HPC cluster Initiation

i-Trop Presentation

DIRECTION

Yves Vigouroux

Jean Benoit Morel/ Gilles Béna

Eric Delaporte/ Maryline Bonnet

Frédéric Simard

Julie ORJUELA-BOUNIOL, IE
Bioinformaticienne

Ndomassi TANDO, IR
Ingénieur systèmes
Animateur plateau, RMQ

Christine TRANCHANT-DUBREUIL, IR
Bioinformaticienne

Aurore COMTE, IE
Bioinformaticienne

Alexis DEREEPER, IE
Bioinformaticien

Bruno GRANOUILLAC, IE
Systèmes d'information

Jacques Dainat, IR
Bioinformaticien

Nicolas Fernandez, IR
Bioinformaticien

i-Trop Presentation

Requests/incidents/Howtos

- Requests form
  https://bioinfo.ird.fr/index.php/cluster-fr/
  - Account
  - Software
  - Projects

- Incidents: contact bioinfo@ird.fr

- Howtos:
  https://bioinfo.ird.fr/index.php/tutorials-fr/howtos-for-hpc-cluster-itrop/

- Slurm Tutorials:
  https://bioinfo.ird.fr/index.php/tutorials-fr/slurm/
- FAQ:
  https://bioinfo.ird.fr/index.php/faq-fr/

ARCHITECTURE

A cluster?

- A logical unit composed of several servers
- A powerful unique machine
- Allow to obtain high computing performance
- A bigger capacity storage
- More reliable
- A better ressources availabilty

A cluster?

- A logical unit composed of several servers
- A powerful unique machine
- Allow to obtain high computing performance
- A bigger capacity storage
- More reliable
- A better ressources availabilty

Cluster components

COMPUTING

- **Master Node**
  Handle ressources and jobs priorities

- **Computing nodes**
  Resources (CPU or RAM memory)

Cluster components

COMPUTING

- **Master Node**
  Handle ressources and jobs priorities
- **Computing nodes**
  Resources (CPU or RAM memory)

STORAGE

- **San**
  Storage

Architecture: components role

- **1 Master node**

bioinfo-master1.ird.fr

Role :

- Launch and prioritize jobs on computing nodes
- Accessible from the Internet
- Connection :

ssh login@bioinfo-master1.ird.fr

Architecture: components role

- **1 Master node**

bioinfo-master1.ird.fr

Role :

- Launch and prioritize jobs on computing nodes
- Accessible from the Internet
- Connection :

ssh login@bioinfo-master1.ird.fr

- **32 nodes**

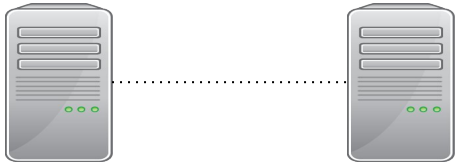

nodeX
X : 0..31

Role :

- Used by the master to execute jobs
- Not accessible from the Internet
- node0 to node31

Practice

Step 1: Connection, srun

1

Go to the Practice 1 *of the github*

Analyses step of the cluster

Connection to bioinfo-master1.ird.fr and resources reservation

Step 1
srun ou sbatch

Partitions

Partitions	Use	Nodes RAM features	Nodes Core features
short	Short jobs < 1 jour	48 to 64 GB	12 cores
normal	Short jobs max 14 days	64 to 512 GB	12 to 112 cores
long	45 days > long jobs > 14 days	48 GB	12 to 24 cores
highmem	Jobs with memory needs	512GB	112 cores
supermem	Jobs with big memory needs	1TB	40 cores
gpu	Need for analyses on GPU cores	192GB	24 cpu et 8 GPU cores

Particular case : gpu partition

- Partition to work on GPUs processors : basecalling, MiniOn etc..
- Restricted access to gpu_account group
- Request access with arguments to do here:

  https://itrop-glpi.ird.fr/plugins/formcreator/front/formlist.php

How do I choose the partition?

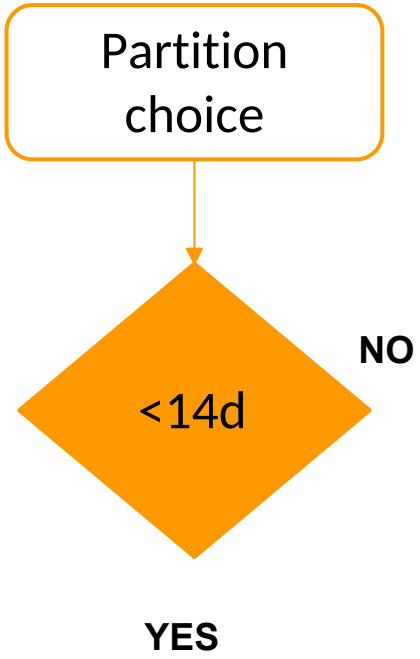

How do I choose the partition?

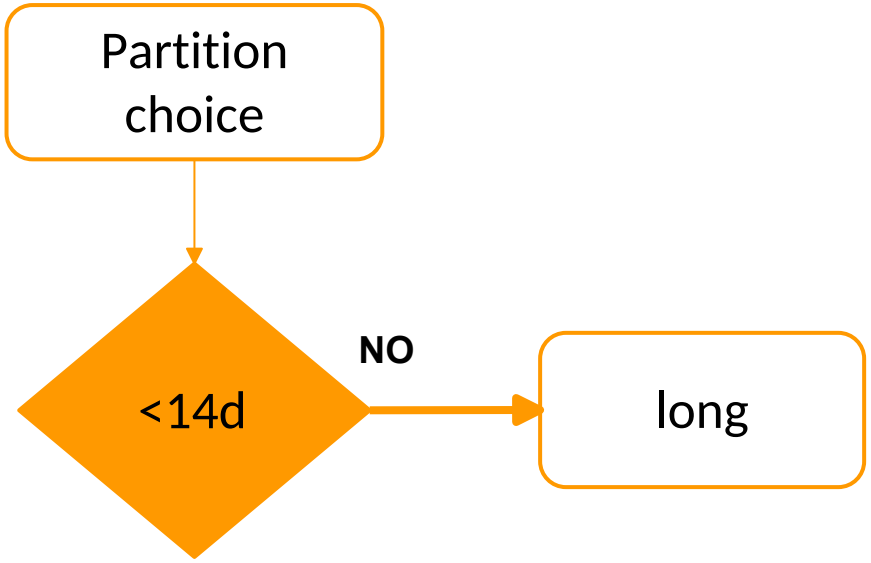

How do I choose the partition?

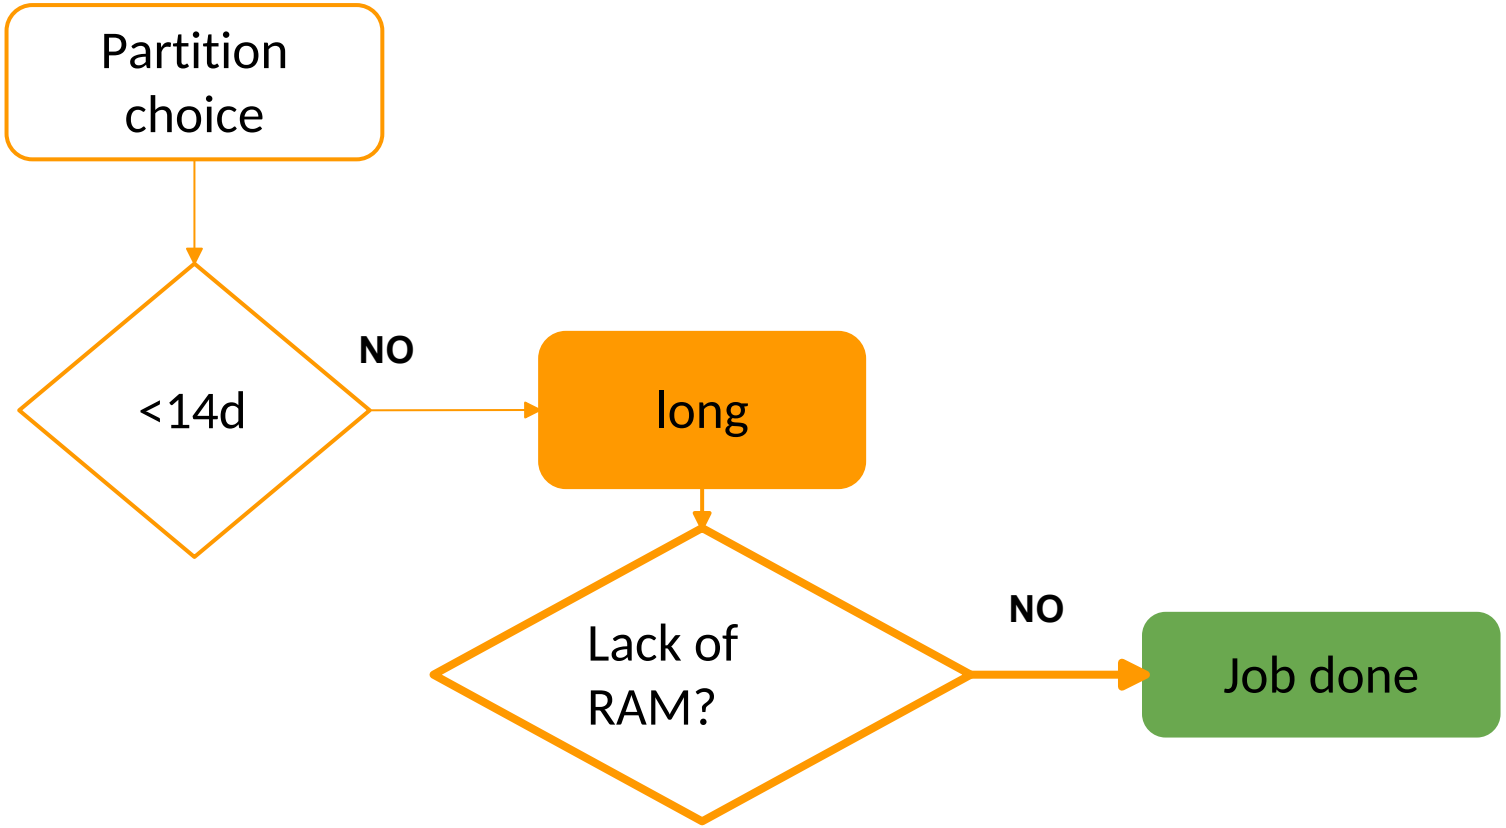

How do I choose the partition?

