

PBSPro - Job Management System

CPU Compute Cluster Resources Available (Jan 2020)

There are 17 CPU compute nodes installed for HPC cluster jobs. All nodes are configured with:

CPU cores	Memory	Network ports	Local SSDs	Operating System
40	384 GiB	2x25Gb/s + 1Gb/s	480GB	RHEL 7.x

GPU Compute Cluster Resources Available (Nov 2018)

There are 2 GPU compute nodes. Each node is configured with:

CPU cores	Memory	GPU cards	GPU memory	SSH Network	NFS Network	Local SSDs	Operating System
24	192 GiB	2 x V100	16GB per card	1Gb/s	10Gb/s	480GB+960GB	Ubuntu 16.04

JCU has purchased GPU capacity from QCIF - access to 36 V100 cards (32GiB of memory per card). The existing GPU servers will be repurposed sometime after we gain access to the UQ managed resource.

Configuration details for job management system (Dec 2018)

Walltime Requested	Queue
0:00:00 - 24:00:00	short
24:00:01 - 168:00:00	normal
168:00:01 - 2160:00:00	long

The maximum walltime for each queue may be changed (to match usage patterns). Note; 2160:00:00 = 90 days.

The values in the table below may be changed (to match usage patterns). Note that the HPC cluster has a maximum of 600 CPU cores available (as of 10-Dec-2019).

Queue	Max. jobs in queue	Max. CPUs in use	Max. job array size
short	1000	540	200
normal	1000	400	120
long	160	80	40

Resource Requirements

More accurate resources request equate to higher return on investment. Organisations such as NCI and AWS charge for resources allocated/requested.

Resource under-specification

PBSPro has been configured to **kill** jobs that consume more resource than they request. In some cases, HPC staff can increase the limits - dependent on resource and situation.

Resource under-specification often leads to inefficiency (impacts worse for memory than CPU). Resource under-specification can also lead to compute node(s) crashing - potentially affecting other users' jobs.

Resource over-specification

The resources you request for a job are dedicated to your job - unused components are not available for other jobs.

Users who repeatedly over-specify CPU and/or memory resource requirements will be contacted by ICT/HPC staff to change their behaviour.

Most jobs will only use 1 CPU core - requesting more will not see your job complete more quickly unless the software you are using is written to support execution on multiple CPU cores.

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.

The more resource you request, the more likely your usage will be scrutinised.