LIÈGE université
GIGA

Clusters and parallel computing

GIGA Doctoral School, 2019

Martin Grignard, Ing
GIGA - CRC In Vivo Imaging

Clusters and parallel computing

What ?

Key concepts

Why ?

Problems it can solve

How ?

Quick start and examples

---

mar.grignard@uliege.be : GIGA Doctoral School 2019 “Clusters and parallel computing”

High-performance computing (HPC)

Distributed system

System whose components are located on different computers, communicating through a network.

Cluster Set of connected computers.

Grid Set of connected clusters.

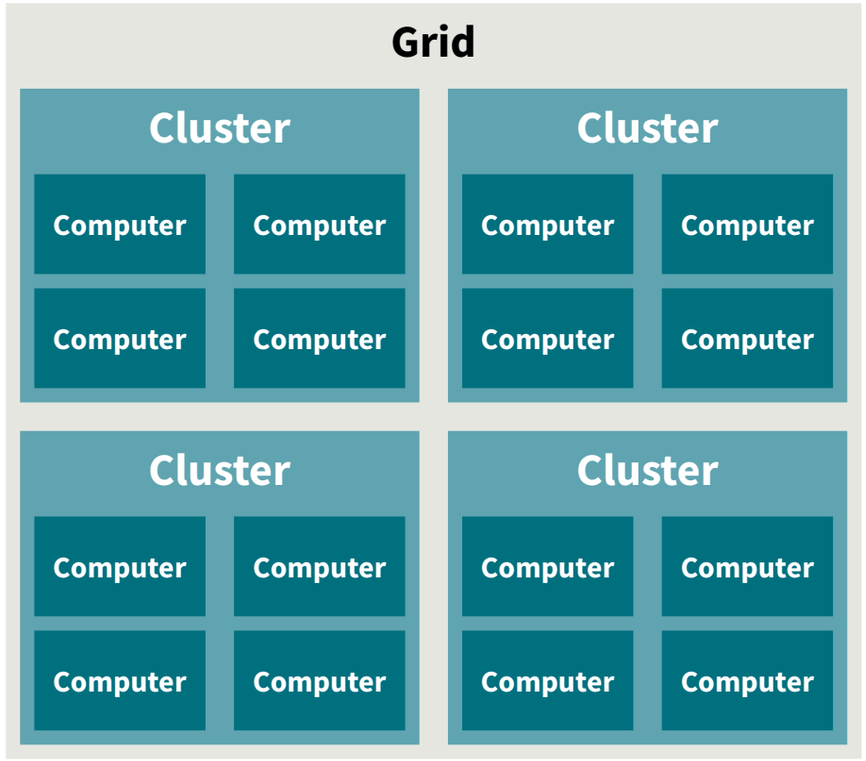

mar.grignard@uliege.be : GIGA Doctoral School 2019 “Clusters and parallel computing”

High-performance computing (HPC)

Nodes Perform the computations.

File system Store and allow data access.

Network Allow nodes to communicate.

Scheduler

Manage jobs as a function of the available ressources.

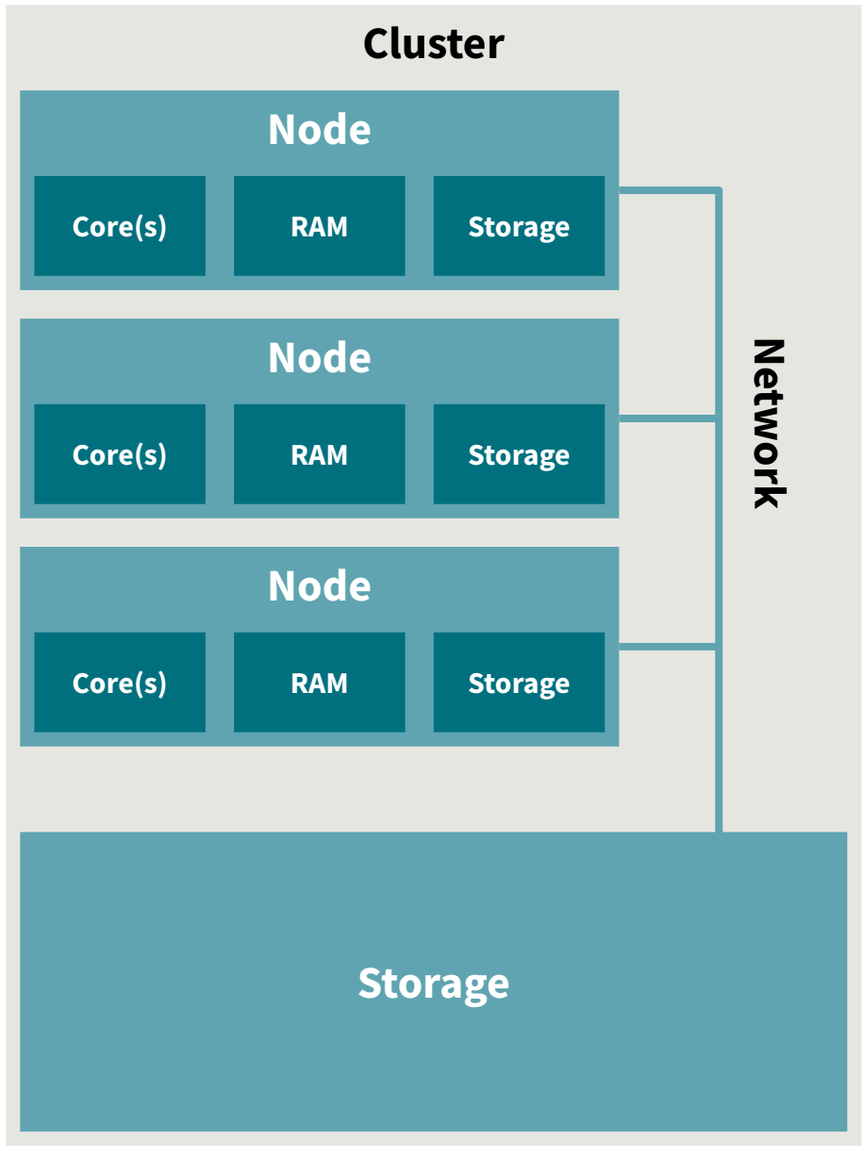

mar.grignard@uliege.be : GIGA Doctoral School 2019 “Clusters and parallel computing”