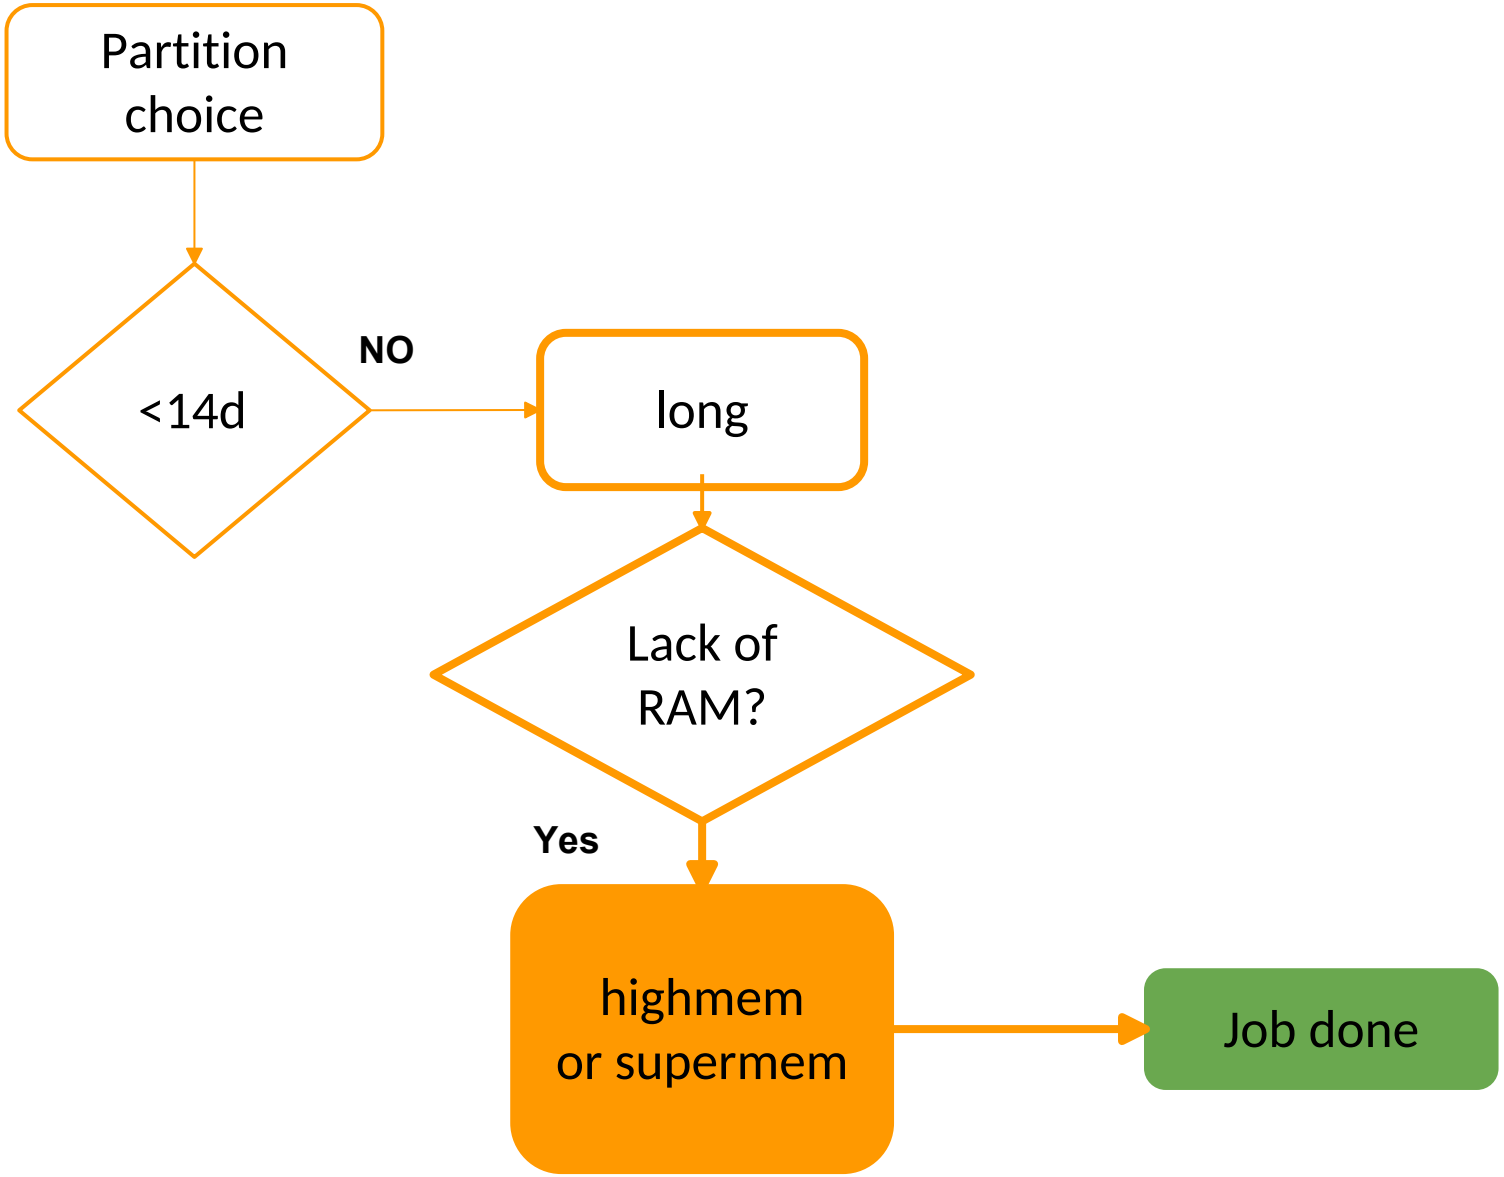

How do I choose the partition?

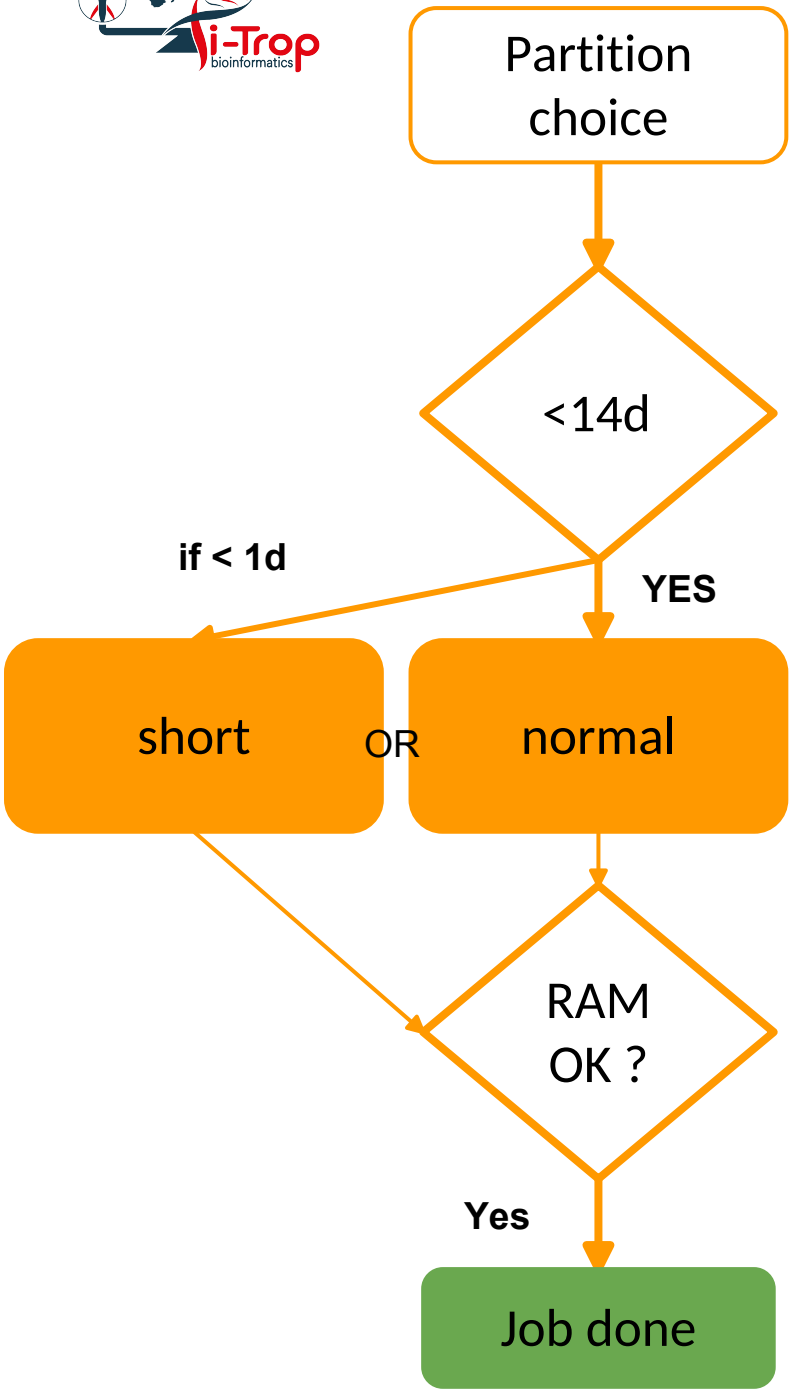

How do I choose the partition?

