

What is SLURM

SLURM is an open source and highly scalable cluster management and job scheduling system for large clusters.

Slurm has three key functions.

- First, it allocates exclusive and/or non-exclusive access to resources (compute nodes) to users for some duration of time so they can perform work.
- Second, it provides a framework for starting, executing, and monitoring work (normally a parallel job) on the set of allocated nodes.
- Finally, it arbitrates contention for resources by managing a queue of pending work.

Currently the cluster has **version 19** of slurm:

The most important concepts within slurm are:

- Nodes.
- Partitions.
- Jobs.
- Task (a task represent a process) in a job.

Login nodes

It is from the login nodes where the user interacts with Slurm and from where they can launch and monitor their jobs. From here, the user accesses their data and the results of the executions.

Info

Remember that the **login nodes** are shared by all users, so that **software execution on these nodes is forbidden**. The Slurm scheduler must be used for this.

Partitions

Partitions can be considered as **job queues**, each of which has an assortment of constraints such as job size limit, job time limit, user permitted to use it, etc. Priority-ordered jobs are allocated nodes within a partition until the resources (nodes, processors, memory, etc) within that partition are exhausted.

The partitions defined in our cluster are the following and their only restrictions are the time and the users who can use it:

Partición	Tiempo máximo	
express	3 horas	<i>all users</i>
batch	24 horas	<i>all users</i>
long	72 horas	<i>authorized users</i>
fatnodes	--	upon request

If no partition is specified, the default partition is the `batch`.

If your work lasts longer than the maximum time established by the partition where you launched it, contact us. support@hpc.iter.es to request an extension of the time limit.

Most common commands in SLURM

The most common slurm commands are listed below. To get more information about the command as options you can always run

```
man <command>
<command> --help
```

- `sbatch <script file>` : launch script
- `squeue` : check the status of job queues
- `scancel <job_id list>` : cancel a job
- `scontrol show job <job_id>` : get information about job
- `sinfo` : view status of system queues
- `salloc <opciones>` : start interactive session (get a node for to use it)

- `srun <aplicacion>` : submit a job to run or start the job steps in real time
- `sacct` : check the accounting of the account itself
- `sstat` : get information about the resources used by a running job

You can see a more detailed summary guide in our section [Useful commands in slurm](#) or in [slurm web](#).