LIÈGE université
GIGA

> # Why would I bother ?
> # I already have a multi-core workstation* !
> * That costs a lot of money

- Any sane PhD student

---

mar.grignard@uliege.be : GIGA Doctoral School 2019 “Clusters and parallel computing”

Use cases : Many-to-many

> I have to apply the same process to many chunks of data.
> I want to run fast(er) over my whole dataset.

Example

Preprocessing of MRI data
Data: MRI data of 360 subjects
Proc: 1 day of CPU time/subject/CPU

Workstation (12 cores) : 30 days

Cluster (360 cores) : 1 day

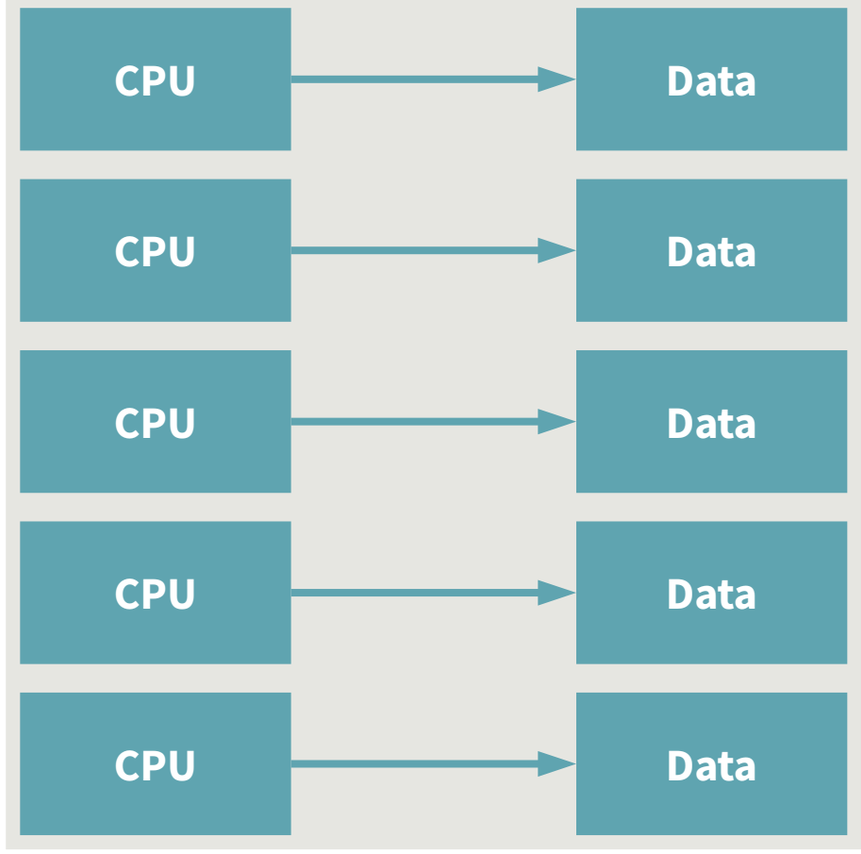

Use cases : Many-to-one

> I have to run the same process with many parameters sets on the same dataset.

Example	Monte Carlo integration
Example	Parameter sweep

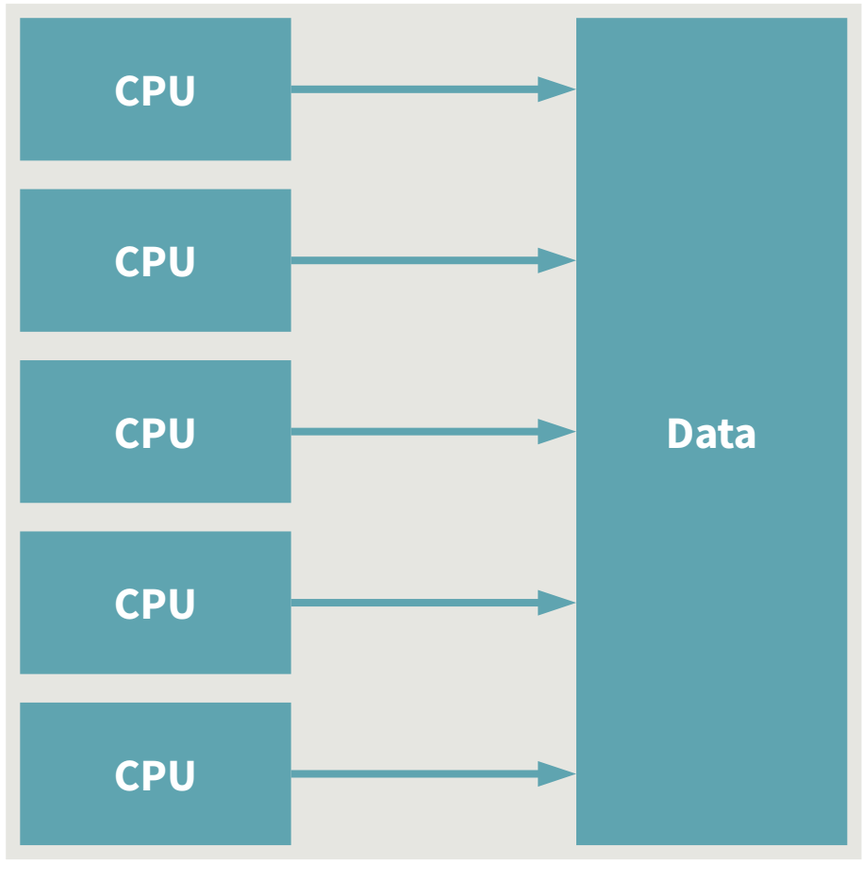

Use cases : (Almost)One-to-one

> My dataset is so huge it does not fit
> in my workstation's RAM.