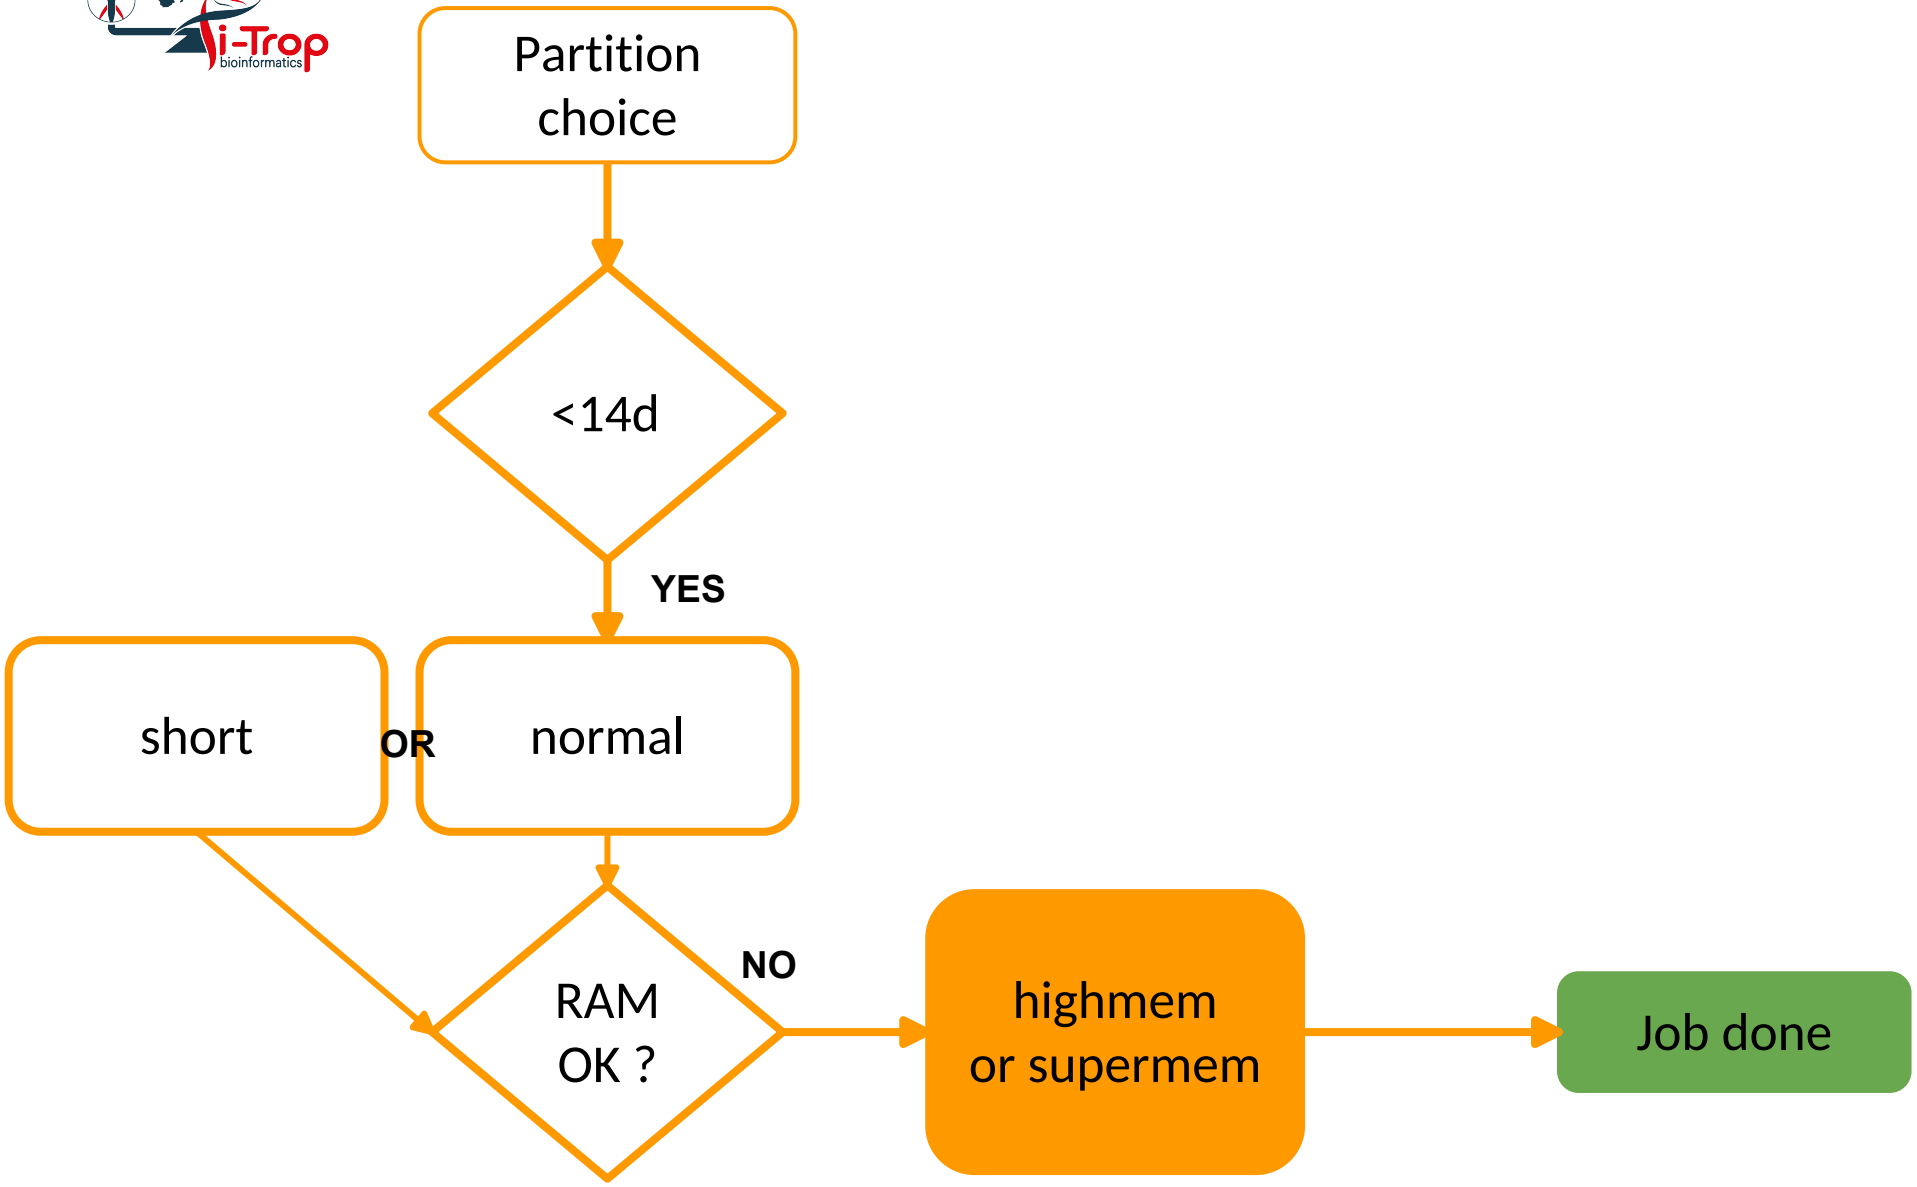

How do I choose the partition?

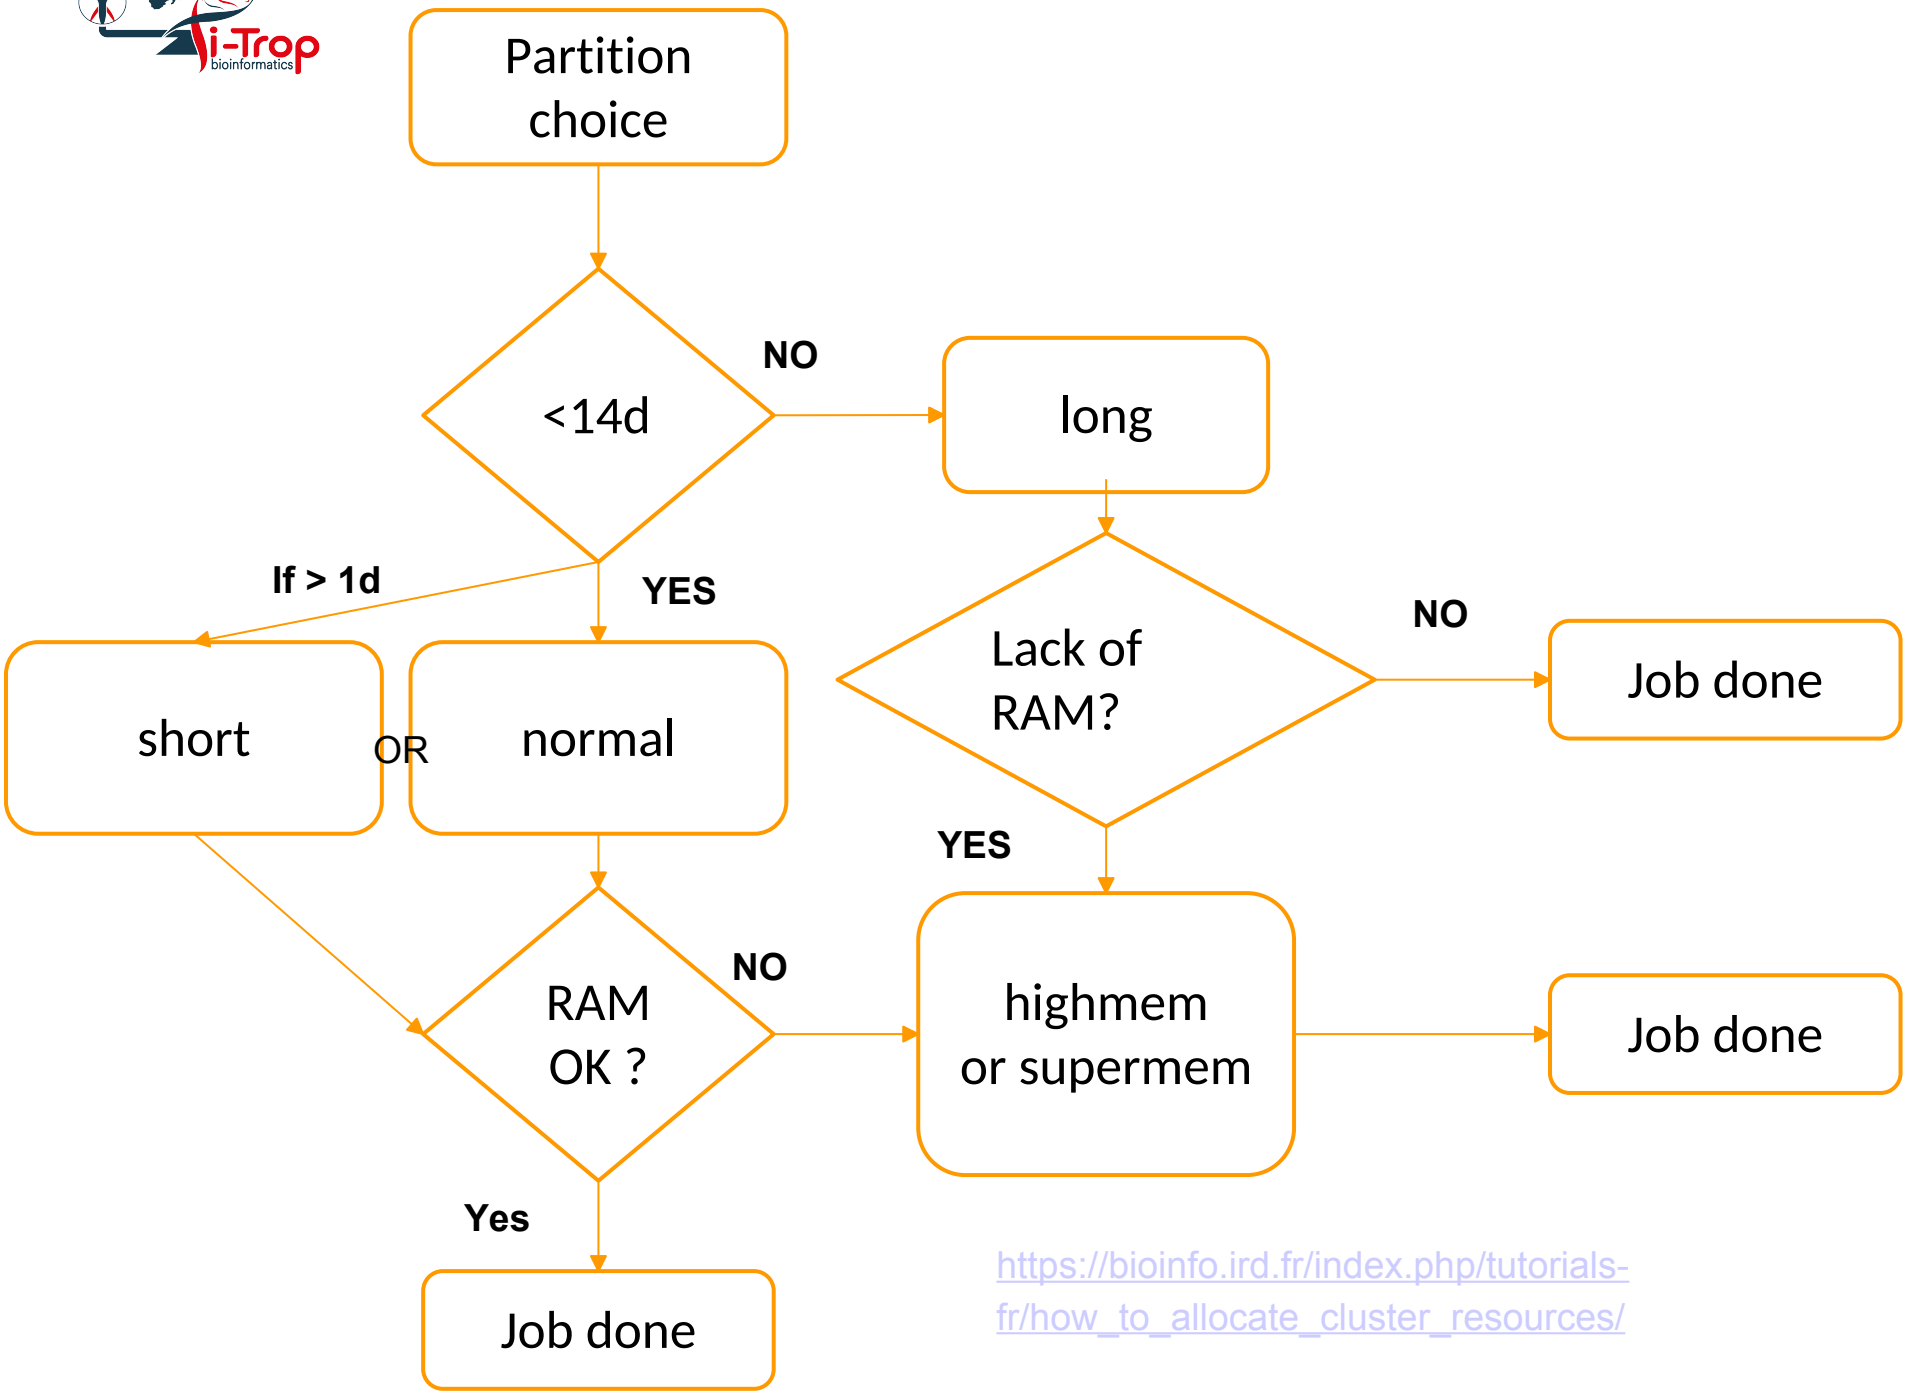

https://bioinfo.ird.fr/index.php/tutorials-fr/how_to_allocate_cluster_resources/

How do I choose a partition?

