

# sbatch

## Description:

---

**sbatch** submits a batch script to Slurm.

The batch script may be given to sbatch through a file name on the command line, or if no file name is specified, sbatch will read in a script from standard input.

The batch script may contain options preceded with "#SBATCH" before any executable commands in the script.

## Options:

---

### **-h, --help**

Display help information and exit.

### **-u, --usage**

Display brief help message and exit.

### **-a, --array=<indexes>**

Submit a job array, multiple jobs to be executed with identical parameters.

The indexes specification identifies what array index values should be used. Multiple values may be specified using a comma separated list and/or a range of values with a "-" separator. For example, "--array=0-15" or "--array=0,6,16-32".

A step function can also be specified with a suffix containing a colon and number. For example, "--array=0-15:4" is equivalent to "--array=0,4,8,12".

A maximum number of simultaneously running tasks from the job array may be specified using a "%" separator. For example "--array=0-15%4" will limit the number of simultaneously running tasks from this job array to 4.

### **--begin=<time>**

Submit the batch script to the Slurm controller immediately, like normal, but tell the controller to defer the allocation of the job until the specified time.

**--comment**=<*string*>

An arbitrary comment enclosed in double quotes if using spaces or some special characters.

**-C, --constraint**=<*list*>

Nodes can have **features** assigned to them by the Slurm administrator.

**-c, --cpus-per-task**=<*ncpus*>

Advise the Slurm controller that ensuing job steps will require *ncpus* number of processors per task (use with shared memory parallelism).

Without this option, the controller will just try to allocate one processor per task.

**-d, --dependency**=<*dependency\_list*>

Defer the start of this job until the specified dependencies have been satisfied completed.

Dependencies list:

**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.

**aftercorr:**<*job\_id[:jobid...]*>

A task of this job array can begin execution after the corresponding task ID in the specified job has completed successfully (ran to completion with an exit code of zero).

**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[:jobid...]*>

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.

### **-D, --workdir=<directory>**

Set the working directory of the batch script to *directory* before it is executed. The path can be specified as full path or relative path to the directory where the command is executed.

### **-e, --error=<filename pattern>**

Instruct Slurm to connect the batch script's standard error directly to the file name specified in the "*filename pattern*". By default both standard output and standard error are directed to the same file.

For job arrays, the default file name is "slurm-%A\_%a.out", "%A" is replaced by the job ID and "%a" with the array index. For other jobs, the default file name is "slurm-%j.out", where the "%j" is replaced by the job ID.

### **--export=<environment variables / ALL / NONE>**

Identify which environment variables are propagated to the batch job. Multiple environment variable names should be comma separated.

### **--gres=<list>**

Specifies a comma delimited list of generic consumable resources. The format of each entry on the list is "name[:type]:count".

### **-H, --hold**

Specify the job is to be submitted in a held state (priority of zero). A held job can now be released using *scontrol release <job\_id>*).

### **-J, --job-name=<jobname>**

Specify a name for the job allocation.

### **-L, --licenses=<license>**

Specification of licenses (or other resources available on all nodes of the cluster) which must be allocated to this job.

**--mail-type=<type>**

Notify user by email when certain event types occur.

Valid *type* values are NONE, BEGIN, END, FAIL, REQUEUE, ALL (equivalent to BEGIN, END, FAIL, REQUEUE, and STAGE\_OUT), TIME\_LIMIT, TIME\_LIMIT\_90 (reached 90 percent of time limit), TIME\_LIMIT\_80 (reached 80 percent of time limit), TIME\_LIMIT\_50 (reached 50 percent of time limit) and ARRAY\_TASKS (send emails for each array task). Multiple *type* values may be specified in a comma separated list.

**--mail-user=<user>**

User to receive email notification of state changes as defined by **--mail-type**.

**--mem=<MB>**

Specify the real memory required per node in megabytes.

**--mem-per-cpu=<MB>**

Minimum memory required per allocated CPU in megabytes. Different units can be specified using the suffix [K|M|G|T].

**-N, --nodes=<minnodes[-maxnodes]>**

Request that a minimum of *minnodes* nodes be allocated to this job.

A maximum node count may also be specified with *maxnodes*. If only one number is specified, this is used as both the minimum and maximum node count.

**-n, --ntasks=<number>**

This option advises the Slurm controller that job steps run within the allocation will launch a maximum of number tasks and to provide for sufficient resources.

Note that sbatch does not launch tasks, it requests an allocation of resources and submits a batch script.

**--ntasks-per-node=<ntasks>**

Request that *ntasks* be invoked on each node.

**-o, --output=<filename pattern>**

Instruct Slurm to connect the batch script's standard output directly to the file name specified in the "*filename pattern*". By default both standard output and standard error are directed to the same file.

For job arrays, the default file name is "slurm-%A\_%a.out", "%A" is replaced by the job ID and "%a" with the array index. For other jobs, the default file name is "slurm-%j.out", where the "%j" is replaced by the job ID.

**-p, --partition=<partition\_names>**

Request a specific partition for the resource allocation.

**--reservation=<name>**

Allocate resources for the job from the named reservation.

**-t, --time=<time>**

Set a limit on the total run time of the job allocation.

**--test-only**

Validate the batch script and return an estimate of when a job would be scheduled to run given the current job queue and all the other arguments specifying the job requirements. No job is actually submitted.

**-w, --odelist=<node name list>**

Request a specific list of hosts.

**--wrap=<command string>**

Sbatch will wrap the specified command string in a simple "sh" shell script, and submit that script to the slurm controller.

When --wrap is used, a script name and arguments may not be specified on the command line; instead the sbatch-generated wrapper script is used.

**-x, --exclude=<node name list>**

Explicitly exclude certain nodes from the resources granted to the job.