

Things to know

when using the paracluster:

1. Always submit your jobs to the queue. Do not ssh into the machines and start jobs.
2. Be aware of how much storage space you are using, and try to store only what you need for your current paracluster usage.
3. Jobs can take more memory than the number of processors you have allotted them, so please monitor your memory usage and adjust your submit scripts accordingly.
4. Jobs can use more cores than allocated. Take care that they should not. if you know a program starts eg 2 sub-programs, edit your job script accordingly and allocate those additional cores.
 1. in some cases you may have to tell your software the maximum number of cores to use, eg numpy, compiled with OpenBLAS, takes as many cores as available and useful for a job; set environment variables OPENBLAS_NUM_THREADS and/or OMP_NUM_THREADS to the maximum number to use (the OPENBLAS variable is for things that use OpenBLAS, obviously; the OMP variable is for OpenMP). There may be other libraries that have similar settings, please check what you are using, and how to properly use multiprocessing features in your software!!

From:

<https://helpdesk.strw.leidenuniv.nl/wiki/> - **Computer Documentation Wiki**

Permanent link:

https://helpdesk.strw.leidenuniv.nl/wiki/doku.php?id=paracluster:things_to_know

Last update: **2017/11/27 09:21**