Rules	Partition	Tools example	comments
basecalling, demultiplexing, correction	***gpu***	medaka, guppy, machine learning tools	Restricted access
assembling >100G RAM	***supermem***	miniasm, flye, raven, smartdenovo	Target genome> 400 Mb (Rice genome doesn't need 100GB)
genomicsbd (gatk) > 100G RAM	***supermem***	GATK genomicsDB	Target genome of more than 400 Mb (>10 samples)
assemblings => 35G et < 100G RAM	***highmem***	miniasm, flye, raven, smartdenovo	Target genome between 100 and 400 Mb
Pops structure	***long***		
simulations	***long***		
metagenomic	***normal***	quiime2, frogs	
mapping	***normal***	bwa, minimap2, hisat2	Need a lot of cores not too many RAM **Tool number of cores = number of cores to reserve**
genotypage	***normal***	GATK haplotypecaller, samtools mpileup, bcftools	Need a lot of cores not too many RAM **Tool number of cores = number of cores to reserve**
stats	***normal***	R	
scripts test	***short***	bash, python, R	

Architecture: elements role

- 1 Master node

bioinfo-master1.ird.fr

Role :
- Launch and prioritize jobs on compting nodes
- Accessible from the Internet

- 32 Nodes

nodeX
X : 0..31

Rôle :
- Used by the master to execute jobs
- Not accessible from the Internet

- 1 san

bioinfo-san.ird.fr (san)

Role :
- Store users data
- Accessible from the Internet
- To transfer data : *via filezilla or scp*

