
Slurm

Thought Experiment

- Executing a program in your PC (say laptop):
  - You just simply run a *fun_app* (./ *fun_app*)

- 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 think about running your *fun_app* in your Souper Computer and **create** (hypothetically) a software ;)

*Figure: https://ianbelcher.me/tech-blog/creating-a-bare-bones-on-premises-k8s-cluster-from-old-hardware

Thought Experiment

S1: How will you run your *fun_app* in all/some laptops (let‘s call it a resource)?

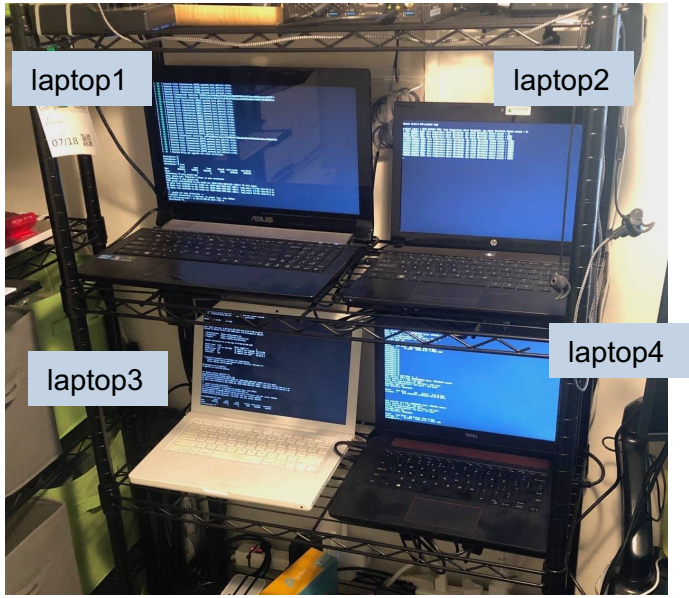

You decided to **create** a **software (bash script?)** for ease of use (what about giving a following name ***Simple Linux Resource Manager***? ***slrm***)

./slrm –n 2 fun_app

Resource Management Software

*Figure: https://ianbelcher.me/tech-blog/creating-a-bare-bones-on-premises-k8s-cluster-from-old-hardware

Thought Experiment

S2: You need to run the *program* only in laptop with GPU (White laptop)

slrm –n 2 fun_app -usegpu

Generic Resources - Resource Management – Resource allocation

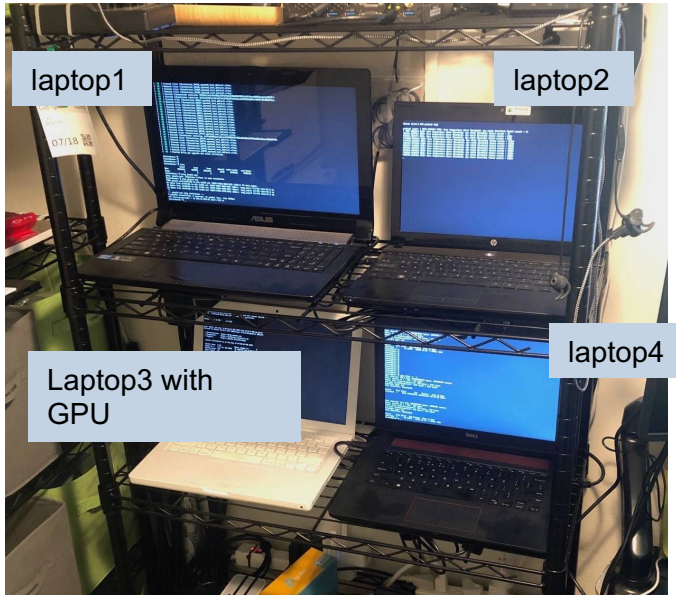

*Figure: https://ianbelcher.me/tech-blog/creating-a-bare-bones-on-premises-k8s-cluster-from-old-hardware

Thought Experiment

S3: You need to give access to your friends, but you don‘t want to give access to your favorite laptop ;)

