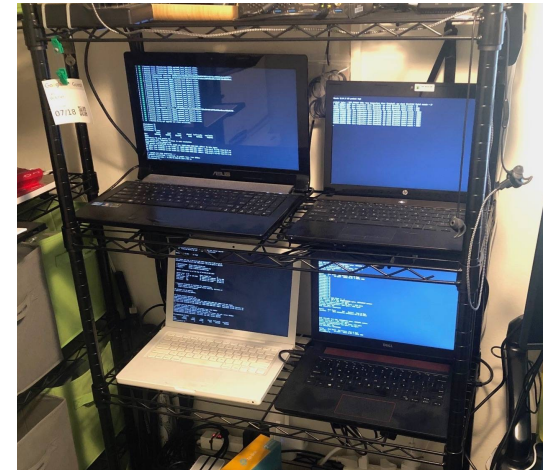




Resource and Job Management

Thought Experiment

- Executing a program in your PC (say laptop):
 - You just simply run a *program* (./program)
- You have 4 old laptops and made your personal cluster (somehow it is magically connected and configured ;)). Your own **“Beowulf Cluster”**! Google **“Stone *Soup* er Computer”**
- Let’s do the **thinking** for below scenarios and **create** (hypothetically) a software ;)
 - S1: How will you run your *program* in all/some laptops (let’s call it a resource)?
 - You must **create** some **software** for ease of use
 - S2: You need to run the *program* only in laptop with GPU (White laptop)
 - S3: You need to give access to your friends, but you don’t want to give access to your favorite laptop ;)
 - S4: You (or your cluster) became popular and now more people want to access it.
 - S5: Two of your best friends are accessing it and you want to give them faster access ;)
 - S6: Now you might be worrying about the electricity costs, health of your laptops!!!