Storage on i-Trop cluster

san server

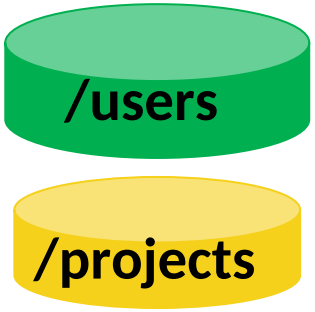

node servers

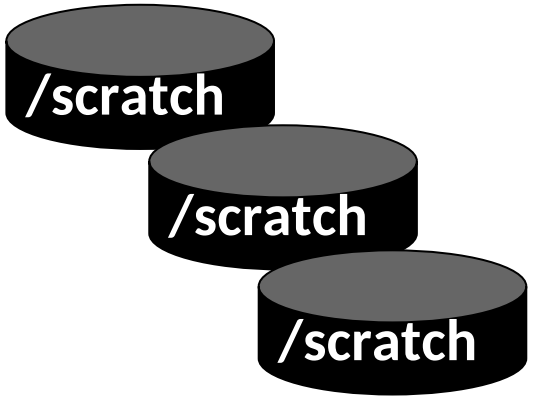

Storage on i-Trop cluster

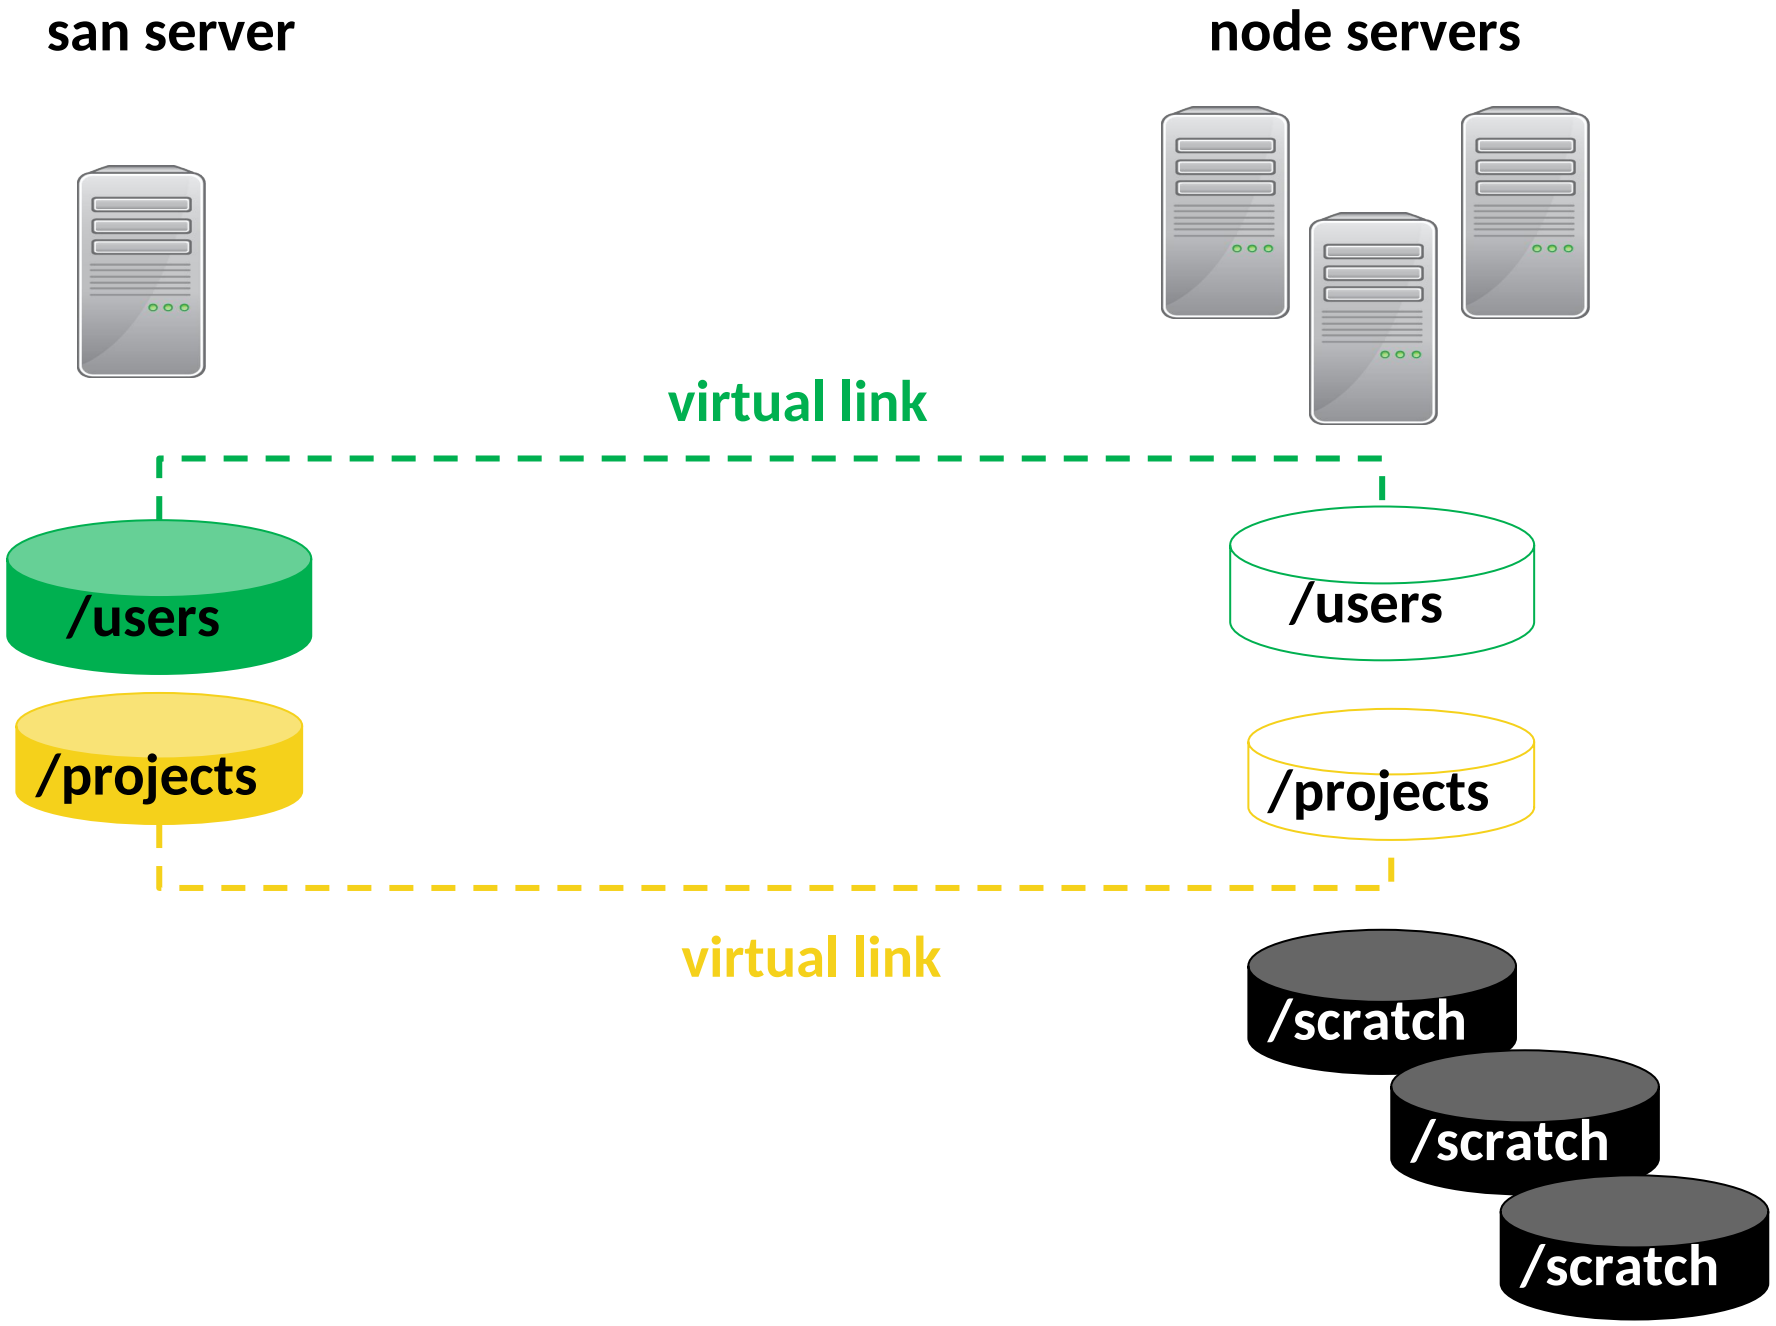

Storage on i-Trop cluster

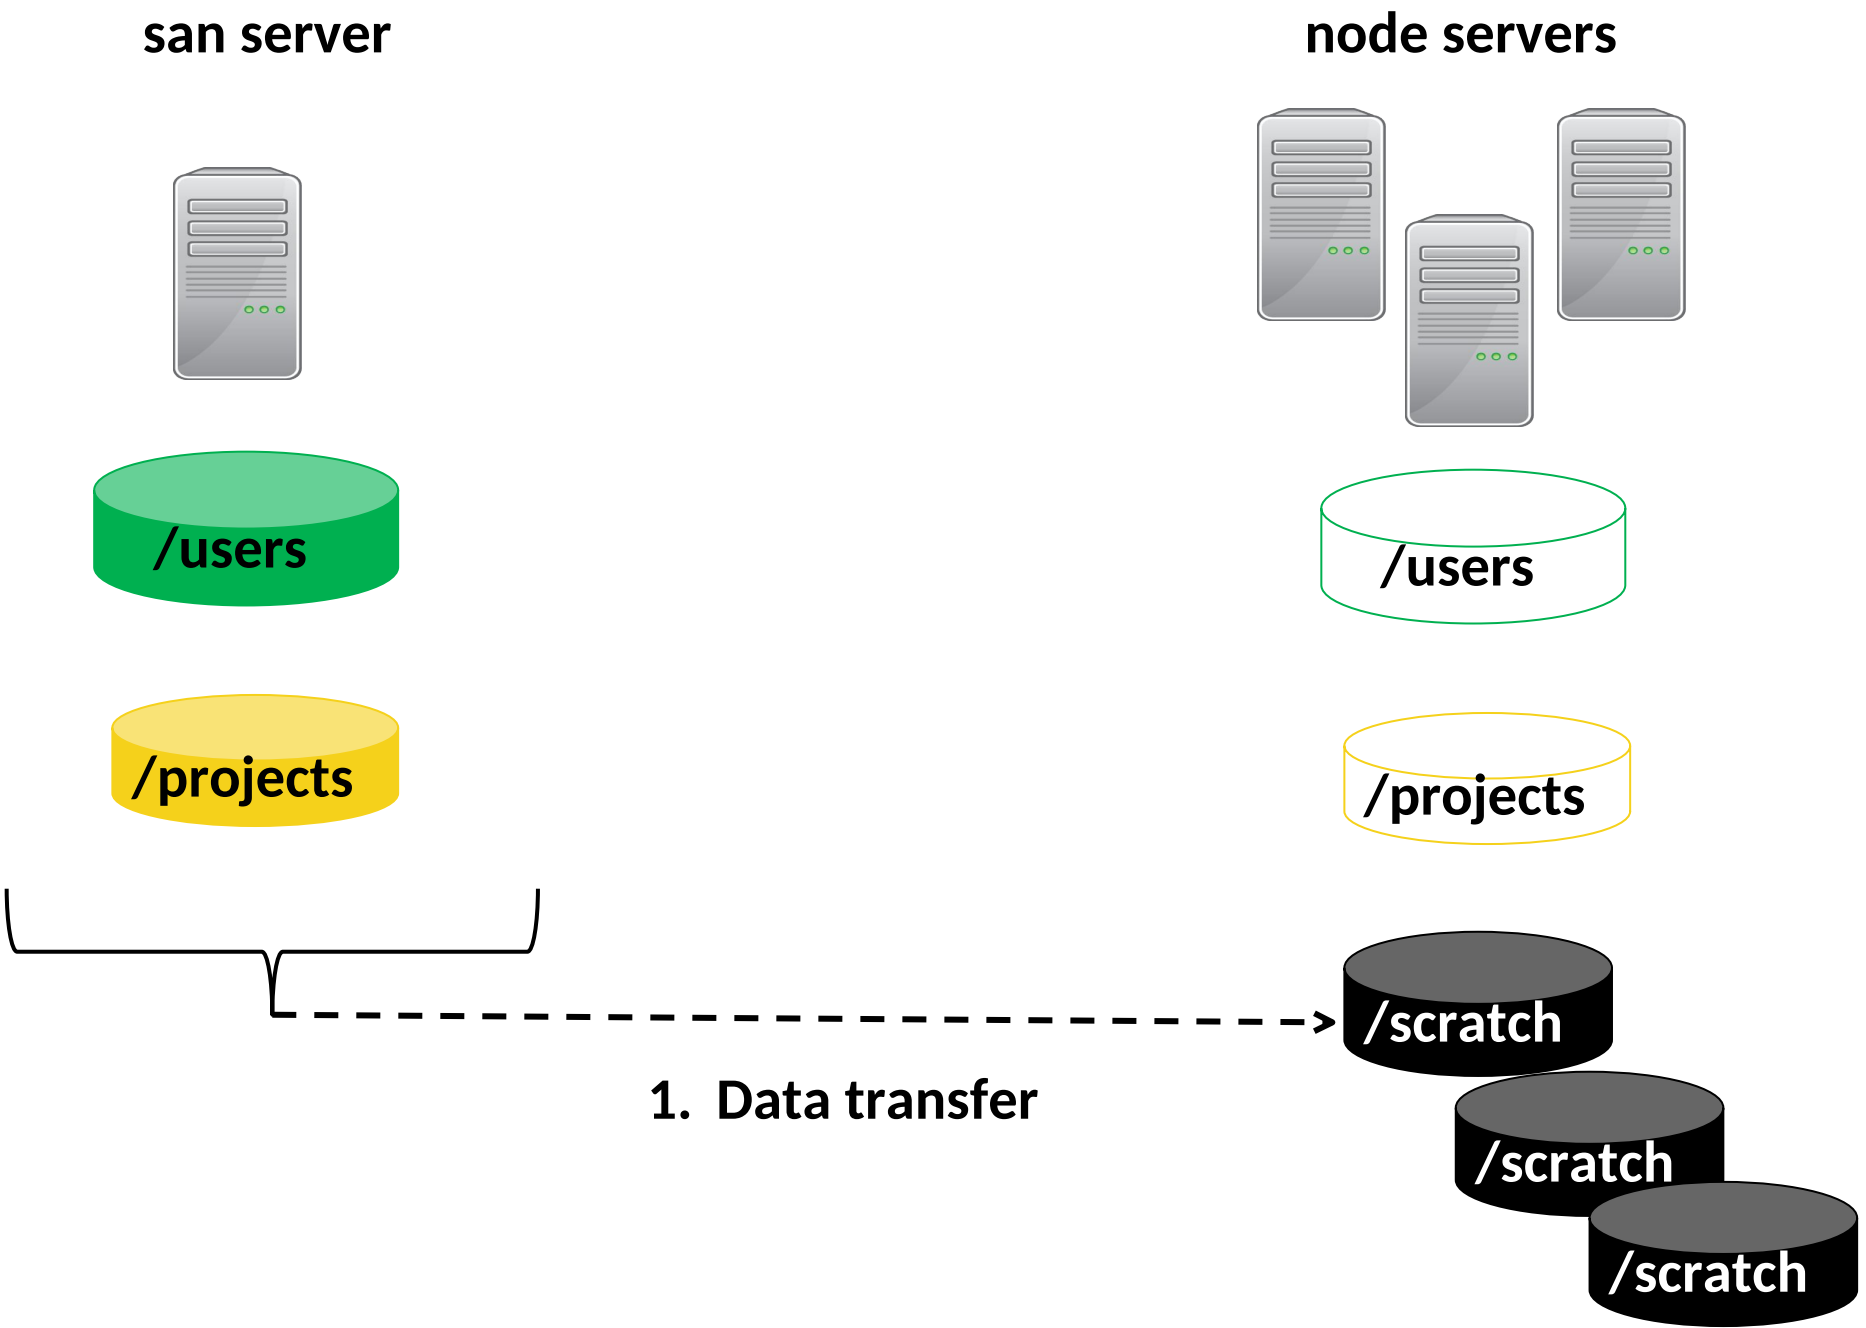

Storage on i-Trop cluster

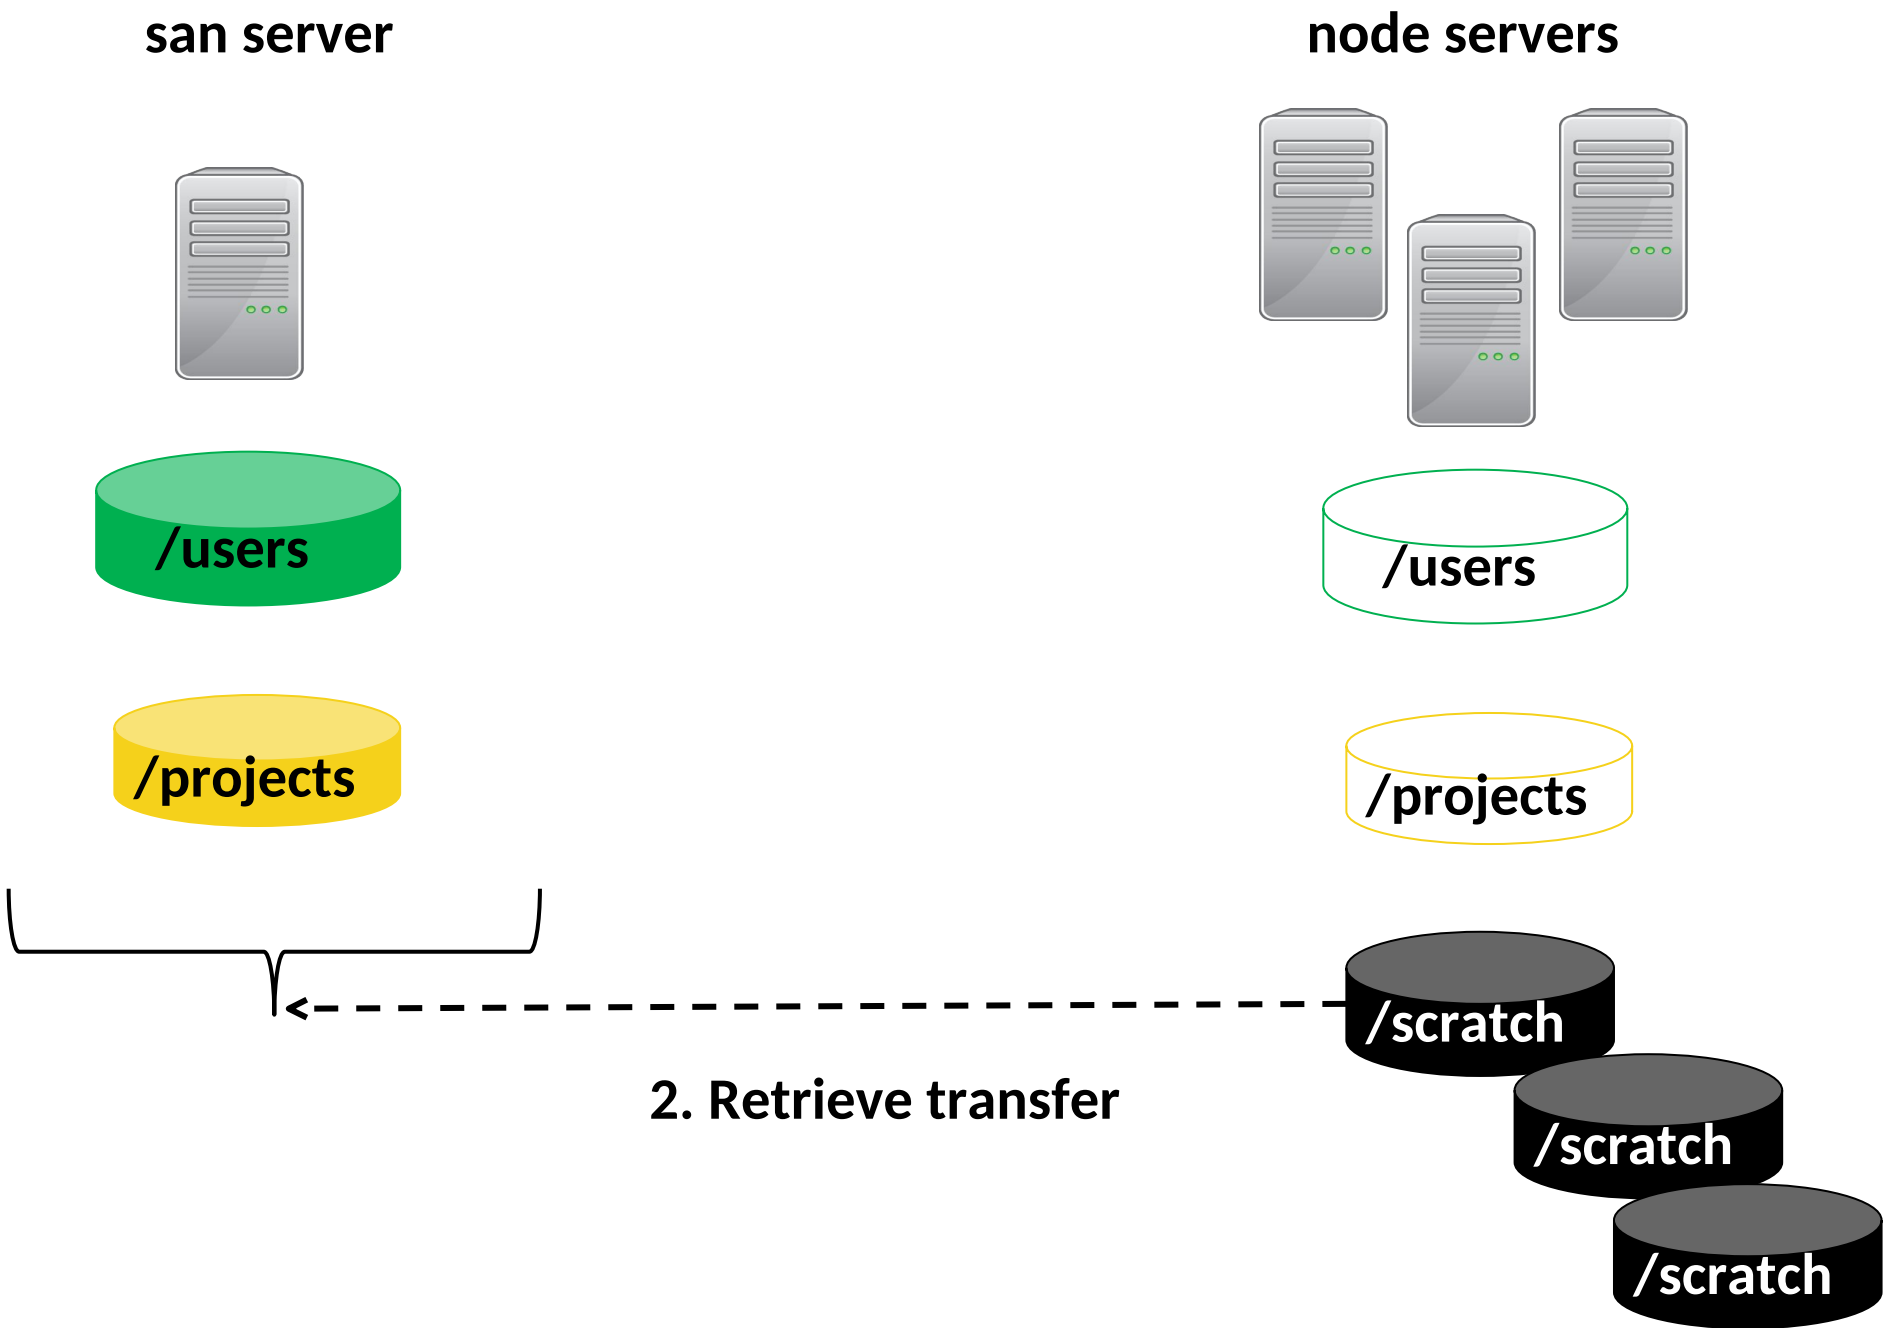

Analyses steps ot the cluster

Connection to bioinfo-master1.ird.fr and resources reservation

Step 1

Creation of the analyses folder in the /scratch of the node

Step 2
mkdir

Practice

Step 2:srun, partition

2

Go to the Practice2 *of the github*

Data transfer on the i-Trop Cluster

Personal computer

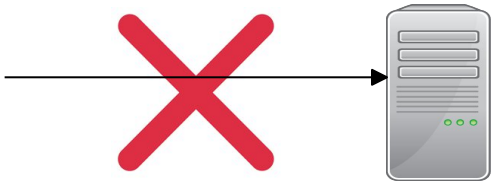

Direct transfer forbidden

bioinfo-master1.ird.fr

Data transfer on i-Trop cluster

/users, /projects/medium/, /projects/large, /projects/xl, projects/xxl ou /share

bioinfo-san.ird.fr

Hostname : bioinfo-san.ird.fr

Login : cluster account

ssh keys
Port : **22**

Analyses steps of the cluster

