

Slurm

This is the front page for information to our compute resource sharing system. We use software called Slurm to fairly share compute resources.

For job submission we will be using a piece of software called [Slurm](#). Simply put, Slurm is a queue management system; it was developed at the Lawrence Livermore National Lab. It currently supports some of the largest compute clusters in the world. The best description of Slurm can be found on its homepage:

"Slurm is an open-source workload manager designed for Linux clusters of all sizes. It provides three key functions. First it allocates exclusive and/or non-exclusive access to resources (computer nodes) to users for some duration of time so they can perform work. Second, it provides a framework for starting, executing, and monitoring work (typically a parallel job) on a set of allocated nodes. Finally, it arbitrates contention for resources by managing a queue of pending work."¹⁾

Slurm is similar to most other queue systems in that you write a batch script, then submit it to the queue manager. The queue manager schedules your job to run on the queue (or partition in Slurm parlance) that you designate. Below is an outline of how to submit jobs to Slurm, how Slurm decides when to schedule your job, and how to monitor progress.

Clusters

Peanut Cluster

Think of these machines as a dumping ground for discrete computing tasks that might be rude or disruptive to execute on the main (shared) shell servers (i.e., linux1, linux2, linux3).

Additionally, this cluster is used for courses that require it.

AI Cluster

This cluster is mainly made up of GPU machines and is used primary for research.

To use this cluster there are specific nodes you need to log into. Please visit the dedicated AI cluster page for more information.

Where to begin

Slurm is a set of command line utilities that can be accessed via the command line from **most** any computer science system you can login to. Using our main shell servers (linux.cs.uchicago.edu) is expected to be our most common use case, so you should start there.

```
ssh user@linux.cs.uchicago.edu
```

If you want to use the AI Cluster you will need to login into:

```
ssh user@fe.ai.cs.uchicago.edu
```

Please read up on the specifics on the cluster you are interested in.

Mailing List

If you are going to be a user of this cluster please sign up for the mailing list. Downtime and other relevant information will be announced here.

[Mailing List](#)

Documentation

The [Slurm website](#) should be your primary source for documentation.

A great way to get details on Slurm commands are the manuals that are already on the cluster. For example, if you type the following command:

```
man sbatch
```

you will get the manual page for the sbatch command.

Resources

- [Common Slurm commands](#)
- [Official Slurm website](#)
- [Official Slurm documentation](#)
- [Slurm tutorial videos](#)
- [LLNL quick start user guide](#)
- [Yale's User Guide](#)

¹⁾

<http://slurm.schedmd.com/>

From:
<https://howto.cs.uchicago.edu/> - **How do I?**

Permanent link:
<https://howto.cs.uchicago.edu/slurm?rev=1609970448>

Last update: **2021/01/06 16:00**

