

Slurm Commands

Start Jobs

srun This is the simplest way to run a job on a cluster. Initiate parallel job steps within a job or start an interactive job (with `-pty`).

salloc Request interactive jobs/allocations. When the job is started a shell (or other program specified on the command line) it is started on the submission host (Frontend). From this shell you should use `srun` to interactively start a parallel applications. The allocation is released when the user exits the shell.

sbatch submit a batch script. The script will be executed on the first node of the allocation. The working directory coincides with the working directory of the sbatch directory. Within the script one or multiple `srun` commands can be used to create job steps and execute parallel applications.

Examples

```
# General:
sbatch --job-name=$jobname -N <num_nodes> --ntasks-per-node=<ppn>
/path/to/sbatch.script.sh

# A start date/time can be set via the --begin parameter:
--begin=16:00
--begin=now+1hour
--begin=now+60 (seconds by default)
--begin=2010-01-20T12:34:00
```

For more information see `man sbatch`. All parameters used there can also be specified in the job script itself using `#SBATCH`.

Check the status of your own jobs

```
squeue
```

Check the status of the nodes

```
sinfo
```

Canceling jobs

On allocation, you will be notified of the job ID. Also, within your scripts and shells (if allocated via `salloc`) you can get the ID via the `$SLURM_JOBID` environment variable. You can use this ID to cancel your submission:

scancel <jobid>

From:
<https://hpc.tu-berlin.de/> - **HPC-Cluster-Dokumentation**

Permanent link:
https://hpc.tu-berlin.de/doku.php?id=hpc:scheduling:slurm_commands&rev=1727355453

Last update: **2024/09/26 14:57**