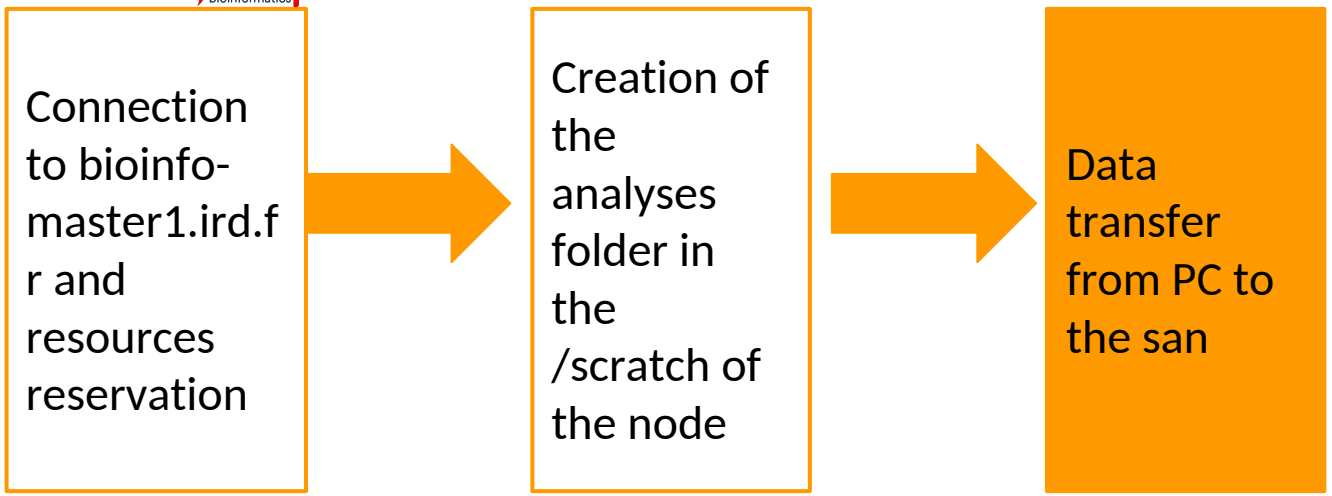

Step 1

Step 2

Step 3
filezilla

Copy your data from your PC/MAC to the san if they are not on the cluster

Practice

Step 3: filezilla

3

Go to the Practice3 *of the github*

Copy with scp

- Copy between 2 remote servers :

  scp -r source destination

- Syntax if the source is remote:

  scp -r server_name:/path/file_to_copy local_folder

- Syntaxe si la destination est distante :

  scp -r /path/file_to_copy nameserver:/path/remote_server

Ex: scp -r san:/home/tando/data/folder /scratch/tando/

Analyses step on the cluster

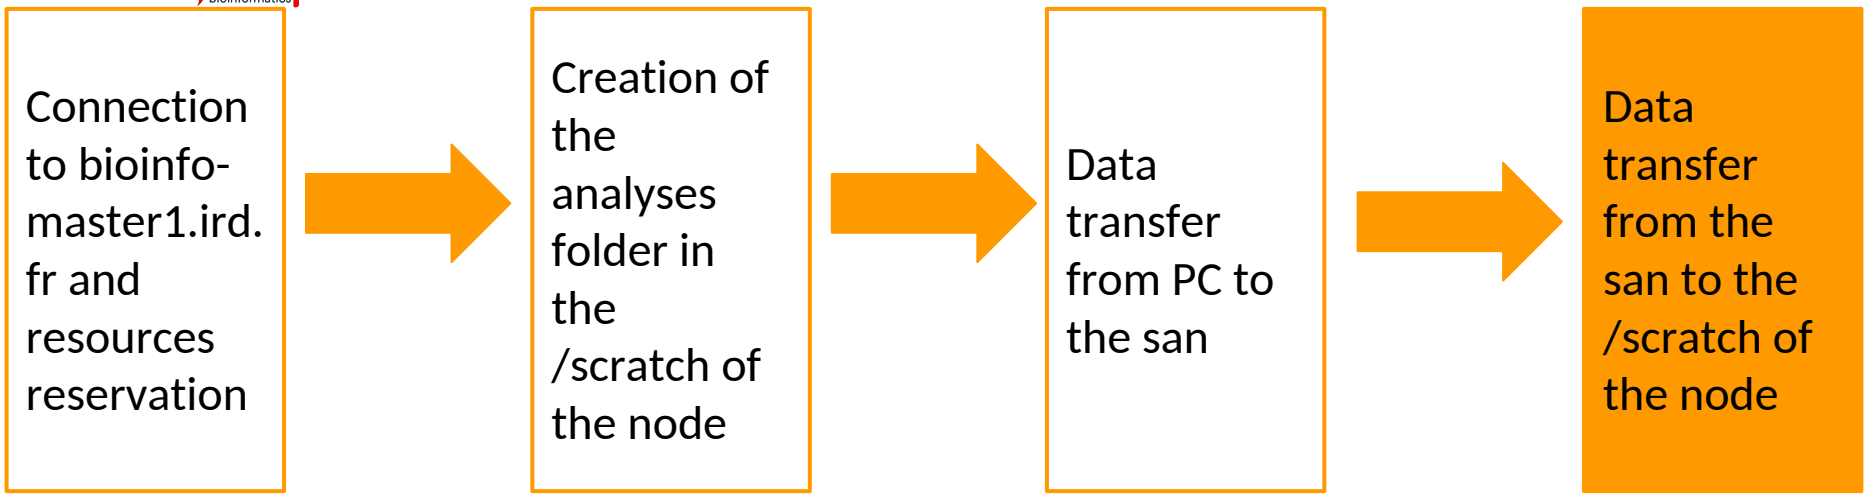

Step 1

Step 2

Step 3

Step 4

scp

Practice

Step 4: scp to nodes

4

Go to the Practice4 *of the github*

Module Environment

- Allow to choose the version of software you want to use
- 2 types of softwares :
  bioinfo : includes all the bioinformatics softwares (example BEAST)
  system : includes all the system softwares(example JAVA)
- Overcome the environment variables

