

# PBSPro - Job Management System

The JCU HPC job management system (PBSPro) works with 2 login nodes, 16 compute nodes, and 1 test node (mostly TS staff only). All nodes have 40 CPU cores, 384GiB of RAM/memory, 480GB of RAID1 SSD storage, ~7TB of RAID0 SSD storage, and 51Gb/s of network connectivity. All nodes are running the RedHat Enterprise Linux 7 operating system.

Three "walltime request" job queues (FIFO) have been configured to accept researcher workflows. On 1-Mar-2021, the following configuration was operational:

Queue	Walltime		Max. CPU cores	
	min.	max.	per user	all jobs
tiny	0:00:01	2:00:00	360	<i>all</i>
short	2:00:01	24:00:00	240	<i>all</i>
normal	24:00:01	168:00:00	160	400
long	168:00:01	720:00:00	80	240

Job array limits match the per-user limits in the table above. The limits mentioned above are reviewed regularly and have been changed on multiple occasions (to match researchers' usage patterns). **Note:** 720:00:00 is equivalent to 30 days.

**Play nice:** The JCU HPC has not be architected for multi-node MPI jobs (no researcher should be submitting job requests that involve more than 1 node (`select=1`)). Researchers who need to run jobs that require more than 40 cores should seek time on QCIF, NCI, or Pawsey HPC facilities (or public cloud if you have sufficient budget). The storage platforms that house JCU HPC filesystems were purchased for capacity, not performance.

Historically, there are idle CPU cores on the HPC cluster about 70% of the time. As a result, a FIFO queues configuration was deemed to be *minimum viable product*. Fairshare queues will be configured if there is evidence of near 100% utilization with jobs waiting in queues for a period of at least 3 months. At maximum compute cluster capacity, we only have 680 CPU cores available to accept jobs/workflows.

## Information relating to resource requests

There are several factors to consider when requesting resources for your job(s).

- The more accurately your resource request is, the higher return on investment JCU gets from operating a HPC cluster.
- Many HPC facilities charge for resources requested. The resources requested are dedicated to a job, regardless of whether they are used or not.
- Requesting multiple CPU cores does not mean that your job will be able to use these cores - most software is only written to use 1 CPU core. Make sure the software you will be using is capable of running on multiple cores before making such requests.
- The resources you request for a job are dedicated to your job - unused components are not available for other jobs (yours and other researchers').
- PBSPro has been configured to **kill** jobs that consume more resources than they request, although there are some margins configured to cater for very short deviations. In some cases, HPC staff can increase the limits on a running job. Under-specification of resource requirements will lead to inefficiency. In the worst case scenarios, your request could result in a server crashing - impacting all jobs running on the server. You could be looking at jobs taking 1000x longer than if insufficient resources were requested.

Researchers who are found to repeatedly under-request or significantly over request job/workflow resources will be contacted in an attempt to change their behaviour. HPC staff realise that many people do not know the memory requirements of their jobs - e.g., memory requirement can vary based on input data or type of analysis performed.