

Best Practices

General Best Practices

1. Before starting, you should always check:
 - Are there any errors in the script?
 - Are the required modules loaded?
 - Is the correct executable used?
2. Check your computer requirements upfront, and request the correct resources in your batch job script.
 - Number of requested cores
 - Amount of requested memory
 - Requested network type
3. Check your jobs at runtime. You could login to the node and check the proper execution of your jobs with, e.g., `top` or `vmstat`. Alternatively you could run an interactive job (`qsub -I`).
4. Try to benchmark the software for scaling issues when using MPI or for I/O issues.
5. Use the scratch file system (`$VSC_SCRATCH_NODE`, which is mapped to the local `/tmp`) whenever possible. Local disk I/O is always much faster as it does not have to use the network.
6. When your job starts, it will log on to the compute node(s) and start executing the commands in the job script. It will start in your home directory `$VSC_HOME`, so going to the current directory with `cd $PBS_O_WORKDIR` is the first thing which needs to be done. You will have your default environment, so don't forget to load the software with `module load`.
7. Submit your job and wait (be patient) ...
8. Submit small jobs by grouping them together. See chapter [Multi-job submission](#) for how this is done.
9. The runtime is limited by the maximum walltime of the queues.
10. Requesting many processors could imply long queue times. It's advised to only request the resources you'll be able to use.
11. For all multi-node jobs, please use a cluster that has an "InfiniBand" interconnect network.
12. And above all, do not hesitate to contact the HPC staff at hpc@ugent.be. We're here to help

you.