Module Environment

- See the available modules :

module avail

- Obtain infos on a particular module:

module whatis + module name

- Load a module :

module load + modulename

- List the loaded module :

module list

- Unload a module :

module unload + modulename

- Unload all the modules :

module purge

Analyses steps of the cluster

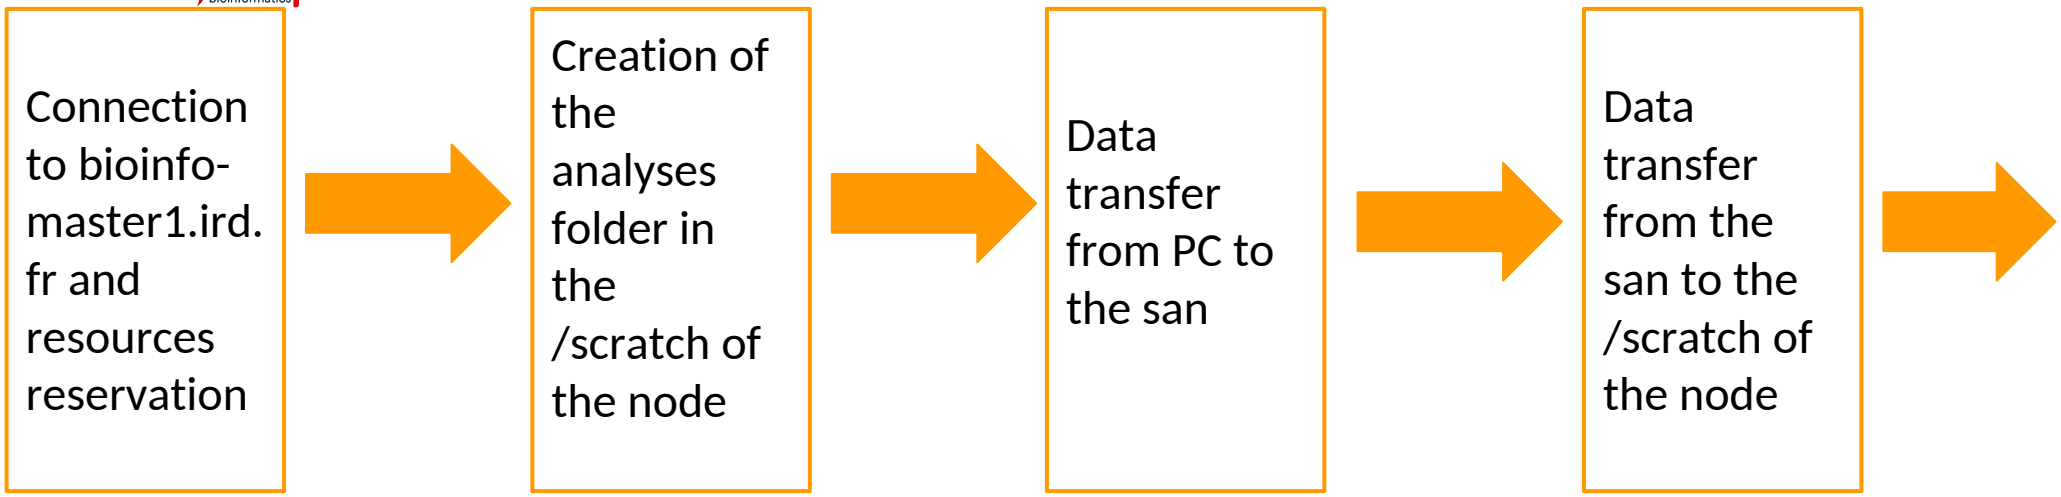

Load the software with module environment

Step 5
module

Practice

Step 5: module environment

5

Go to the Practice5 of the github

Analyses steps of the cluster

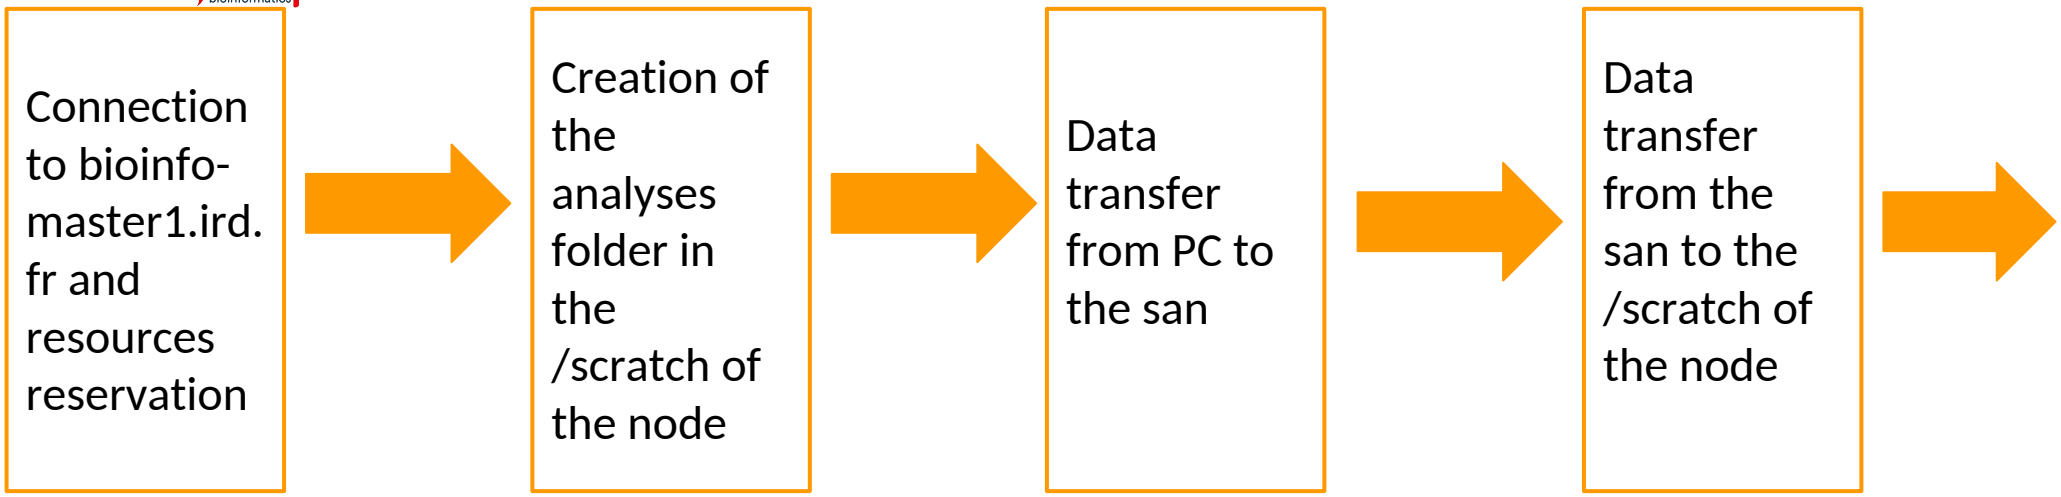

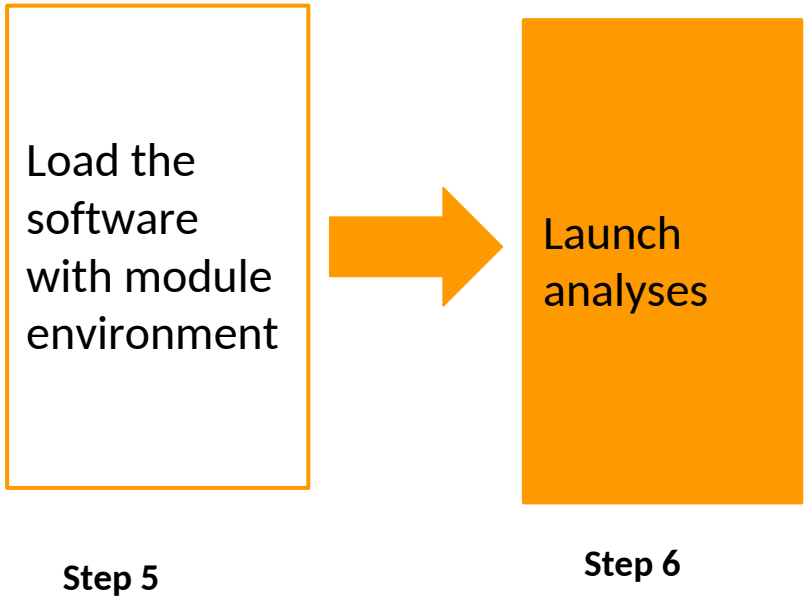

Launch a command from the prompt

- Load the software version to launch
- Launch the data analysis

```
$~ command <options> <arguments>
```

With *command*: the command to launch

Practice

Step 6: launch the analysis

6

Go to the Practice6 of the github

Results transfer to the san

- Copy between 2 remote servers :

  scp -r source destination

- Syntax if the source is remote:

  scp -r server_name:/path/file_to_copy local_folder

- Syntax if the destination is remote :

  scp -r /path/file_to_copy nameserver:/path/remote_server

Analyses steps of the cluster

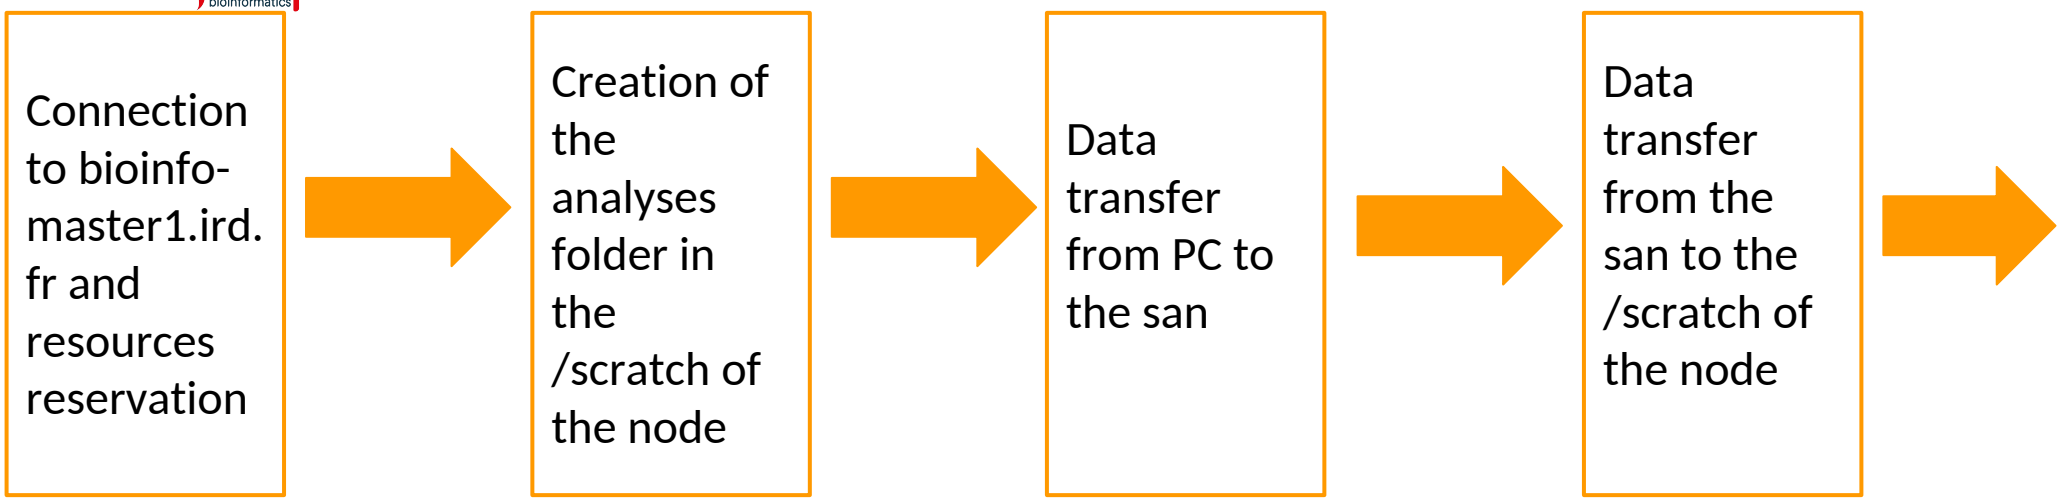

Step 1

Step 2

Step 3

Step 4

Step 5

Step 6

Step 7

scp

Practice

Step 7: Retrieve results

7

Go to the Practice7 of the github

Delete results from scratchs

- Scratch = temporary spaces
- Verify that the copy is OK
- Use rm command

```
cd /scratch
rm -rf nom_rep
```