Example	Data mining
Example	Linear algebra
Example	Machine learning

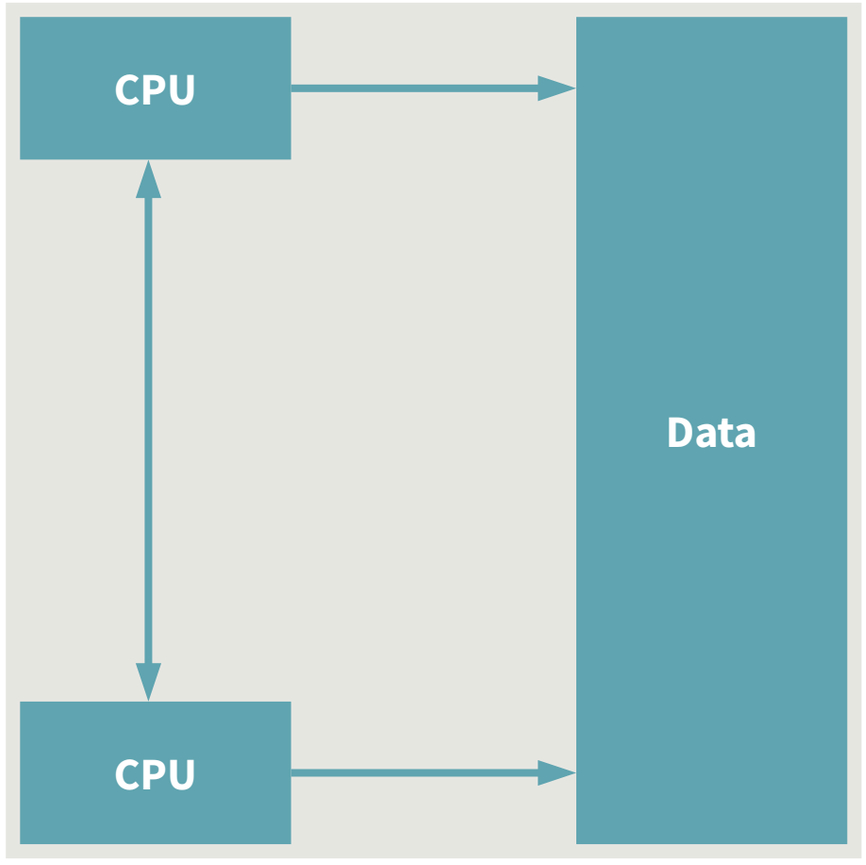

CECI clusters

Members

Uliège, UCLouvain, ULB,
Umons, UNamur

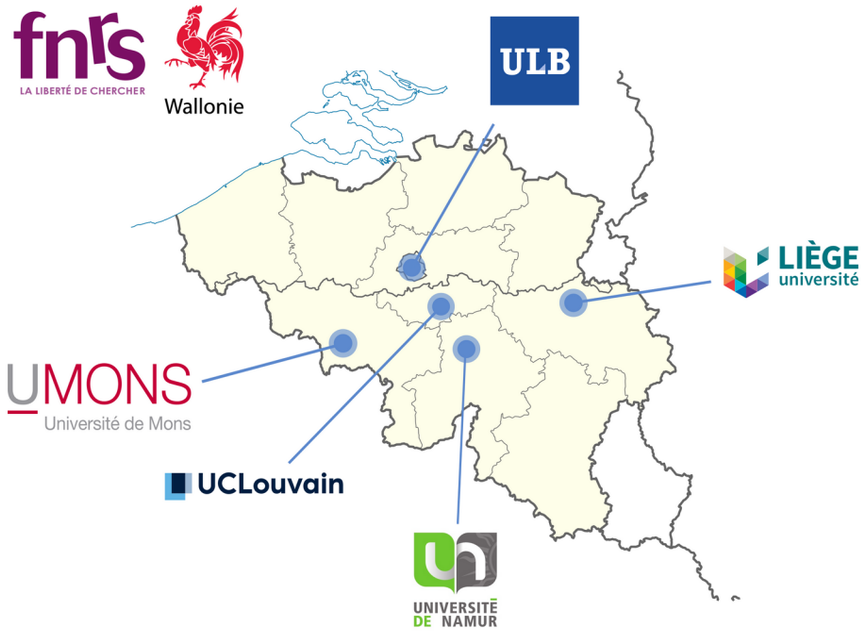

Website http://www.ceci-hpc.be/

Services

Support, training sessions,
documentation…

LIÈGE université
GIGA

CECI clusters

Cluster	Host	CPU type	CPU count*	RAM/node	Network	Filesystem**	Accelerator	Max time	Preferred jobs***
Hercules2	UNamur	Epyc 2 GHz SandyBridge 2.20 GHz	1024 (30 x 32 + 2 x 64) 512 (32 x 16)	64 GB..2 TB	10 GbE	NFS 20 TB	None	15 days	⋮ serial / ≡ SMP
Dragon2	UMons	SkyLake 2.60 GHz	592 (17 x 32 + 2 x 24)	192..384 GB	10 GbE	RAID0 3.3 TB	4x Volta V100	41 days	⋮ serial / ≡ SMP
Lemaitre3	UCL	SkyLake 2.3 GHz Haswell 2.6 GHz	1920 (80 x 24) 112 (4 x 28)	95 GB 64 GB	Omnipath	FHGFS 580 TB	None	2 days 6 hours	⋮⋮⋮ MPI