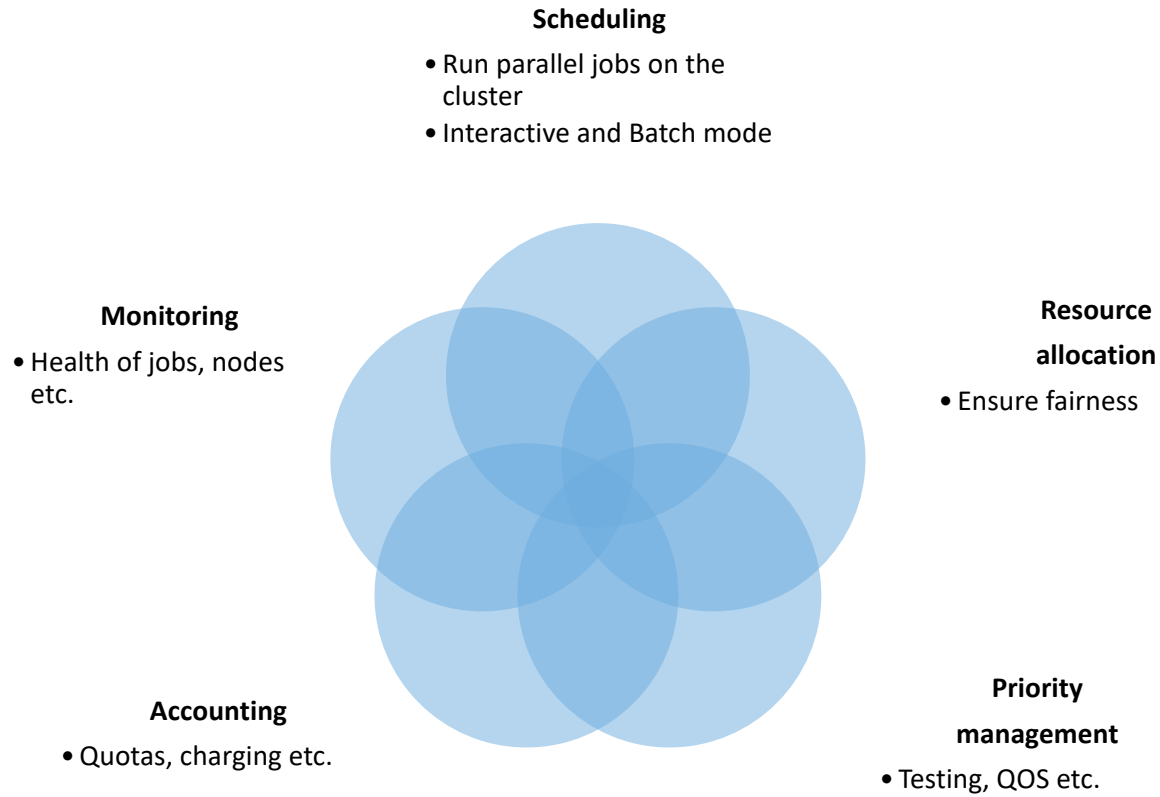



Thought Experiment



- S1: How will you run your program in all/some laptops?
 - **Resource Management Software and Runtime System**
- S2: You need to run the program only in laptop with GPU (White laptop)
 - **Generic Resources - Resource Management – Resource allocation**
- S3: You need to give access to your friends, but you don't want to give access to your favorite laptop ;)
 - **Partitions and Access permissions**
- S4: You (or your cluster) became popular and now more people want to access it.
 - **Job Scheduling and Job Queues**
- S5: Two of your friends are accessing it and you want it to give faster access to your best friend ;)
 - **Job Priority**
- S6: Now you might be worrying about the electricity costs, health of your laptops!!!
 - **Monitoring and Accounting**








“Slurm is an open source workload management system and job scheduler. It is the users entrypoint for job scheduling on HPC systems”

This allows sharing computing resources efficiently and fairly.

