



GPU at CC-IN2P3


January 22-23, 2020

- ▶ **Hardware**
- ▶ **Software environment**
- ▶ **Usage**
- ▶ **Conclusion & perspectives**




K80 farm

- ▶ **10 workers:**
 - 2 Intel(R) Xeon(R) CPU E5-2640 (8 cores)
 - 128GB RAM
 - SSD disk
- ▶ - 2 Nvidia **Tesla K80** cards
(4 GPU Nvidia GK210, 12 Go DDR5 each)
- ▶ **40 GPU au total**
- ▶ Network
 - Infiniband interconnection

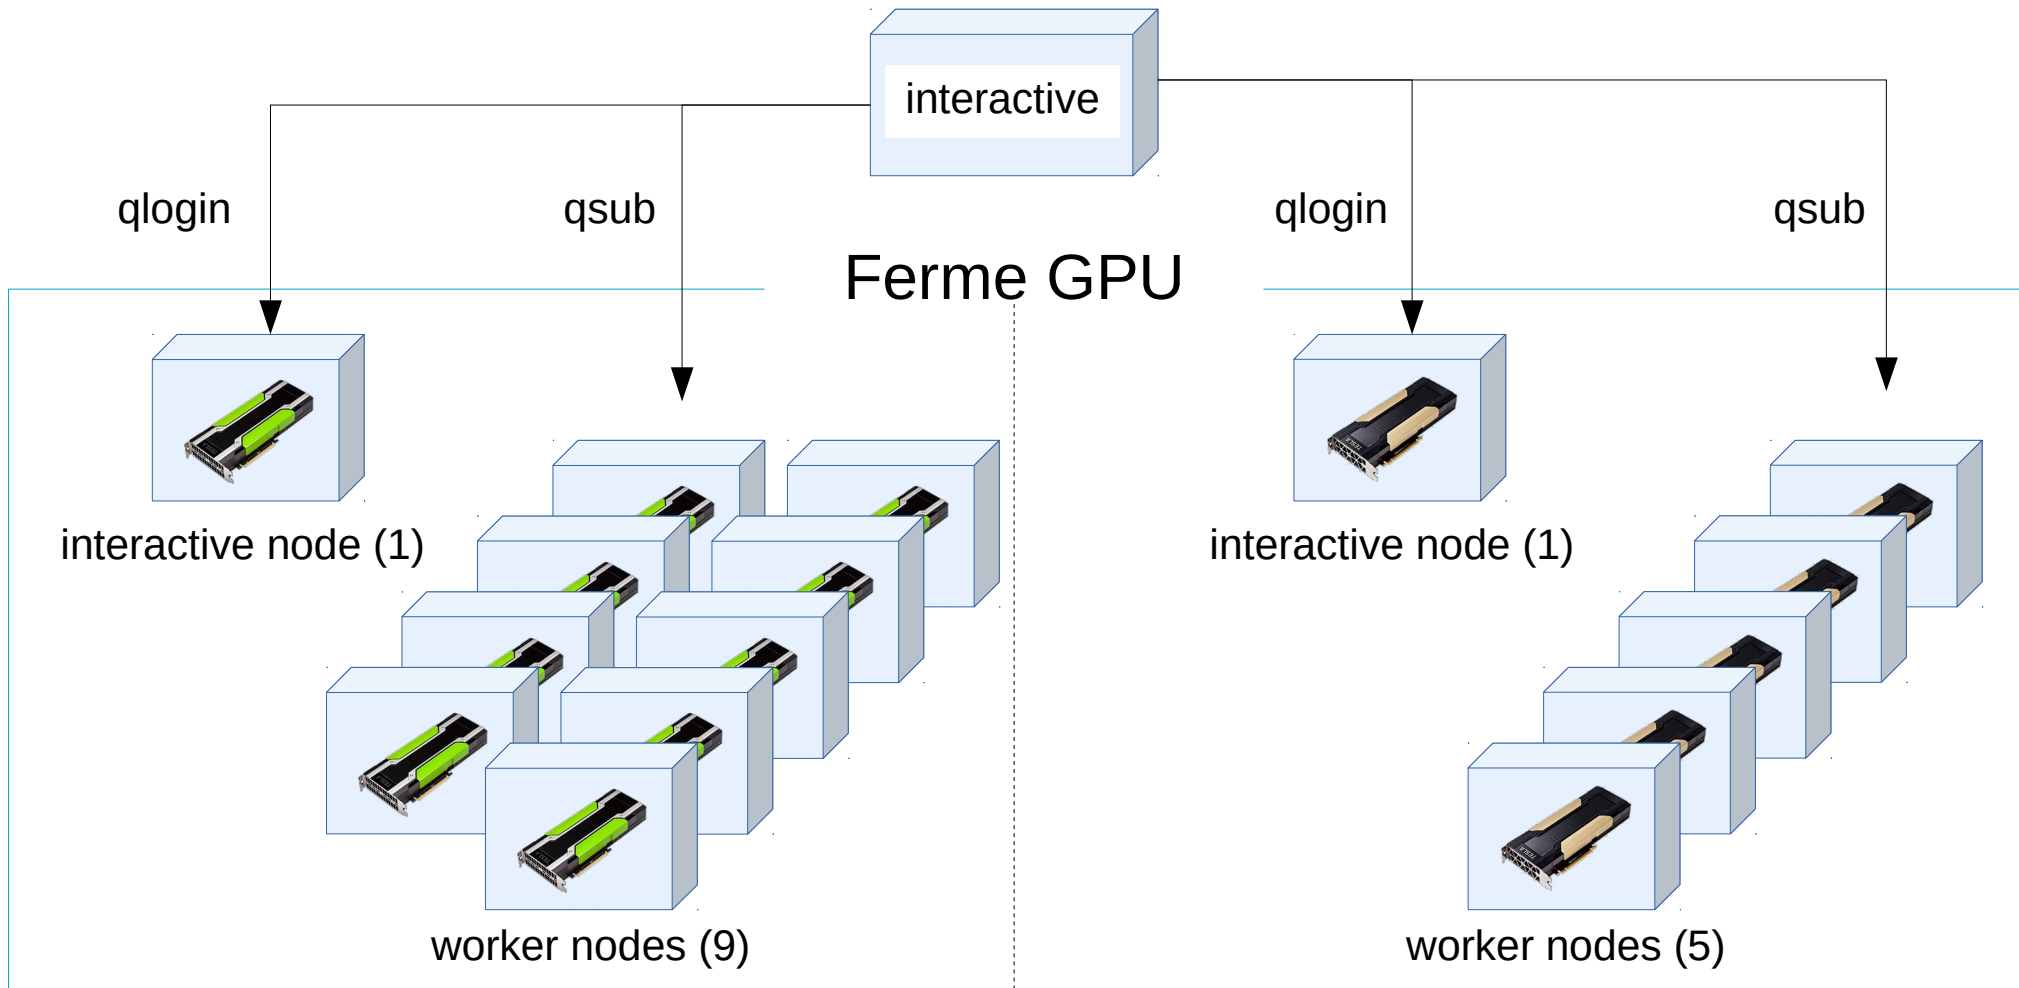


V100 farm

- ▶ **6 workers:**
 - 2 Intel(R) Xeon(R) Silver 4114 (10 cores)
 - 192GB RAM
 - SSD M2 disk
- ▶ - 4 Nvidia **Tesla V100** 32GB PCIe cards
- ▶ **24 GPU (total)**
- ▶ Network
 - NO Infiniband interconnection!