NIC4	ULiège	SandyBridge 2.0 GHz IvyBridge 2.0 GHz	2048 (120 x 16 + 8 x 16)	64 GB	QDR Ib	FHGFS 144 TB	None	3 days	⋮⋮⋮ MPI
Vega	ULB	Bulldozer 2.1 GHz	2752 (43 x 64)	256 GB	QDR Ib	GPFS 70 TB	2x Tesla M2090	14 days	⋮ serial / ≡ SMP / ⋮⋮⋮ MPI
Hercules	UNamur	SandyBridge 2.20 GHz	512 (32 x 16)	64..128 GB	GbE	NFS 20 TB	None	63 days	⋮ serial / ≡ SMP
Dragon1	UMons	SandyBridge 2.60 GHz	416 (26 x 16)	128 GB	GbE	RAID0 1.1 TB	4x Tesla C2075	41 days	⋮ serial / ≡ SMP
Lemaitre2	UCL	Westmere 2.53 GHz	1380 (115 x 12)	48 GB	QDR Ib	Lustre 120 TB	3x Quadro Q4000	3 days	⋮⋮⋮ MPI
Hmem	UCL	MagnyCours 2.2 GHz	816 (17 x 48)	128..512 GB	QDR Ib	FHGFS 30 TB	None	15 days	≡ SMP

CECI clusters : Access

> Would you be my new best friend ?

Terminal

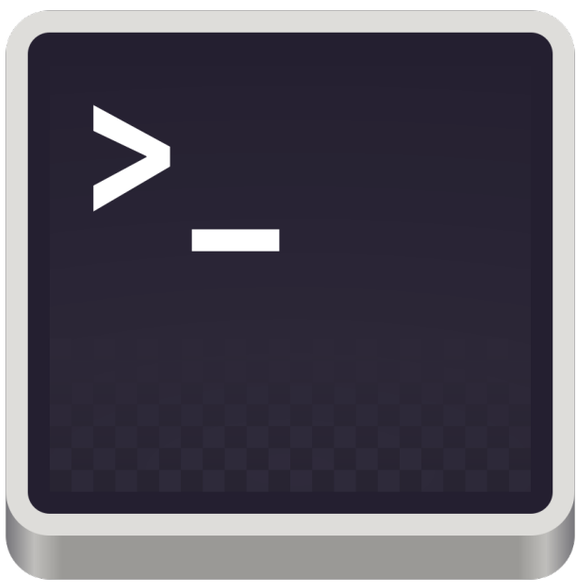

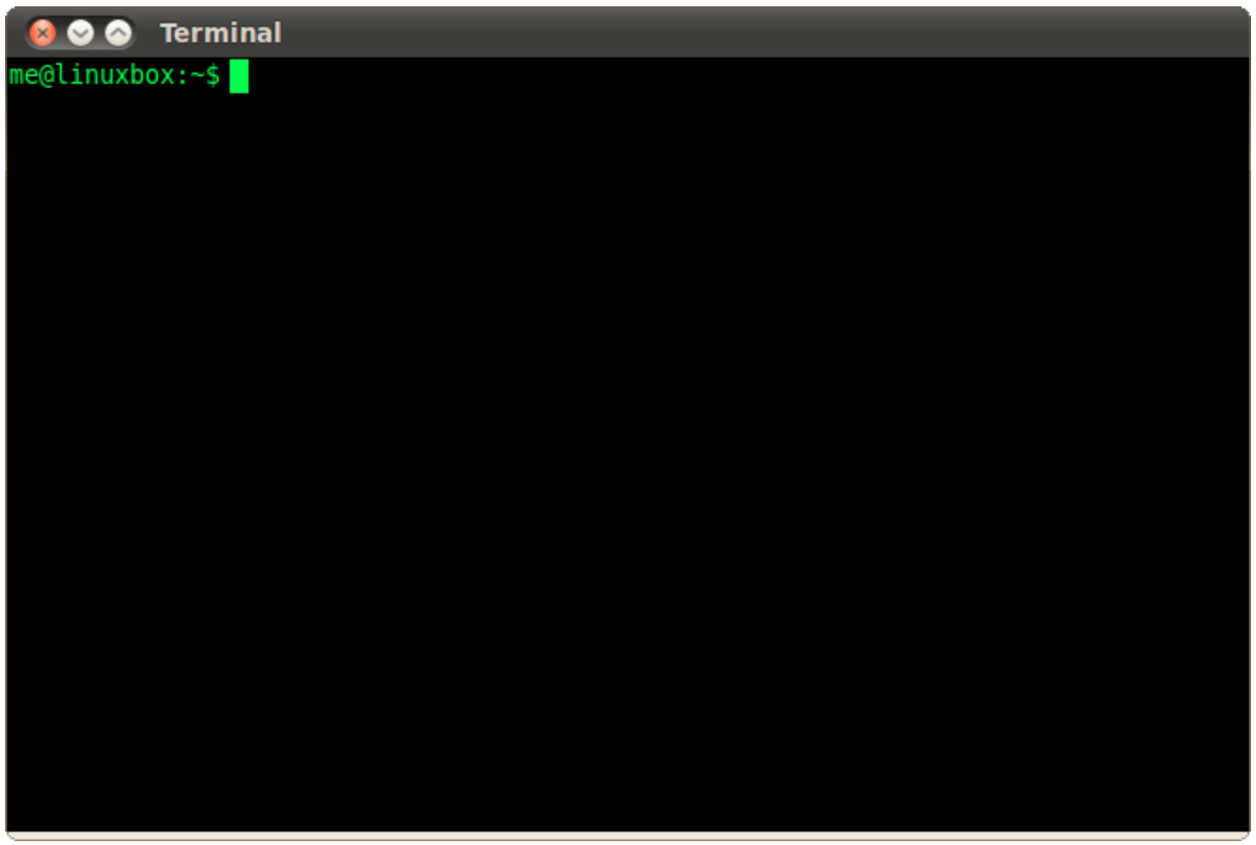

mar.grignard@uliege.be : GIGA Doctoral School 2019 “Clusters and parallel computing”