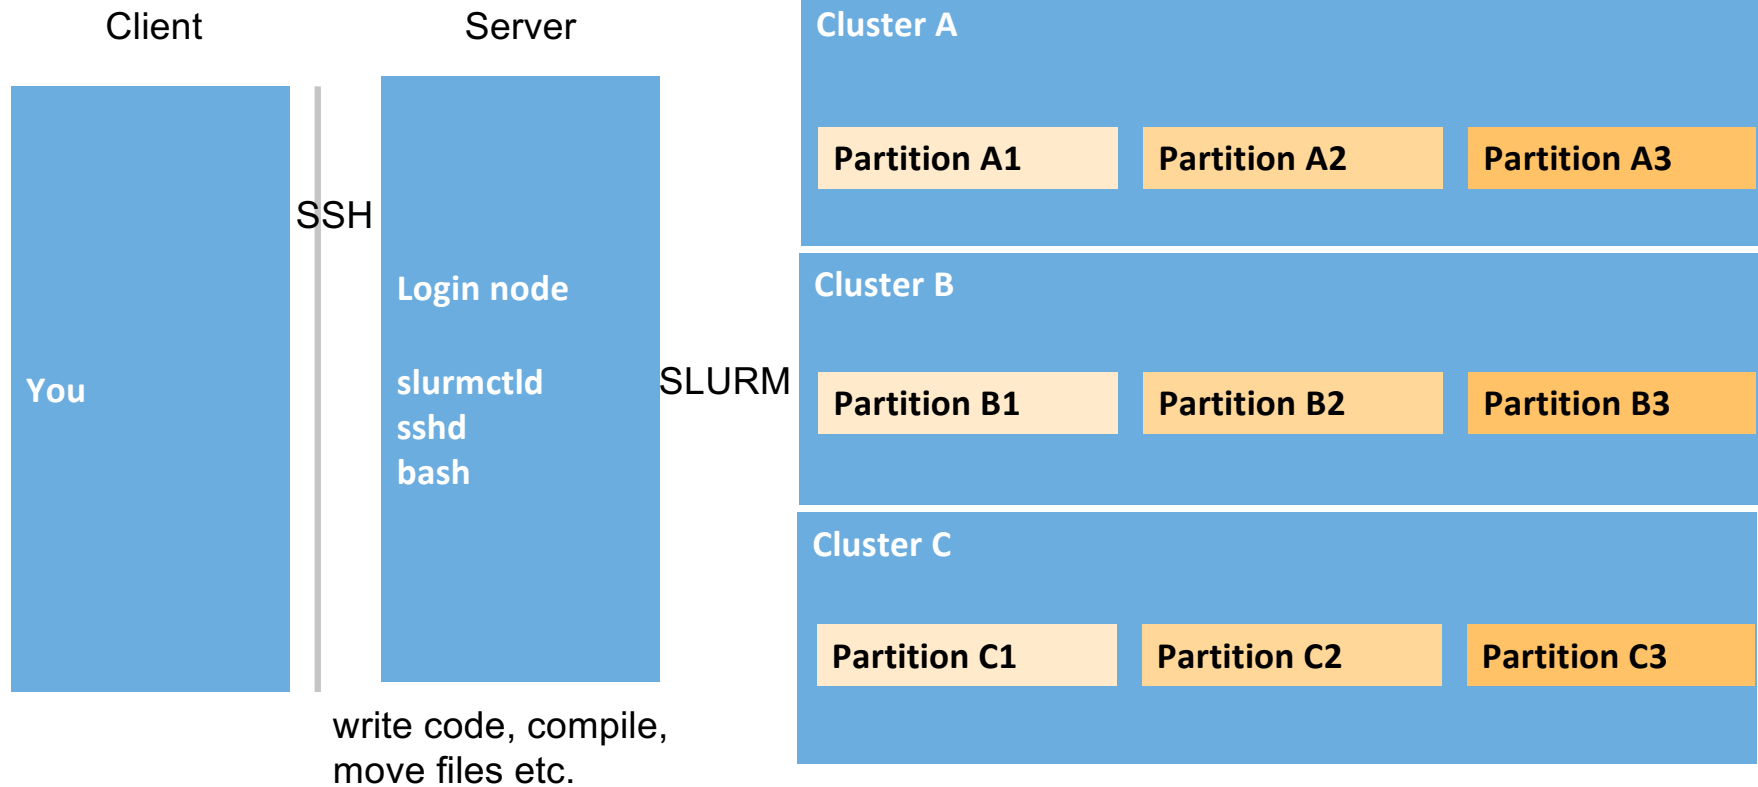
It is fun to use.



*image source: <https://www.syfy.com/futurama/season-2/blogs/episode-recap-fry-and-the-slurm-factory>



Minimal Cluster Architecture





Workload Management
Commands

Slurm

- salloc
- sbatch
- scancel
- sinfo
- srun
- sacct
- squeue





Workload Management

Slurm commands



Scheduling

- Run parallel jobs on the cluster
- Interactive and Batch mode

srun, salloc, sbatch, scancel

Monitoring

- Health of jobs, nodes etc.

sinfo, squeue, scontrol

Accounting

- Quotas, charging etc.


sacct

Resource allocation

- Ensure fairness

Priority management

- Testing, QOS etc.



Upcoming Courses



AI Training Series

- <https://doku.lrz.de/data-analytics-big-data-and-ai-training-courses-at-lrz-10333405.html>

Information on further HPC courses:

- by LRZ: <http://www.lrz.de/services/compute/courses/>
- by the Gauss Centre of Supercomputing (GCS): <http://www.gauss-centre.eu/training>
- by German Centres (collected by the Gauß-Allianz): <https://hpc-calendar.gauss-allianz.de/>
- by the EuroCC National Competence Centres: <https://www.eurocc-access.eu/services/training/>
- by European Centres: <https://hpc-portal.eu/upcoming-events-courses>