slrm –n 2 fun_app –partition1 –user admin

Partitions and Access permissions

Partition 2

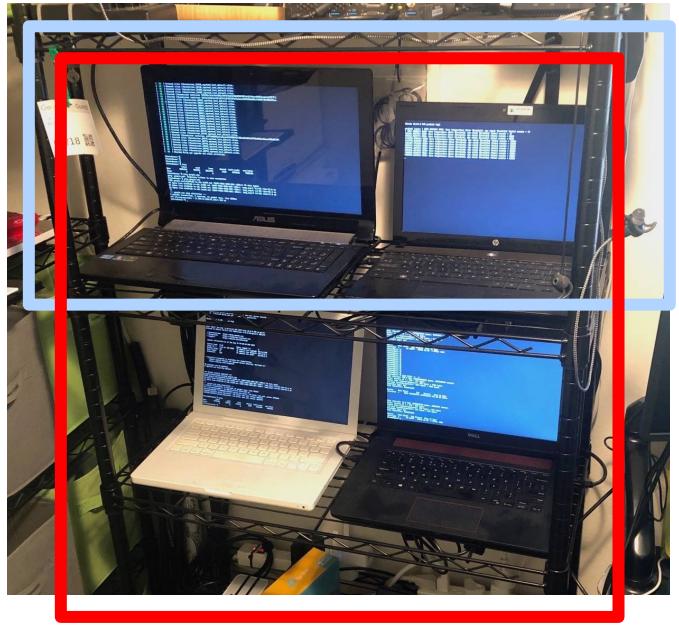

Partition 1

*Figure: https://ianbelcher.me/tech-blog/creating-a-bare-bones-on-premises-k8s-cluster-from-old-hardware

Thought Experiment

S4: You (or your cluster) became popular and now more people want to access it.

Job Scheduling and Job Queues

*Figure: https://ianbelcher.me/tech-blog/creating-a-bare-bones-on-premises-k8s-cluster-from-old-hardware

Thought Experiment

S5: Two of your best friends are accessing it and you want to give them faster access ;)

slrm –n 2 fun_app –partition1 –user priorityuser

Job Priority

*Figure: https://ianbelcher.me/tech-blog/creating-a-bare-bones-on-premises-k8s-cluster-from-old-hardware

Thought Experiment

S6: Now you might be worrying about the electricity costs, health of your laptops!!!

Monitoring and Accounting

*Figure: https://ianbelcher.me/tech-blog/creating-a-bare-bones-on-premises-k8s-cluster-from-old-hardware

Workload Management

Features

Scheduling

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

Monitoring

- Health of jobs, nodes etc.

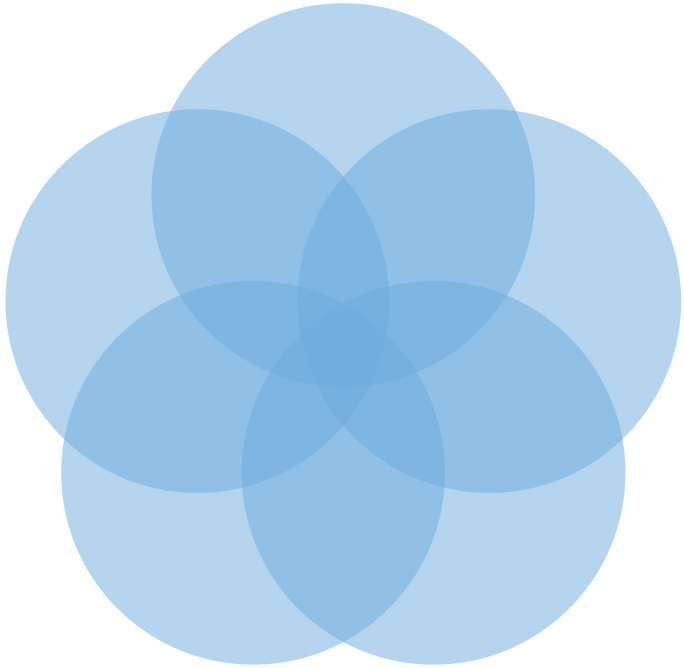

Resource allocation

- Ensure fairness

Accounting

- Quotas, charging etc.

Priority management

- Testing, QOS etc.

(Simple Linux Utility for Resource Management)

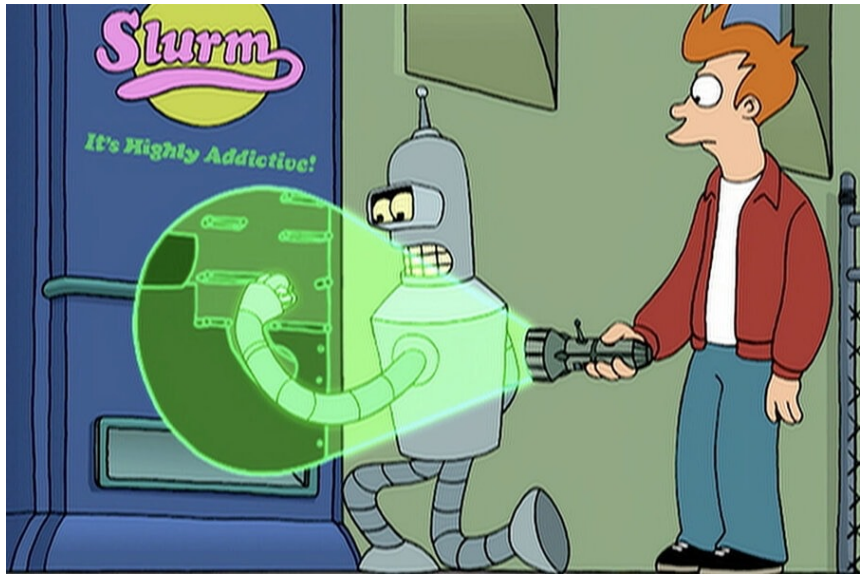

"Slurm is an open source workload management system and job scheduler. It is the users entrypoint for job scheduling on HPC systems"

