Paracluster usage rules

- Never use more CPUs than requested in your job submission.
  - WARNING: It is possible to use more cpu cores than allocated through the queuing system. Please do not do that! It will lock up a node for other users. When this happens, we will kill the offending jobs.
  - Some programs internally spawn multiple threads. Check your program for this, and adjust your job scripts accordingly. Python's numpy and multitasking modules can do this. Amuse uses a main thread plus one per worker code used in your program.
- Set reasonable limits on your jobs. Don't reserve more CPUs than you are going to use. A good estimate of the expected execution time helps scheduling a lot, and will make your jobs process more efficiently (and with some luck, everyone else's jobs too).
- Memory use is not used in scheduling, but the total jobs on a node should not exceeed the total available memory. If a node crashes, all the work done there will be interrupted and job output may be lost, so it is in everyone's best interest to check how much memory your programs will require.
- Don't use big data sets from a disk outside the para cluster (eg from your desktop). Access to such files will be quite slow and will slow down all execution on the node. Copying data locally to one of the /net/para*/data* disks before starting the job, will usually be much faster.
- Don't forget to free up disk space when you are done. The local disks on the cluster are for data processing, not for permanent storage

Request for access

Users who have read these usage rules and agree to follow them, can request access by contacting the computer group. Access can be revoked if a user repeatedly and willingly violates the usage guidelines. Being part of any research group does not automatically grant you access.

From:
https://helpdesk.physics.leidenuniv.nl/wiki/ - **Computer Documentation Wiki**

Permanent link:
https://helpdesk.physics.leidenuniv.nl/wiki/doku.php?id=paracluster:access&rev=1454408338

Last update: **2016/02/02 10:18**