Analyses steps of the cluster

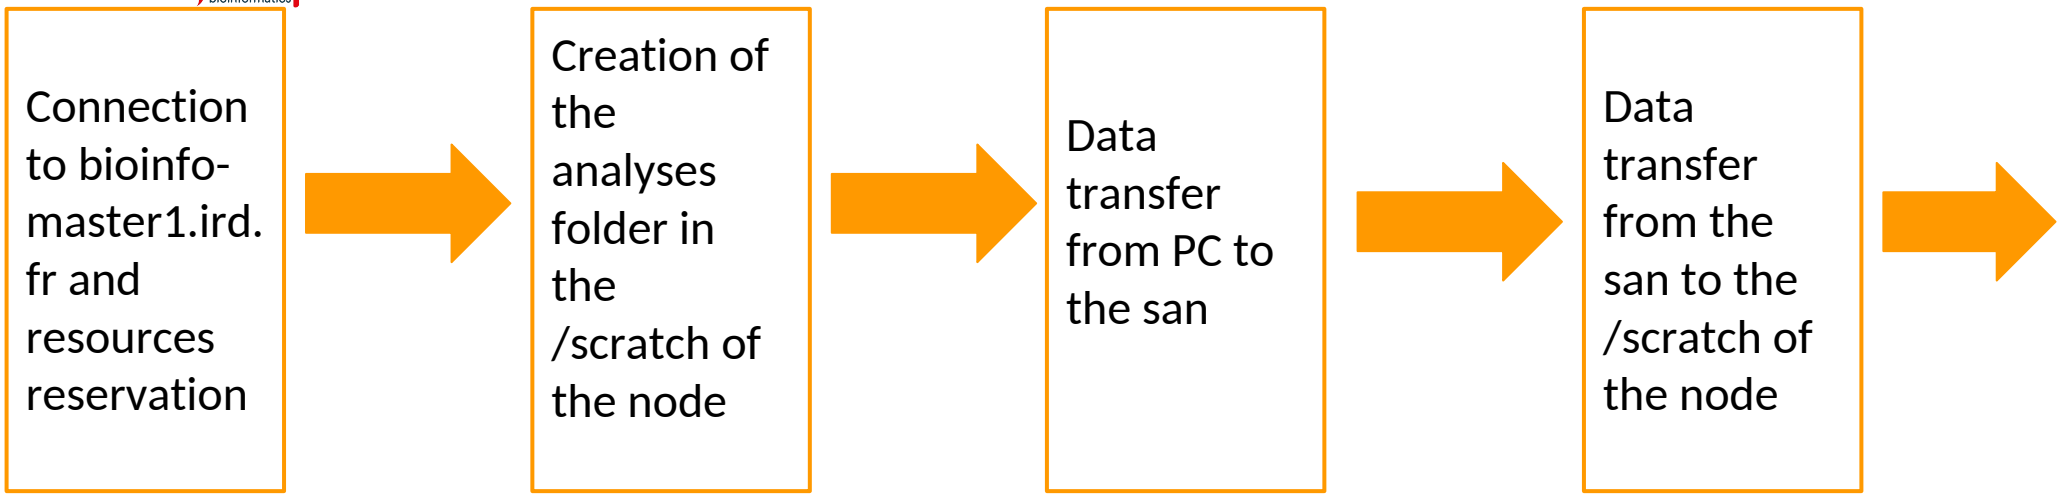

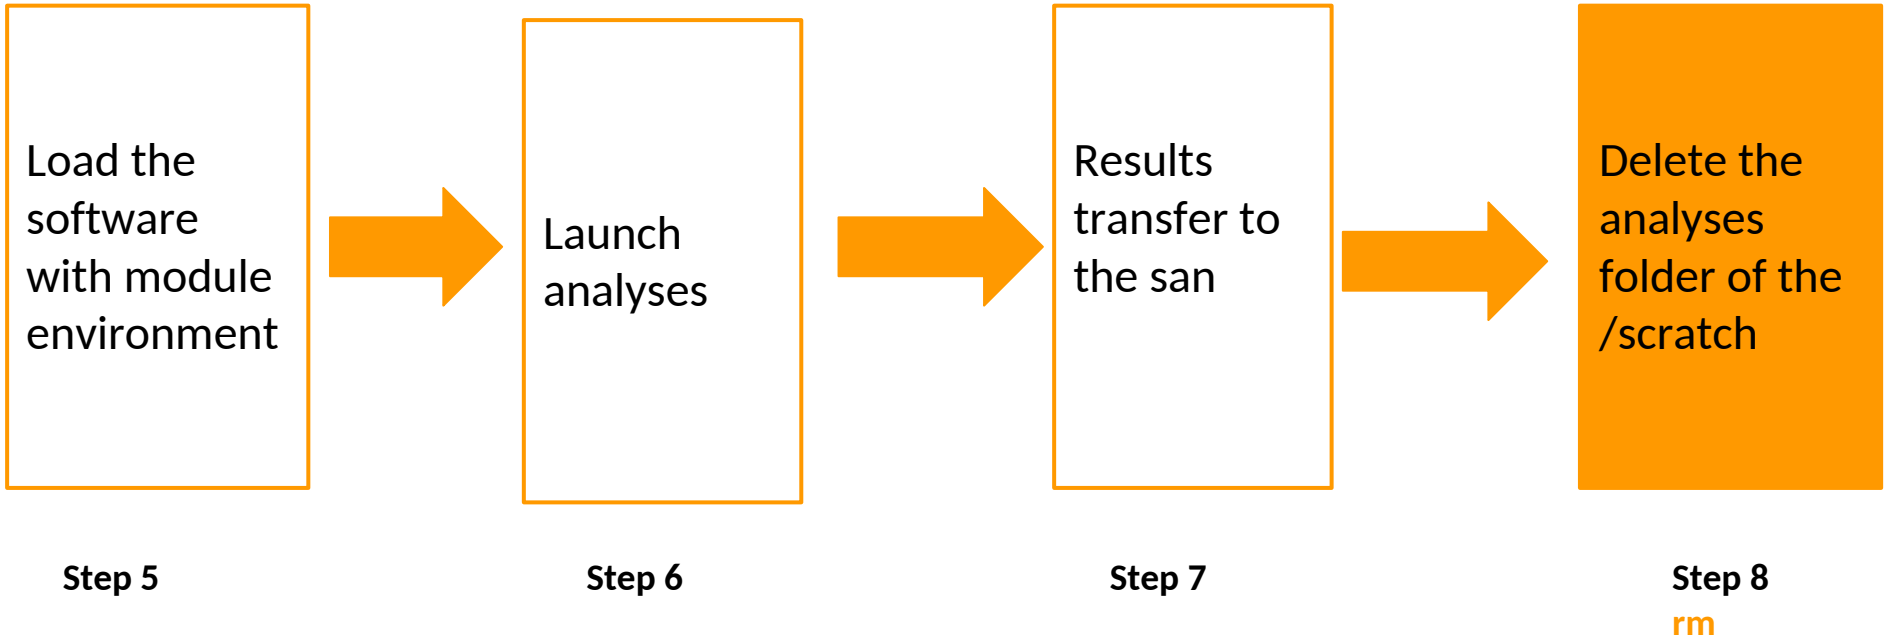

Practice

Step 8: data deletion

8

Go to the Practice8 *of the github*

Scripts to visualize/delete temporary data

- Scripts location: /opt/scripts/scratch-scripts
- Visualize data on scratchs: scratch_use.sh

```
sh /opt/scripts/scratch-scripts/scratch_use.sh
```

- Delete data on scratch: clean_scratch.sh

```
sh /opt/scripts/scratch-scripts/clean_scratch.sh
```

Main Slurm commands

Commande	Description	Example
srun --time=0X:00 --pty bash -i	Interactive way to connect to a node for X minutes	srun --time=02:00:00 --pty bash -i Connection for 2 hours
sbatch	Launch a analyses in background via a script	sbatch script.sh
sinfo	Informations on partitions	sinfo
scancel	Deletion of jobs <job_id>	scancel 1029
squeue	Infos onjobs	squeue -u tando
scontrol show job <job_id>	Infos on active job <job_id>	scontrol show job 1029
sacct -j <job_id>	Infos on finished job <job_id>	sacct -j 1029

More Slum infos here: https://bioinfo.ird.fr/index.php/tutorials-fr/slurm/#part-2

Options of sbatch, srun, salloc

Options	Description	Example
--job-name=<name>	Name the job	sbatch --job-name=tando_blast
-p <partition>	Choose a partition	sbatch -p highmem
--nodelist=<nodeX>	Choose a particular node	sbatch -p normal --nodelist=node14
-n <nbre_taches>	Launch several instance of a command	srun -n 4 hostname
-c <nb_cpu_par_tache>	Allocate the number of cpus per task	srun -n 4 -c 2 hostname
--mail-user=<emailaddress>	Send a email	sbatch --mail-user=ndomassi.tando@ird.fr
--mail-type=<event>	Send a email when: END: end of the job FAIL: abort BEGIN: beginning ALL: all	sbatch ---mail-type=BEGIN

LAUNCH A JOB

Advantages

- Scheduler choose resources automatically
- Use up to 24 cores at the same time
- Possibility to configure this choice
- Jobs launch in background

→ possibility to turn off your PC/MAC

→ automatic results retrieving

Launch a job in batch mode

- Execute a script via Slurm
- Use:

```
$~ sbatch script.sh
```

With script.sh : the name of the script

bash script syntax

First part of the script (in green): sge execution options with the key word #SBATCH

```
#!/bin/bash

###### Configuration SLURM#################################
## On définit le nom du job
#SBATCH --job-name=test
## On définit le nom du fichier de sortie
#SBATCH --output=res.txt
## On définit le nombre de tâches
#SBATCH --ntasks=1
## On définit le temps limite d'éxécution
#SBATCH --time=10:00
##########################################################
```

bash scripts syntax

In the 2nd part of the script: the command to execute

```
###############Partie exécution des commandes ###################

nom_variable1=”valeur_variable1”
nom_variable2=”valeur_variable2”

sleep 30
hostname
```

Practice

Launch a script with sbatch

9

Go to the Practice9 of github

Satisfaction survey

It is mandatory for you to fill this form to have your account extend :

https://itrop-survey.ird.fr/index.php/432222?lang=fr

Citations

If you use i-Trop Bioinformatics resources.

Thank you for citing us with:

“The authors acknowledge the ISO 9001 certified IRD i-Trop HPC (South Green Platform) at IRD montpellier for providing HPC resources that have contributed to the research results reported within this paper.
URL: https://bioinfo.ird.fr/- http://www.southgreen.fr
”

Projects

- Include a budget for bioinformatics resources in your answer to projects funding

- A need in hard drives, renew machines etc…
- Available quotations

- Contact : help, needs definition, quotations…

Thank you for you attention !