HPC series intro

Author:
Sami Ait Ali Oulahcen

Rabat, Morocco
07 Aug 2023

Credit: Frontier Supercomputer https://www.flickr.com/photos/olcf/52117623843/

HPC intro

Personal computer

vs

large computer OR server

vs

HPC cluster OR supercomputer

HPC intro

TOP 500 supercomputer “Frontier” currently at 1.2 exaFLOPS

1 to 10s of TeraFLOPS

10s to 100s of TeraFLOPS

100s of Tera FLOPS to 100s of PetaFLOPS

HPC intro

$$FLOPS = racks \times \frac{nodes}{rack} \times \frac{sockets}{node} \times \frac{cores}{socket} \times \frac{cycles}{second} \times \frac{FLOPs}{cycle}$$

How does it work ?

- Cluster contains login/master nodes + compute nodes + storage nodes + GPU nodes …
- User talks to login node
- Login node distributes compute tasks to compute nodes with the help of scheduler (usually SLURM).

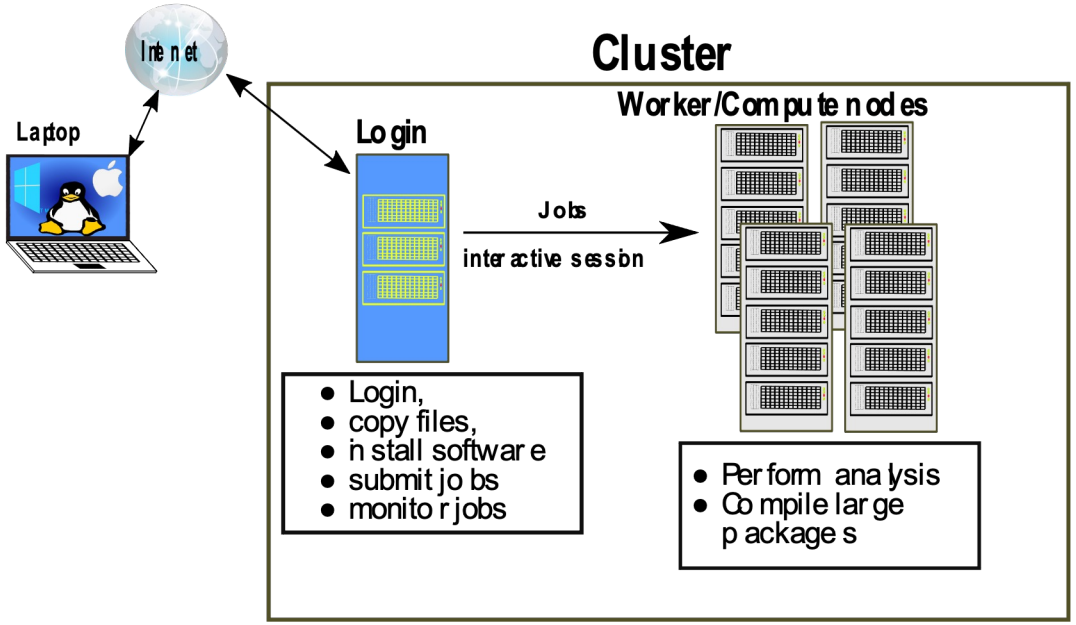

Accessing the cluster

- Use SSH to connect to HPC (login node)
  => natively in linux & macOS

```
ssh username@hpc.marwan.ma
```

  => for windows use **putty** or **mobaXterm**

- Explore resources on all compute nodes:

```
$ sinfo (-N nodes --long)
```

- Compare with resources on login node:

```
$ nproc --all
$ free -m
$ df -Th
```

Transfering files

- Filezilla (sftp) can be used to transfer files to and from the HPC cluster
  => available for linux/macOS/windows
- Many alternatives exist in every OS

- For Linux, the command line **scp** can be used

```
$ scp -r source destination
```

- We also use scp to transfer files between servers

Choosing software/modules

- Working with modules

See the available modules : `module avail`

Obtain info on a particular module: `module whatis + something`

Load a module : `module load + something`

List the loaded modules : `module list`

Unload a module : `module unload + something`

Unload all the modules : `module purge`

- Check if running the correct version of software:

`which somesoft`

Running a simulation

- See available partitions on an HPC cluster:

```
$ sinfo
```

- Reserve a compute node with one of these commands:

```
$ srun -p <partition> --pty bash
$ salloc -N 1 -n 1 -p <partition> bash
```

Here choose a short jobs partition for quick access

- Run your simulations:

```
$ ./myDemoLab.sh
```

- Release the compute node when done:

```
$ exit
```

Launch a simulation as a job using slurm

- Scheduler chooses ressources automatically
- Possiblity to choose resources manually like number of CPUs, amount of RAM, GPU
- Jobs are launched in background

  => session can be closed (you can turn off your computer)

  => results are automatically retrieved
- Run a job:

```
$ sbatch myDemoLab.sh
```

- See the queue:

```
$ squeue (-u username)
```

- Cancel a job:

```
$ scancel <job_id>
```

Structure of a slurm batch script

```
#!/bin/bash
#SBATCH --job-name=FIRST-SIMULATION
#SBATCH --ntasks=1 --cpus-per-task=2
#SBATCH --mem-per-cpu=2G
#SBATCH --time=00:15:00
#SBATCH --mail-user=<email>
echo ”hello world”
sleep 5
echo “done”
```

The *#!/bin/bash* is called a shebang, and it indicates on which shell the script should be run.

Slurm in action

Credit: https://hpc-wiki.info/hpc/Admin_Guide_Scheduling_Algorithms

Case study

- Benchmark: computing electronic structures and nano-scale material modelisation using Quantum Espresso

Number of CPU	Execution time
8	18h30m
16	9h22m
32	7h12m
64	6h21m
128	1h47m
190	1h47m
256	1h49m

Lab

- Task 1: login to your HPC environment and change your password
- Task 2: transfer a file from your laptop to the /home/$USER directory in your HPC environment
- Task 3: execute commands from previous slides to check the number of compute nodes and their resources (CPU, RAM, GPU, storage), and check the available modules/software on the cluster