This allows sharing computing resources efficiently and fairly.

It is fun to use.

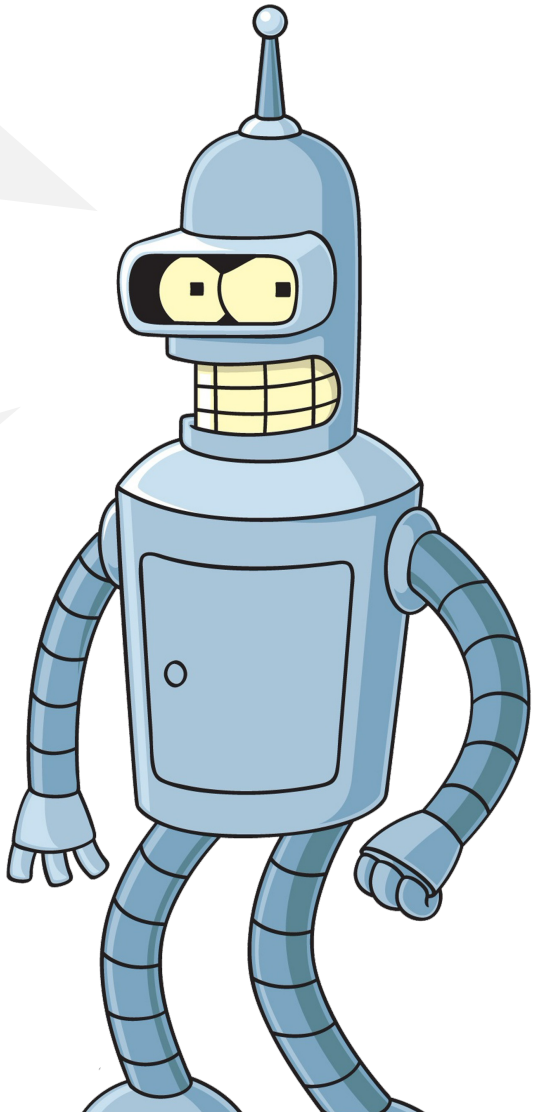

*image source: https://www.syfy.com/futurama/season-2/blogs/episode-recap-fry-and-the-slurm-factory


Slurm
Workload Management
Commands
Slurm
sinfo
salloc
srun
sbatch
scancel
sacct
squeue

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

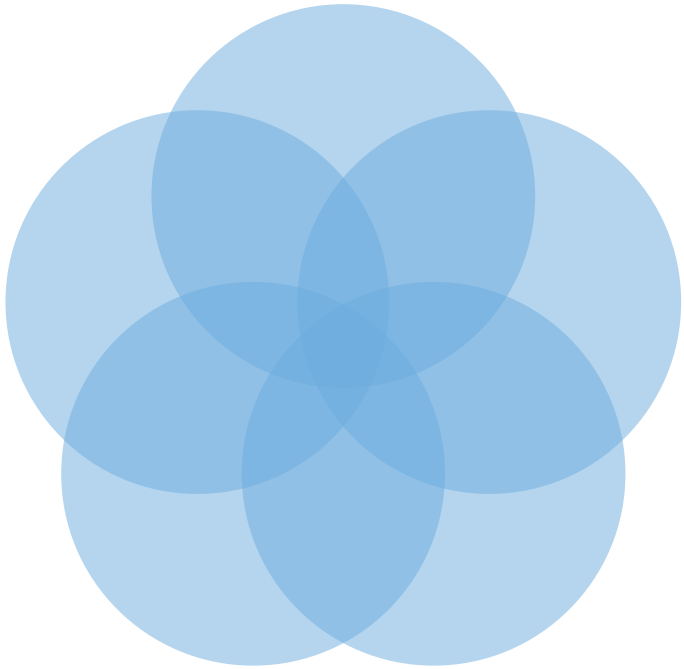

Resource allocation

- Ensure fairness

Accounting

- Quotas, charging etc.

sacct

Priority management

- Testing, QOS etc.

Minimal Cluster Architecture