- ▶ First, request an access (authorisation required)
- ▶ Classical submission on Grid Engine (qsub) in multicore or parallel mode (https://doc.cc.in2p3.fr/jobs_gpu)
- ▶ Accounting rules still to be defined (K80 vs V100)



Libraries (drivers) available



OpenCL



OPEN MPI

Customized software

- ▶ Customized software provided thanks to **Singularity** containers



- ▶ Interactive Worker and Batch Worker nodes are the same in terms of architecture (same cpus, gpus, memory)
- ▶ Batch scheduler provides access to Interactive Worker nodes in shell mode



Interactive Worker access (qlogin)



Specifying the use of sps resources

Nb of gpus <1 - 4>

Specifying K80 or V100 farm

```
qlogin -l GPU=1,sps=1,GPUtype= 

|      |
|------|
| K80  |
| V100 |

 -q mc_gpu_interactive -pe multicores_gpu 4
```

Custom parameters

Queue

Nb of cpu cores

- ▶ Batch scheduler provides access on Batch Worker nodes to run a program



Batch submission (qsub)



Specifying the use of sps resources

Nb of gpus <1 - 4>

Specifying K80 or V100 farm

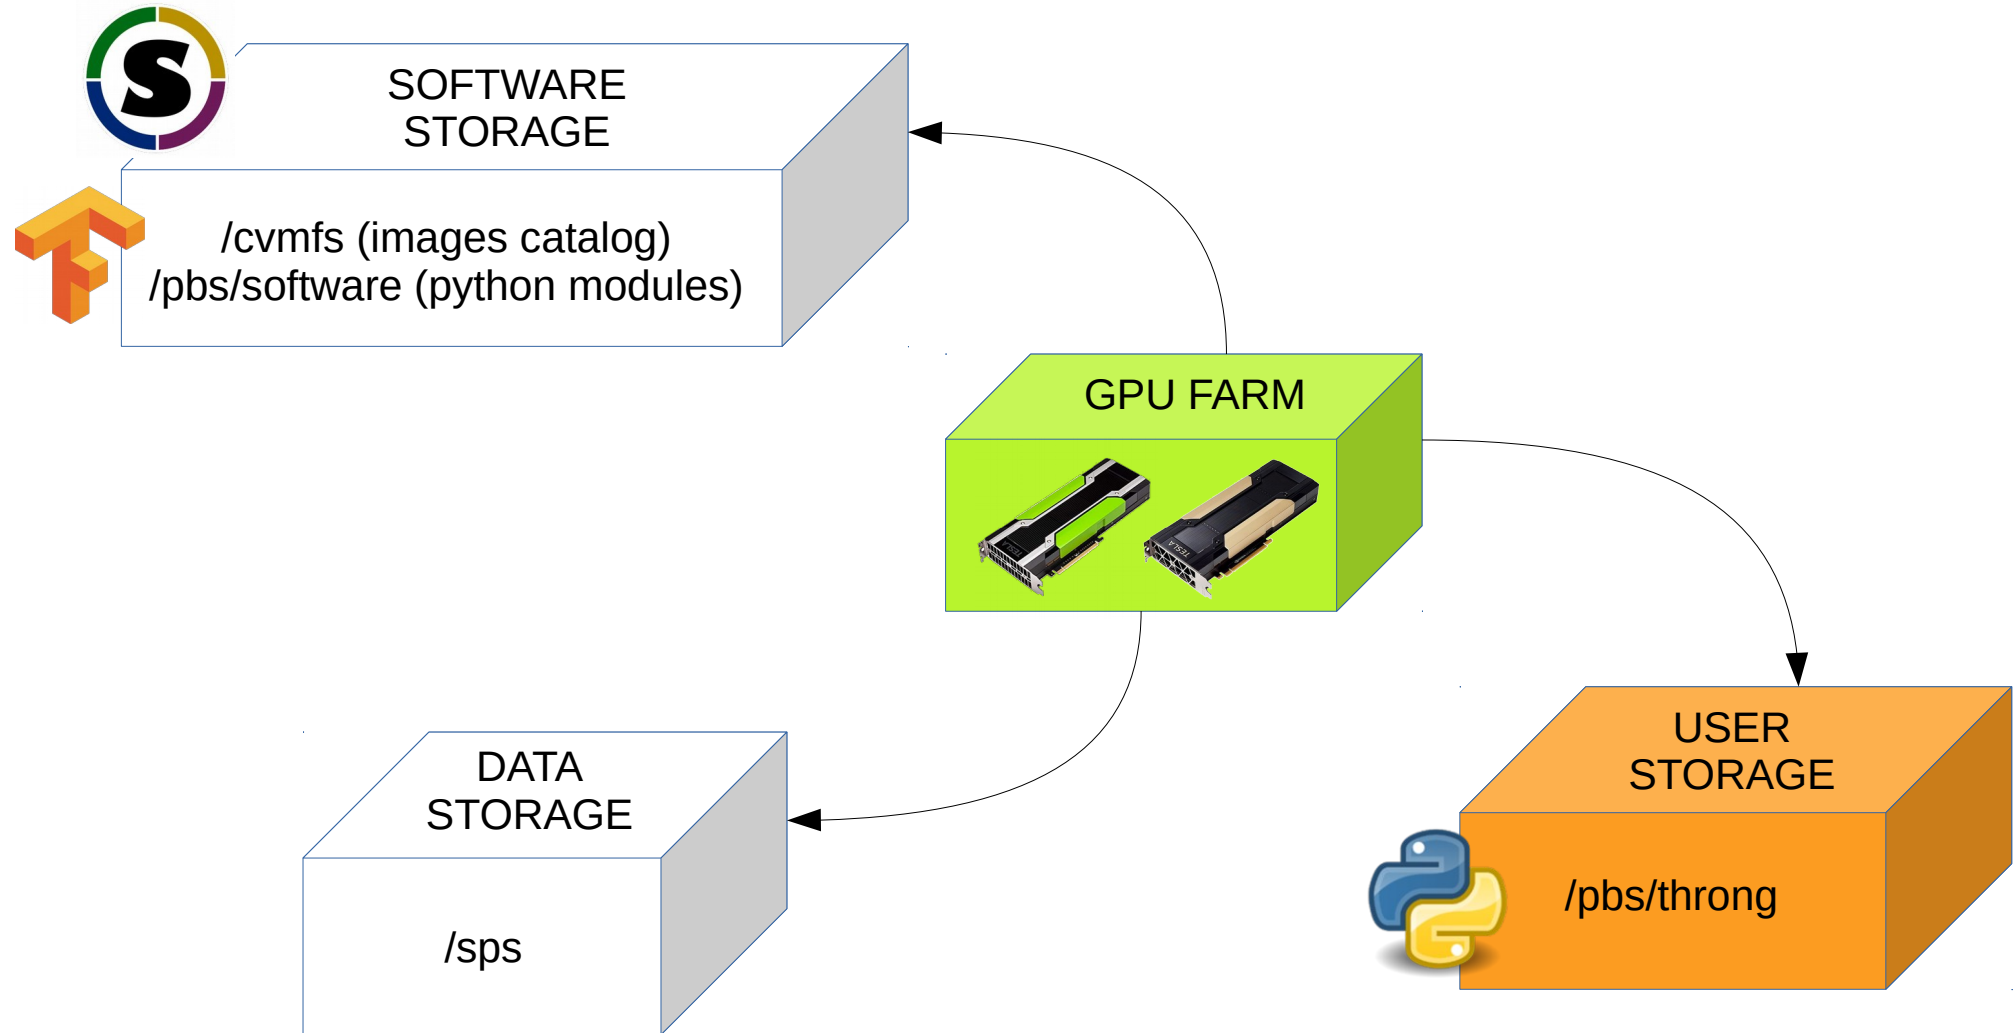
```
qsub -l GPU=1,sps=1,GPUtype=

|      |
|------|
| K80  |
| V100 |

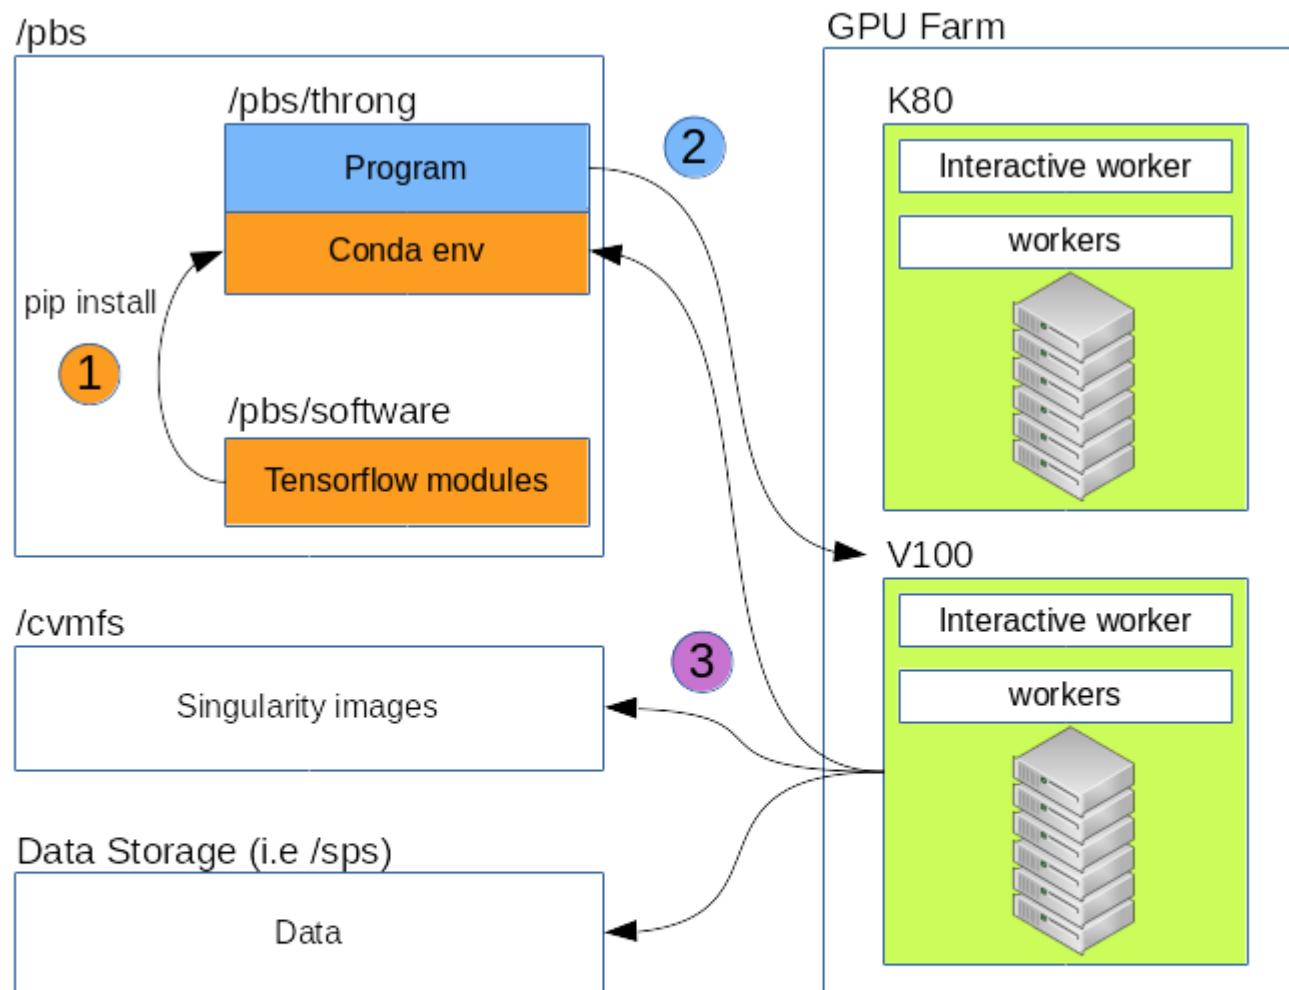
 [ options ] <file_to_execute>
```