CECI clusters : Connection

Account	https://login.ceci-hpc.be/init/

SSH

```
$ ssh -i mysshkey.ceci mylogin@my-ceci-cluster.example.com
>> Enter passphrase for key 'mysshkey.ceci': mypwd
```

Quick start	https://support.ceci-hpc.be/doc/_contents/QuickStart

CECI clusters : Connection

Demo : NIC4

mar.grignard@uliege.be : GIGA Doctoral School 2019 “Clusters and parallel computing”

CECI clusters : Access (GUI)

> I use Windows, how do I connect ?

MobaXterm

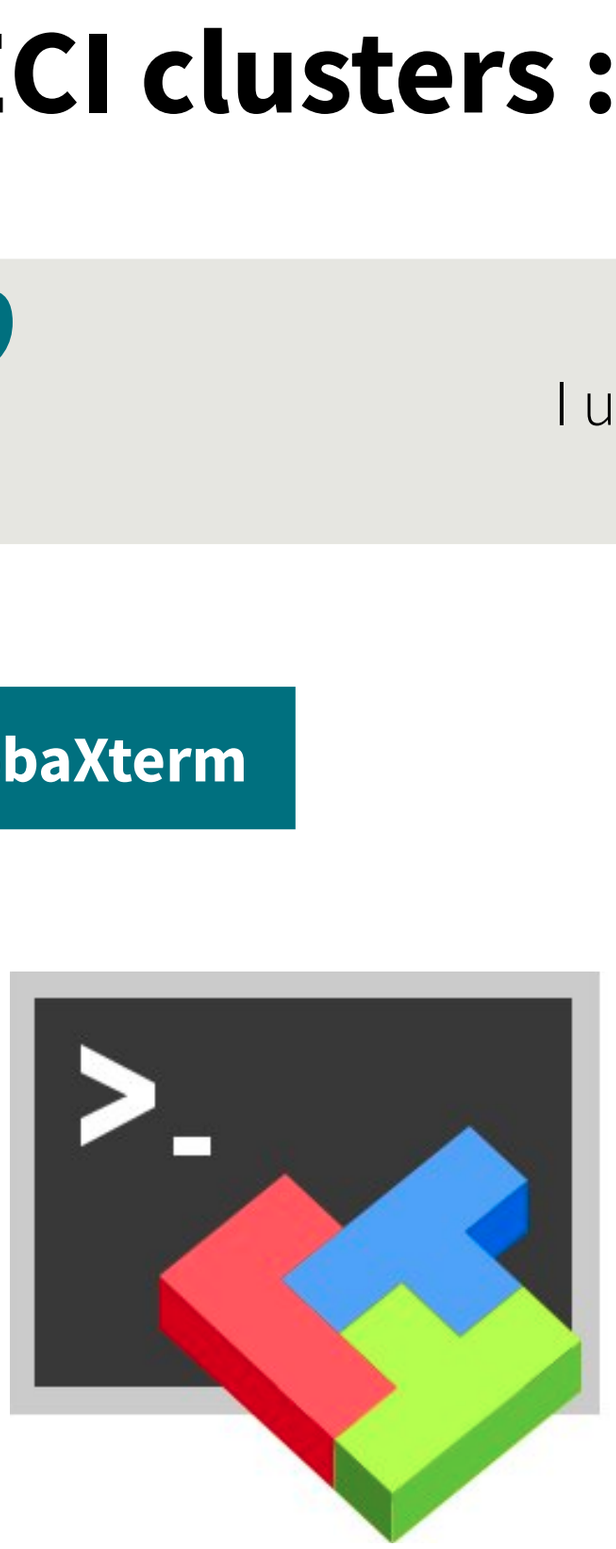

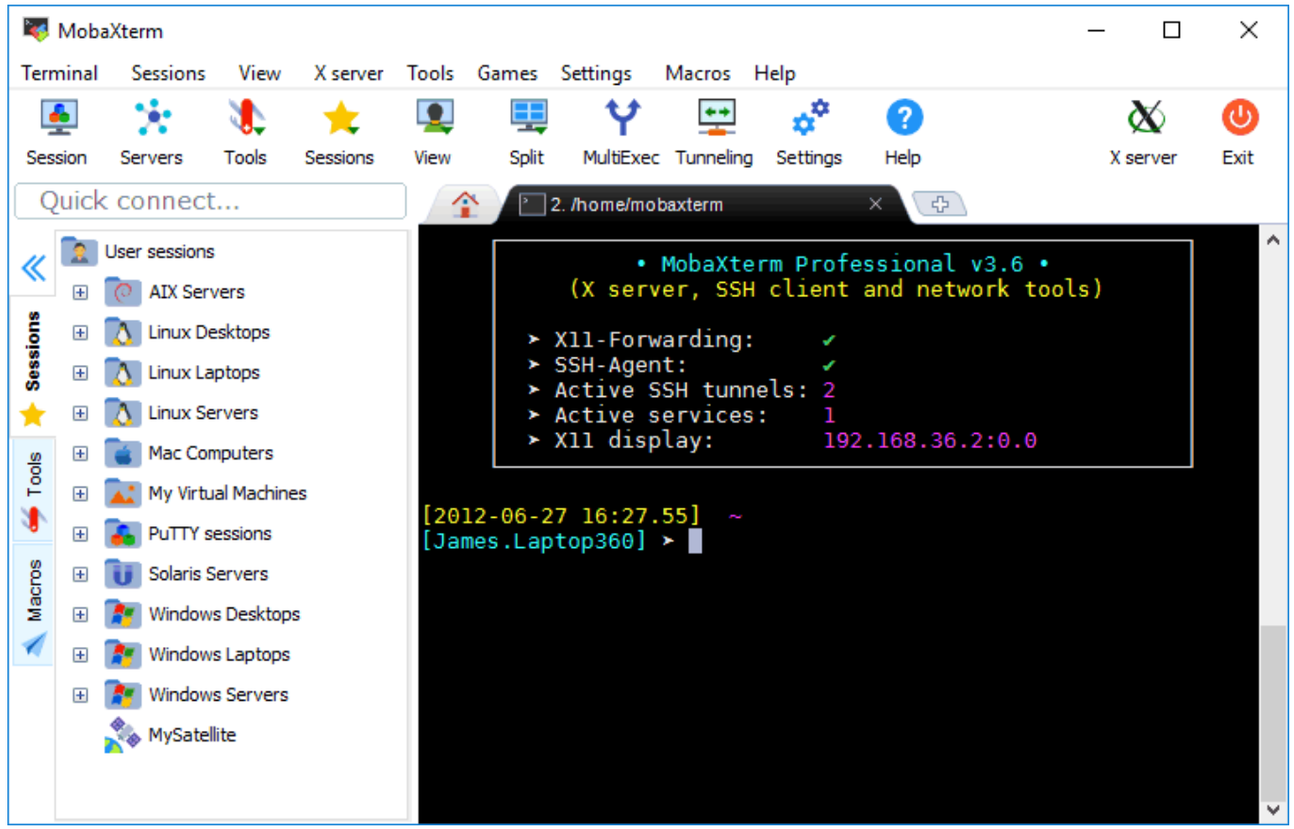

