@casper-pbs -W depend=afterok:$JID run_postprocess.pbs
```