

IFB-core cluster in practice

Frédéric Jarlier

September 2025



Table of Contents


- 1 HPC clusters
- 2 In practice
- 3 Slurm usage
- 4 Modules
- 5 Conclusion

The slides are adapted from previous presentations by Alix Silvert, Magali Hennion and Julien Rey and inspired by IFB training material.

What is a HPC cluster for ?

- High hardware resource needs
- Long-running analyses
- A lot of similar analyses
- Shared work between users
- Free your desktop from heavy tasks
- Improve reproducibility

What is a HPC cluster ?



How to talk to it?

SSH Connect to the cluster using the following line in your local terminal, replacing 'username' with your personal username. Then enter your password.

```
ssh username@core.cluster.france-bioinformatique.fr
```

Open OnDemand web interface Open a web browser and go to <https://ondemand.cluster.france-bioinformatique.fr>.
Type in your login and password. Grant access.

sftp

- To navigate in your remote folders with a local file manager
- To edit files with your local editor

Open a file manager and go to **sftp://core.cluster.france-bioinformatique.fr**.

Note about security

Best practices, change regularly your password

```
username@clust-slurm-client2:~$ ~ : passwd
```

Where you can go, write, or execute

User environment

/shared/home/username

Computations (hot data)

/shared/projects/projectname

Common banks (read-only)

/shared/bank/

Note about quotas

The disk space available for projects is limited by quotas. To check the space available for your projects, use the following commands:

```
fjarlier@clust-slurm-client2:~$ statusBars
2417_wf4bioinfo [##-----]      25 /    250 GB
2521_wf4bioinfo [-----]        0 /    250 GB

      HOME [-----]          1 /    100 GB
Update: 2025-09-08 14:00 - Your current default account is 2521_wf4bioinfo - More info: statusBars --help
```

```
fjarlier@clust-slurm-client2:~$ gid=$(getent group 2417_wf4bioinfo | cut -f 3 -d ':')
fjarlier@clust-slurm-client2:~$ lfs quota -h -p $gid /shared/projects/2417_wf4bioinfo
Disk quotas for prj 166486 (pid 166486):
  Filesystem      used  quota  limit  grace  files  quota  limit  grace
 /shared/projects/2417_wf4bioinfo
                   25.12G   250G   375G      -   27158  250000  375000      -
```

All jobs **must** be launched from a **project** directory.

About Slurm




Slurm is the job scheduling system.

It is what will take your code and distribute it on the computing nodes, while ensuring they have the CPU(s) and RAM that you asked for.

It requires specific commands to run (srun, sbatch, salloc, etc...).

You **must** launch jobs with Slurm.

Flowsheet



srun

Launch a (simple) interactive job.

```
fjarlier@clust-slurm-client2:~$ srun hostname
```

Some parameters can be added to the command line:

--partition/-p: request a specific partition

--account/-A: select the (project) account

--cpus-per-task/-c: request that ncpus be allocated (default: 1 cpu)

--mem-per-cpu: specify the required memory per cpu (default: 2GB)

Example

```
fjarlier@clust-slurm-client2:~$ srun -A 2521_wf4bioinfo -p fast hostname
```

sbatch

Launch more complex jobs using scripts.

myscript.sbatch

```
#!/bin/bash
#SBATCH --partition=fast
#SBATCH --account=2521_wf4bioinfo
#SBATCH --cpus-per-task=1
#SBATCH --mem-per-cpu=4GB
#SBATCH --output=resultat.log
hostname
```

Example

```
fjarlier@clust-slurm-client2:~$ sbatch myscript.sbatch
```

About associations

Cluster, account and partition must match your user's associations.

To check your associations:

```
fjarlier@clust-slurm-client2:~$ sacctmgr show user $USER withassoc
```

User	Def Acct	Admin	Cluster	Account	Partition	Share	Priority	MaxJobs
fjarlier	2521_wf4b+	None	core	2417_wf4b+	fast	1	normal	normal
fjarlier	2521_wf4b+	None	core	2417_wf4b+	long	1	normal	normal
fjarlier	2521_wf4b+	None	core	2521_wf4b+	fast	1	normal	normal
fjarlier	2521_wf4b+	None	core	2521_wf4b+	long	1	normal	normal

squeue : list submitted jobs on the cluster

```
fjarlier@clust-slurm-client2:~$ squeue
```

JOBID	PARTITION	NAME	USER	ST	TIME	NODES	NODELIST (REASON)
61189751	bigmem	test_mod	mrobin	R	16:55:20	1	cpu-node-100
61185490	bigmem	binning	plecoq	R	1-15:11:07	1	cpu-node-69
61185491	bigmem	binning	plecoq	R	1-15:11:07	1	cpu-node-69
61170104	bigmem	iqtFULL	romainl	R	3-04:45:29	1	cpu-node-69
61199999	bigmem	nf-egapx	rguyot	R	1:48:13	1	cpu-node-69
61199992	bigmem	bash	rguyot	R	1:50:40	1	cpu-node-69

Some parameters can be added to filter jobs or show more infos:

