

## File Systems

Path	Shell variable	Purpose	Quota	Back-ups
/home/<lab>/<username>	\$HOME	non-reproducible data, data with high reproduction cost	200 GB	nightly
/scratch/<lab>/<username>	\$HPC_LOCAL	temporary data, large files, intermediate results, easily re-computable results	2 TB	none

## Default

Users are required to use partitions that best fit their use case. Refer to the cluster wiki (<https://justhpc.hpc.uni-giessen.de/trac/wiki/SlurmExamples>) for example job scripts.

## Access

**Only members of the JLU Giessen are eligible to obtain a cluster account.**

Login is only possible from within the university's network. From outside connect to the internal network via VPN.

Connect to the cluster's front-end node via SSH:

```
$ ssh -XC <username>@justhpc.hpc.uni-giessen.de
```



## Partitions/ Runtime constraints

Partition	Time limit	User case
single	5 days	single <i>node</i> jobs
debug	30 minutes	testing and debugging
serial	5 days	single <i>core</i> jobs
regular	5 days	parallel production jobs
bigmem	5 days	jobs with large RAM requirements

## Module System

**Hint:** Load *all* required modules inside your job scripts! `module purge` should be your first command!

Common commands:

- `module [un]load <name of module>`: Load / unload a particular module
- `module purge`: (use in job script as first module command)
- `module list`: List all currently loaded modules
- `module avail`: List all currently loadable modules

Please refer to the cluster wiki (<https://justhpc.hpc.uni-giessen.de/trac/wiki/ModuleFiles>) to learn more about the details of our module system.

## Help

In case you have questions or run into some issues (e.g. required software not installed) file a ticket in our ticket system: <https://justhpc.hpc.uni-giessen.de/trac>.

Please note that the ticket system is only accessible from the JLU internal network. From the outside you need to connect to this network via VPN.

## SLURM commands

For details refer to the man pages or the Slurm website: <https://slurm.schedmd.com>

Most important commands:

- `sbatch`, `scancel`, `squeue`, `srun`

`sbatch` examples

- Submit a jobscript:  
`$ sbatch <job script>`
- Submit a jobscript as a specific account (other than default):  
`$ sbatch --account <account> <job script>`
- Submit a jobscript that only starts after the other given jobs have finished successfully:  
`$ sbatch --dependency=afterok:<jobid>[:<jobid>[:...]] <job script>`

`scancel` examples

- `$ scancel <jobid>`
- **Note:** `skill` is *not* a slurm command!

`srun` examples

- **Note:** `srun` is used inside a job script to start subtasks.  
`srun [-N <number of tasks>] /path/to/application`

`squeue` examples

- Show all running and pending jobs  
`$ squeue`
- Show running and pending jobs of a user (for yourself, you can also use `squeue -u $USER`)  
`$ squeue -u <username>`
- Show a time *estimate* when this job will be starting (can change over time!)  
`$ squeue --start -j <jobid>`