

Slurm

Slurm's architecture

Slurm is made of a slurmd daemon running on each compute node and a central slurmctld daemon running on a management node.

Node

In slurm a node is a compute resource, usually defined by particular consumable resources, i.e. cores, memory, etc...

Partitions

A partition (or queue) is a set of nodes with usually common characteristics and/or limits. Partitions group nodes into logical sets. Nodes are shareable between partitions.

Jobs

Jobs are allocations of consumable resources from the nodes and assigned to a user under the specified conditions.

Job Steps

A job step is a single task within a job. Each job can have multiple tasks (steps) even parallel ones.

Common user commands:

- [**sacct**](#): report job accounting information about running or completed jobs.
- [**salloc**](#): allocate resources for a job in real time. Typically used to allocate resources and spawn a shell. Then the shell is used to execute commands to launch parallel tasks.
- [**sbatch**](#): submit a job script for later execution. The script typically contains the tasks plus and the environment definitions needed to execute the job.
- [**scancel**](#): cancel a pending or running job or job step.
- [**sinfo**](#): overview of the resources (node and partitions).
- [**squeue**](#): used to report the state of running and pending jobs.

- [**srun**](#):submit a job for execution or initiate job steps in real time. The srun allows users to requests consumable resources.
-

Revision #8

Created 2019-11-28 10:54:19 UTC by Joao Machado

Updated 2021-05-25 23:43:09 UTC by Joao Machado