- partition/-p**: specify the partition to view
- user/-u**: request jobs from a list of users
- format/-o**: specify the information to be displayed

```
fjarlier@clust-slurm-client2:~$ squeue --me -p fast
fjarlier@clust-slurm-client2:~$ squeue -o "%10i %.9P %.8j %.12u %.15T %.15M %.15l %.6D %.20R %.6C %.10b
```

scancel

Kill a job

```
fjarlier@clust-slurm-client2:~$ scancel job_id
```

Cancel all my jobs

```
fjarlier@clust-slurm-client2:~$ scancel -u <my_user_name>
```

or

```
fjarlier@clust-slurm-client2:~$ scancel --me
```

sacct

Display accounting data for your running and completed jobs:

```
fjarlier@clust-slurm-client2:~$ sacct --format=JobID,JobName,Start,Elapsed,CPUTime,NCPUS,NodeList,MaxRSS,ReqMem,State
```

JobID	JobName	Start	Elapsed	CPUTime	NCPUS	NodeList	MaxRSS	ReqMem	State
61200956	hostname	2025-09-08T15:19:48	00:00:00	00:00:00	1	cpu-node-92		2000M	COMPLETED
61200956.0	hostname	2025-09-08T15:19:48	00:00:00	00:00:00	1	cpu-node-92	3756K		COMPLETED
61200963	hostname	2025-09-08T15:21:27	00:00:01	00:00:01	1	cpu-node-97		2000M	COMPLETED
61200963.0	hostname	2025-09-08T15:21:28	00:00:00	00:00:00	1	cpu-node-97	3708K		COMPLETED

reportseff

Report a job efficiency

```
fjarlier@clust-slurm-client2:~$ module load reportseff
fjarlier@clust-slurm-client2:~$ reportseff 61200956
```

JobID	State	Elapsed	TimeEff	CPUEff	MemEff
61200956	COMPLETED	00:00:00	0.0%	---	0.2%

Analyse a folder of slurm outputs

```
fjarlier@clust-slurm-client2:~$ reportseff --format "+Start,CPUTime,NCPUS,NodeList,MaxRSS,ReqMem" --modified-sort slurm_output/
```

JobID	State	Elapsed	TimeEff	CPUEff	MemEff	Start	CPUTime	NCPUS	NodeList	MaxRSS	ReqMem
multiqc--41907184.out	COMPLETED	00:00:07	2.3%	57.1%	1.6%	2024-09-24T12:23:17	00:00:07	1	cpu-node-45	3184K	191M
fastqc-sample=SRR3099585_chr18-41907179.out	COMPLETED	00:00:25	8.3%	64.0%	1.1%	2024-09-24T12:22:47	00:00:25	1	cpu-node-38	3284K	286M
fastqc-sample=SRR3105697_chr18-41907180.out	COMPLETED	00:00:15	5.0%	80.0%	1.1%	2024-09-24T12:22:47	00:00:15	1	cpu-node-39	3260K	286M
fastqc-sample=SRR3099587_chr18-41907181.out	COMPLETED	00:00:14	4.7%	85.7%	1.1%	2024-09-24T12:22:47	00:00:14	1	cpu-node-40	3200K	286M
fastqc-sample=SRR3105698_chr18-41907183.out	COMPLETED	00:00:13	4.3%	69.2%	1.1%	2024-09-24T12:22:47	00:00:13	1	cpu-node-40	3192K	286M
fastqc-sample=SRR3099586_chr18-41907178.out	COMPLETED	00:00:12	4.0%	66.7%	1.1%	2024-09-24T12:22:47	00:00:12	1	cpu-node-45	3208K	286M
fastqc-sample=SRR3105699_chr18-41907182.out	COMPLETED	00:00:12	4.0%	66.7%	1.1%	2024-09-24T12:22:47	00:00:12	1	cpu-node-40	3268K	286M
multiqc--41907169.out	OUT_OF_MEMORY	00:00:06	2.0%	10.7%	3.3%	2024-09-24T12:21:06	00:00:06	1	cpu-node-45	3240K	95M
fastqc-sample=SRR3099585_chr18-41907159.out	COMPLETED	00:00:33	11.0%	54.5%	75.7%	2024-09-24T12:20:25	00:00:33	1	cpu-node-38	221708K	286M
fastqc-sample=SRR3105697_chr18-41907160.out	COMPLETED	00:00:19	6.3%	73.7%	1.1%	2024-09-24T12:20:25	00:00:19	1	cpu-node-39	3212K	286M
fastqc-sample=SRR3105698_chr18-41907163.out	COMPLETED	00:00:17	5.7%	82.4%	1.1%	2024-09-24T12:20:25	00:00:17	1	cpu-node-40	3224K	286M
fastqc-sample=SRR3099587_chr18-41907161.out	COMPLETED	00:00:15	5.0%	80.0%	1.1%	2024-09-24T12:20:25	00:00:15	1	cpu-node-40	3192K	286M
fastqc-sample=SRR3105699_chr18-41907162.out	COMPLETED	00:00:14	4.7%	71.4%	1.1%	2024-09-24T12:20:25	00:00:14	1	cpu-node-40	3244K	286M
fastqc-sample=SRR3099586_chr18-41907158.out	COMPLETED	00:00:15	5.0%	53.3%	1.1%	2024-09-24T12:20:25	00:00:15	1	cpu-node-45	3304K	286M
multiqc--41907141.out	COMPLETED	00:00:24	0.3%	25.0%	0.6%	2024-09-24T12:15:57	00:00:24	1	cpu-node-38	3228K	500M
fastqc-sample=SRR3105698_chr18-41907138.out	COMPLETED	00:00:18	0.2%	72.2%	0.6%	2024-09-24T12:15:35	00:00:18	1	cpu-node-34	3244K	500M
fastqc-sample=SRR3105699_chr18-41907137.out	COMPLETED	00:00:16	0.2%	68.8%	0.6%	2024-09-24T12:15:35	00:00:16	1	cpu-node-34	3228K	500M
fastqc-sample=SRR3099586_chr18-41907133.out	COMPLETED	00:00:15	0.2%	66.7%	0.6%	2024-09-24T12:15:35	00:00:15	1	cpu-node-39	3144K	500M
fastqc-sample=SRR3099585_chr18-41907134.out	COMPLETED	00:00:15	0.2%	73.3%	0.6%	2024-09-24T12:15:35	00:00:15	1	cpu-node-40	3272K	500M
fastqc-sample=SRR3099587_chr18-41907136.out	COMPLETED	00:00:14	0.2%	71.4%	0.6%	2024-09-24T12:15:35	00:00:14	1	cpu-node-40	3208K	500M
fastqc-sample=SRR3105697_chr18-41907135.out	COMPLETED	00:00:15	0.2%	66.7%	0.6%	2024-09-24T12:15:35	00:00:15	1	cpu-node-40	3232K	500M