mar.grignard@uliege.be : GIGA Doctoral School 2019 “Clusters and parallel computing”

CECI clusters : Access (GUI)

> Can I transfer my files using a GUI
> even on Linux?

FileZilla

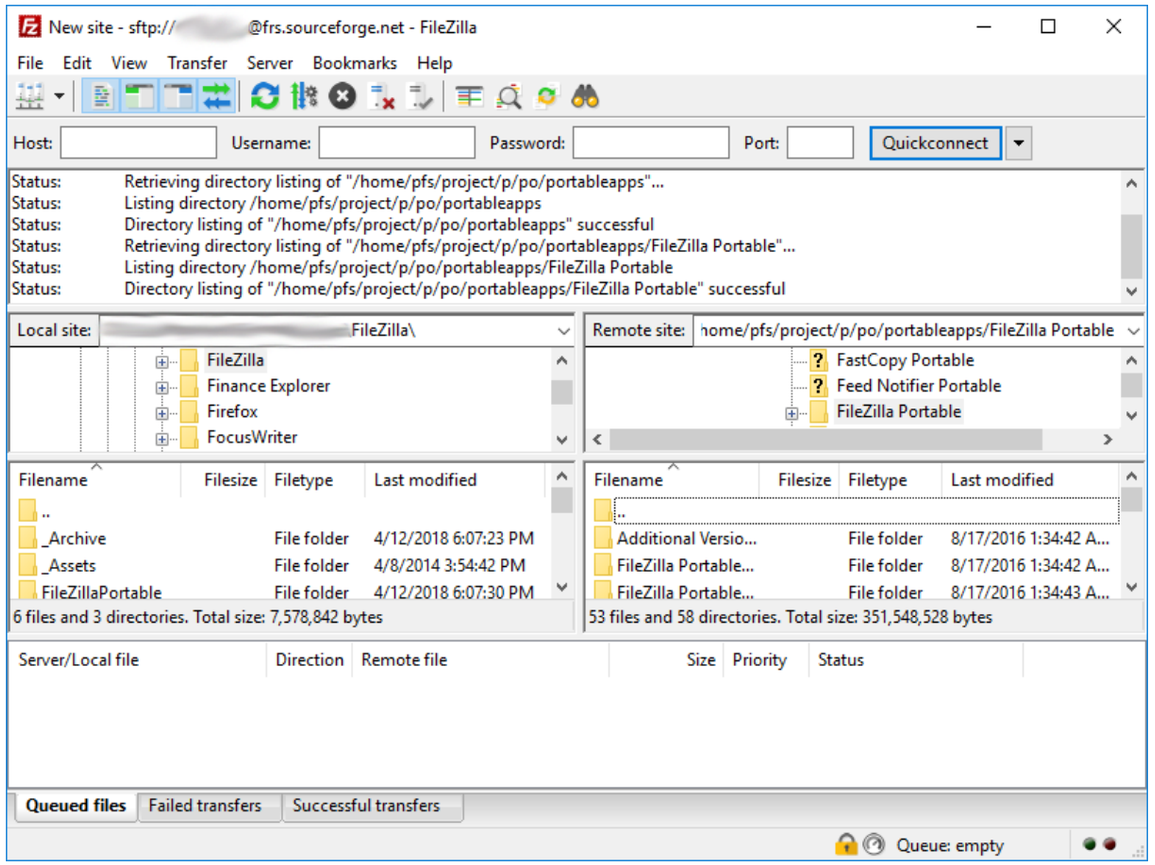

mar.grignard@uliege.be : GIGA Doctoral School 2019 “Clusters and parallel computing”

CECI clusters : File systems

Disk space	Scope	Environment variables (depends on cluster)	
Home	cluster	$HOME	Local
Workdir	cluster	$GLOBALSCRATCH	
Local scratch	node	$TMPDIR,$LOCALSCRATCH *	
Global Home	CÉCI	$CECIHOME	Global
Transfer	CÉCI	$CECITRSF	
Long-term	external		

* Specific to each node.