Custom parameters

- ▶ Worker nodes can access different types of storage systems




Workflow example



- 1** Install the tensorflow module you need depending on your python environment (python 2.7 or python 3.6) and the GPU type you want (K80 or V100)
- 2** Submit your code to the GPU farm, specifying which image you want to run it, and your python environment
- 3** The GPU farm computes your code through the specified environment

Interactives (K80/V100)



K80

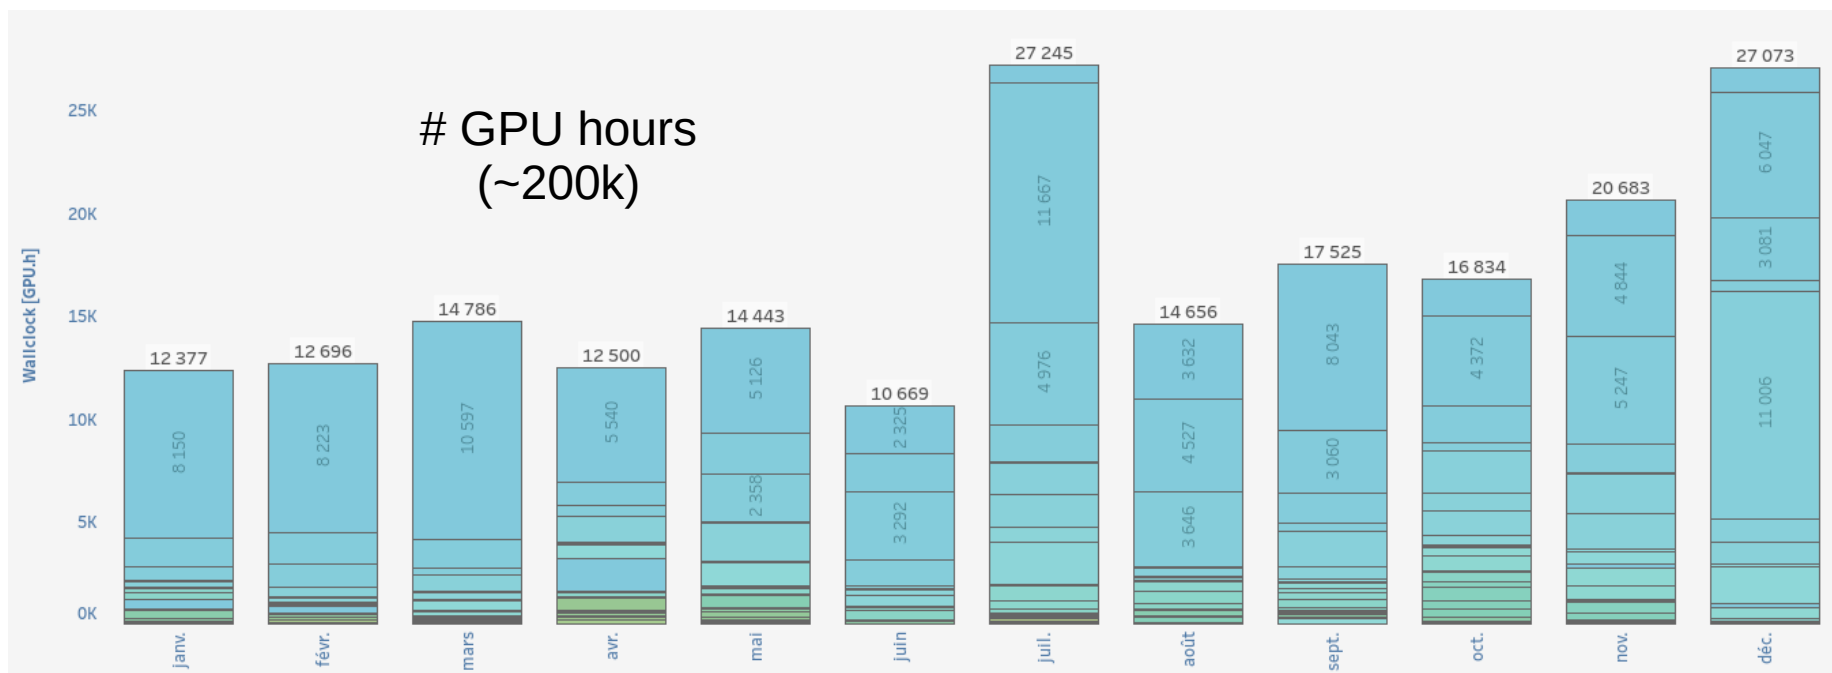

V100 in
production

August
holidays

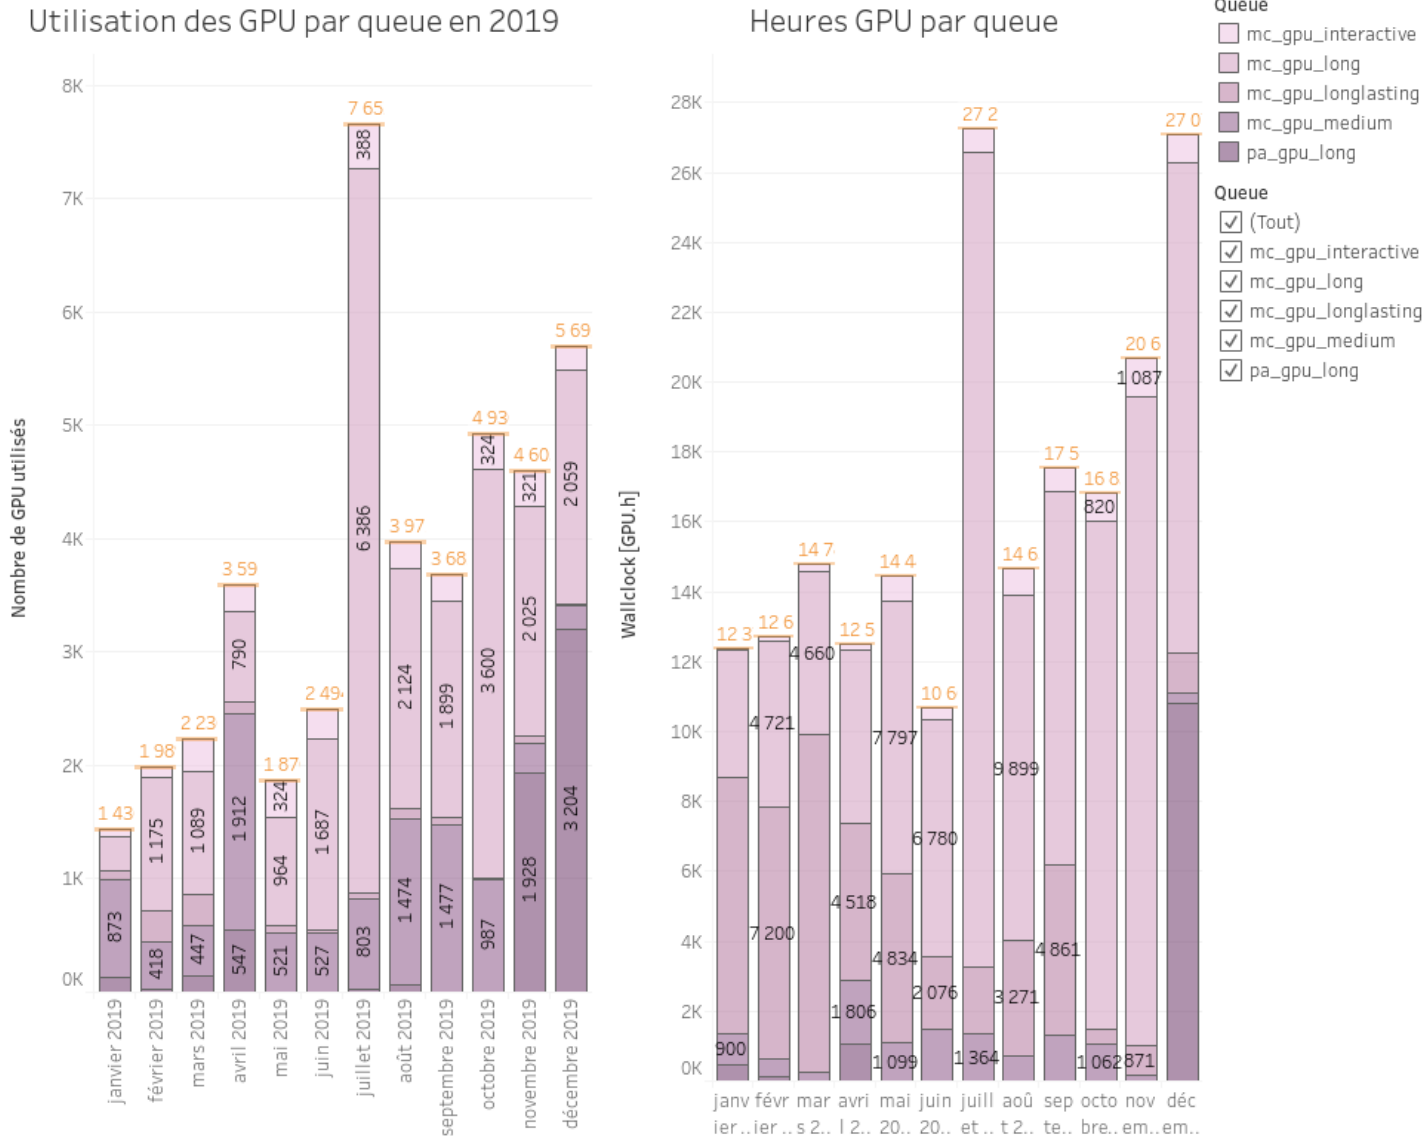
Christmas
holidays

- ▶ Not overloaded
- ▶ 2 interactive nodes still under used