Job arrays

Job arrays offer a mechanism for launching a large number of tasks simultaneously. Each task of the job will have the environment variable `$SLURM_ARRAY_TASK_ID` set to its array index value.

`myscript.sbatch`

```
#!/bin/bash
#SBATCH --partition=fast
#SBATCH --account=2521_wf4bioinfo
#SBATCH --output=resultat_%a.log
#SBATCH --array=1-3
case "$SLURM_ARRAY_TASK_ID" in
  1) fruit='orange';;
  2) fruit='apple';;
  3) fruit='banana';;
esac
echo $fruit
```

Job arrays

```
fjarlier@clust-slurm-client2:~$ sbatch myscript.sbatch
```

Results:

```
fjarlier@clust-slurm-client2:~$ ls
resultat_1.log  resultat_2.log  resultat_3.log

fjarlier@clust-slurm-client2:~$ tail resultat_*
==> resultat_1.log <==
orange

==> resultat_2.log <==
apple

==> resultat_3.log <==
banana
```

Job arrays

Use Bash lists

myscript.sbatch

```
#!/bin/bash
#SBATCH --partition=fast
#SBATCH --account=2521_wf4bioinfo
#SBATCH --output=resultat_%a.log
#SBATCH --array=0-200%50
#SBATCH --cpus-per-task=2
INPUTS=(*.fq.gz)
fastqc ${INPUTS[$SLURM_ARRAY_TASK_ID]}
```

Job arrays

```
#SBATCH --array=0-15
```

= 16 jobs (\$SLURM_ARRAY_TASK_ID: from 0 to 15 included).

```
#SBATCH --array=10-16:2
```

= 4 jobs (\$SLURM_ARRAY_TASK_ID: 10,12,14,16).


```
#SBATCH --array=2,3-7:2,11,13
```

= 6 jobs (\$SLURM_ARRAY_TASK_ID: 2,3,5,7,11,13).

```
#SBATCH --array=1-10000%32
```

= 10 000 jobs, max 32 simultaneous jobs

Complex workflows



Use workflow managers such as Snakemake or Nextflow.

Modules



Tools are installed on the cluster in virtual environments

- each tool has its own dependencies (libraries) and it's not possible to make them all coexist in the same environment
- reproducibility: keep several versions of a tool



Conda environments



Containers (Apptainer)

Modules

They can be loaded with the module command.

Look for the different versions of multiqc:

```
fjarlier@clust-slurm-client2:~/test2$ module avail multiqc
----- /shared/software/modulefiles -----
multiqc/1.3  multiqc/1.6  multiqc/1.7  multiqc/1.9  multiqc/1.11  multiqc/1.12  multiqc/1.13
```

Load an environment:

```
fjarlier@clust-slurm-client2:~$ module load multiqc/1.29
```

List loaded environments:

```
fjarlier@clust-slurm-client2:~/test2$ module list
Currently Loaded Modulefiles:
 1) reportseff/2.7.6   2) multiqc/1.29
```


Useful resources

To find out more, the SLURM manual : `man sbatch` or
<https://slurm.schedmd.com/sbatch.html>

Ask for help or signal problems on the cluster : <https://community.france-bioinformatique.fr>

IFB cluster documentation:
<https://ifb-elixirfr.gitlab.io/cluster/doc/>

Thanks