Local Specific to one cluster, high bandwith from node to disk.

Global Shared between all CECI clusters, duplicate latency, not for storage.

SLURM

Monitor resources

How many CPUs/How much RAM is used ? For how long ?

Allocate resources

Number of CPUs, RAM, time.

```
$ sbatch submitmyjob.sh
```

Manage jobs

List jobs states, report info about running/finished jobs.

```
$ squeue -u $USER
$ sstat -j myjob -o myfmt
$ sacct -j myjob -o myfmt
```

mar.grignard@uliege.be : GIGA Doctoral School 2019 “Clusters and parallel computing”

SLURM : Job submission

myjob.sh

Resources requested
(http://www.ceci-hpc.be/scriptgen)

```
#!/bin/bash
# Submission script for NIC4
#SBATCH --time=01:00:00
#
#SBATCH --ntasks=4
#SBATCH --cpus-per-task=8
#SBATCH -mem-per-cpu=500
#SBATCH -partition=defq
#
#SBATCH --mail-user=my@email.com
#SBATCH --mail-type=ALL
```

Instructions
(Shell script, Python...)

```
# Do some stuff
echo Job start !
foo
bar
echo Job done !
```

Submit

```
$ sbatch myjob.sh
```

mar.grignard@uliege.be : GIGA Doctoral School 2019 “Clusters and parallel computing”

SLURM : Job submission

Demo : helloworld.sh

mar.grignard@uliege.be : GIGA Doctoral School 2019 “Clusters and parallel computing”

SLURM : Job submission

Demo : preproc.sh

mar.grignard@uliege.be : GIGA Doctoral School 2019 “Clusters and parallel computing”

