

[Pages](#) / ... / [Advanced HPC Skills](#)

Job dependencies with SLURM

Created by Michael Burrell (NBI), last modified by Sam Gallop (NBI) on Feb 04, 2016

Many scientific computing tasks consist of serial processing steps. A genome assembly pipeline, for example, may require sequence quality trimming, assembly, and annotation steps that must occur in series. Launching each of these jobs without manual intervention can be done by repeatedly polling the controller with `squeue / sacct` until the State is `COMPLETED`. However, it's much more efficient to let the SLURM controller handle this using the **`sbatch --dependency`** option.

`-d, --dependency=<dependency_list>`

Defer the start of this job until the specified dependencies have been satisfied. The `<dependency_list>` is of the form `<type:job_id[:job_id][,type:job_id[:job_id]]>`. Many jobs can share the same dependency and these jobs may even belong to different users. The value may be changed after job submission using the `scontrol` command.

after:job_id[:jobid...]

This job can begin execution after the specified jobs have begun execution.

afterany:job_id[:jobid...]

This job can begin execution after the specified jobs have terminated.

afternotok:job_id[:jobid...]

This job can begin execution after the specified jobs have terminated in some failed state (non-zero exit code, node failure, timed out, etc).

afterok:job_id[:jobid...]

This job can begin execution after the specified jobs have successfully executed (ran to completion with an exit code of zero).

expand:job_id

Resources allocated to this job should be used to expand the specified job. The job to expand must share the same QOS (Quality of Service) and partition. Gang scheduling of resources in the partition is also not supported.

singleton

This job can begin execution after any previously launched jobs sharing the same job name and user have terminated.

 [Like](#) Be the first to like this

[slurm](#) [advanced](#)