- ▶ A batch system means regular production and not production by peaks
- ▶ Reservations (of GPU resources) not always well used

Consumption over 2019



Queues utilisation over 2019



- ▶ Long, long-lasting and parallel queues more and more used
- ▶ Medium one less and less used

- Total GPU-hours currently available
 - K80 ~351k hours & V100 ~210k hours → 561k hours available
- Resources requests (some are not yet validated, rough numbers)
 - IN2P3's labs / experiments asks for ~1140k hours
 - IJCLAB [ex-IMNC] (~1M), LSST (87k) and ATLAS (35k) being the top 3
 - Other experiments (no-IN2P3) asks for more than 1500k hours...
- Possible extension of the current farm but we should not be able to cope all the needs (budget constraints) → **IN2P3's groups will be favoured.**

- Current GPU farms overall well used
 - Better usage: more and more multi-GPU and multi-nodes jobs
 - Still « holes » though, especially in the K80 farm
- Some basic software environments provided
 - Conda & Python + TensorFlow & PyTorch + Singularity images
 - Do you need something else ? Please let us know
- Current discussion about extending the GPU farm
 - Come back to us if you would like to use some specific GPU

Questions ?

Bertrand.rigaud@cc.in2p3.fr

More questions? Ask us :
[OTRS ticketing System](#)

Thank you for your attention.