Workflow tips

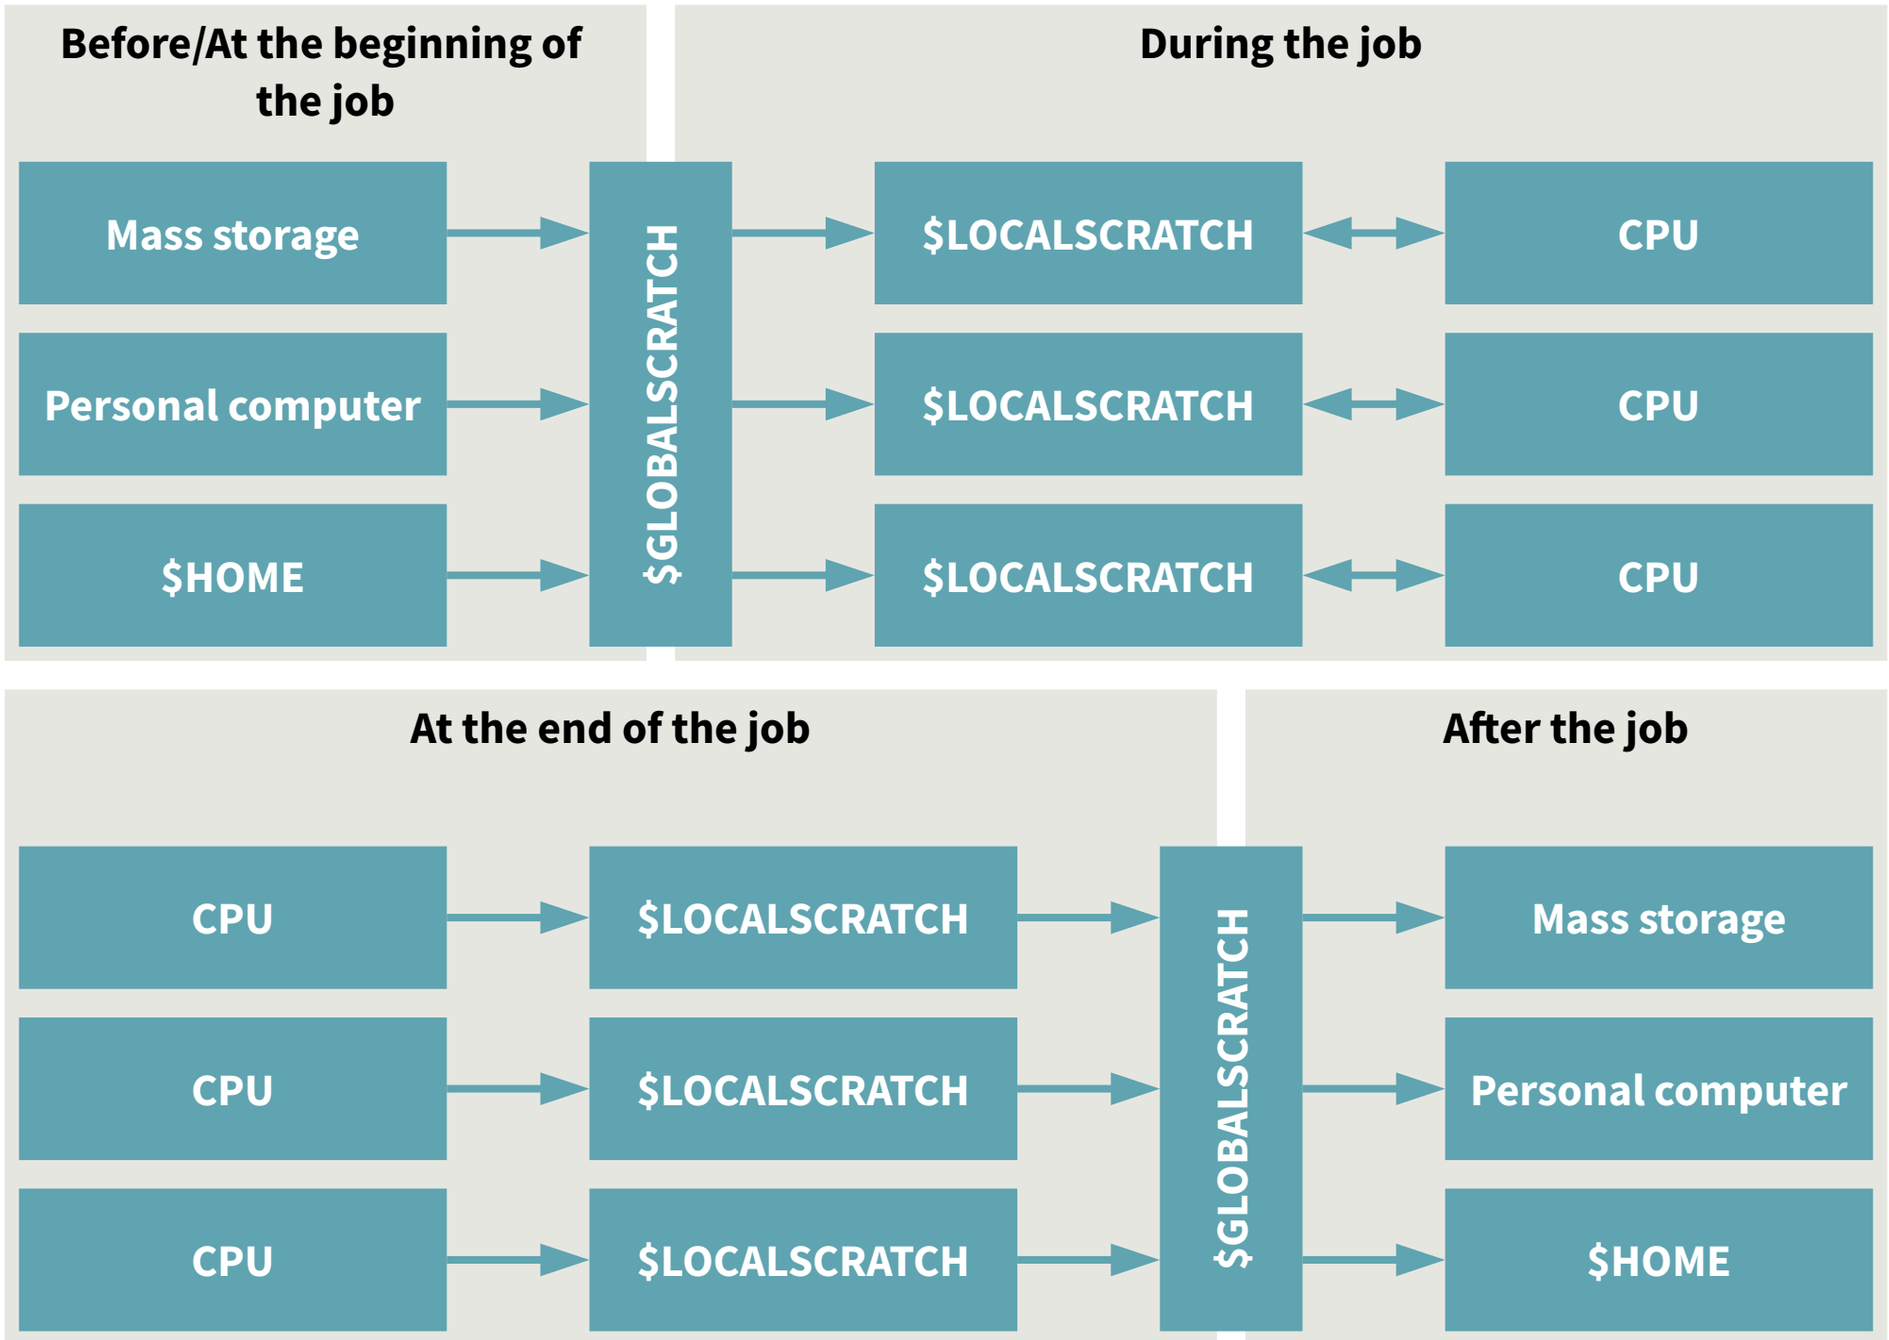

Workflow tips

Demo : workflow.sh

mar.grignard@uliege.be : GIGA Doctoral School 2019 “Clusters and parallel computing”

High-performance computing (HPC)

> This course must be taken
> with a grain of salt.

Disclaimer

User point of view, focused on many-to-many,
CECI clusters, SLURM

But

SLURM is broadly used (GIGA cluster)

---

mar.grignard@uliege.be : GIGA Doctoral School 2019 “Clusters and parallel computing”

Afterword

Training

Learn how to use the clusters might be time consuming
but remember **« 30 to 1 day »**. (http://www.ceci-hpc.be/training)

Fundings

Annual reports, public fundings…
Use them, otherwise they might disappear.

mar.grignard@uliege.be : GIGA Doctoral School 2019 “Clusters and parallel computing”

LIÈGE université
GIGA

> Using clusters is not so difficult, after all !

- Any PhD student who followed this course

---

mar.grignard@uliege.be : GIGA Doctoral School 2019 “Clusters and parallel computing”

Matlab on clusters

> I just learned to use Matlab this week,
> can I run my scripts on the cluster ?

MCC	https://nl.mathworks.com/
MCR	https://nl.mathworks.com/
Octave	https://www.gnu.org/

mar.grignard@uliege.be : GIGA Doctoral School 2019 “Clusters and parallel computing”