

Job Dependencies

This page describes the SLURM dependency feature.

This feature is used when you need to chain jobs, due to dependencies. For example:

- A preprocessing job with 1 core should be followed by a simulation with 16 cores. The results should be post-processed with a single main job.
- A post-processing job must be sent after all tasks in a job matrix are completed.

Usage

The dependent job must be launched using the `sbatch` command by specifying the `--dependency=<type>:<listOfJobIDs>` option. The type can be:

- `after` : the job is launched after the first job enters into execution, i.e. it is RUNNING.
- `afterok` : the job is launched after the first job has been successfully completed.
- `afterany` : the job is launched after the first job has finished, regardless of whether it has failed or not.
- `afternotok` : the job is launched after the first job has finished on failure.
- `aftercorr` : a job in an array of jobs can start executing after the ID of another job in the array has been successfully completed.
- `singleton` : the job can start running when any other job with the same name and from the same user has finished.

Example:

```
sbatch --dependency=afterok:111111,111112 my_job_script.sh
```

The underlying job (on which this job depends) must be submitted first. The related job ID can be captured by collecting the output of the `sbatch` command with the `--parsable` option:

```
first_jobid=$(sbatch --parsable my_first_job.sh)
```

Info

[more information, check the official documentation](#)