

Service GPU

Tuesday, October 13th

Adrien, Carlos & Julien

Status:

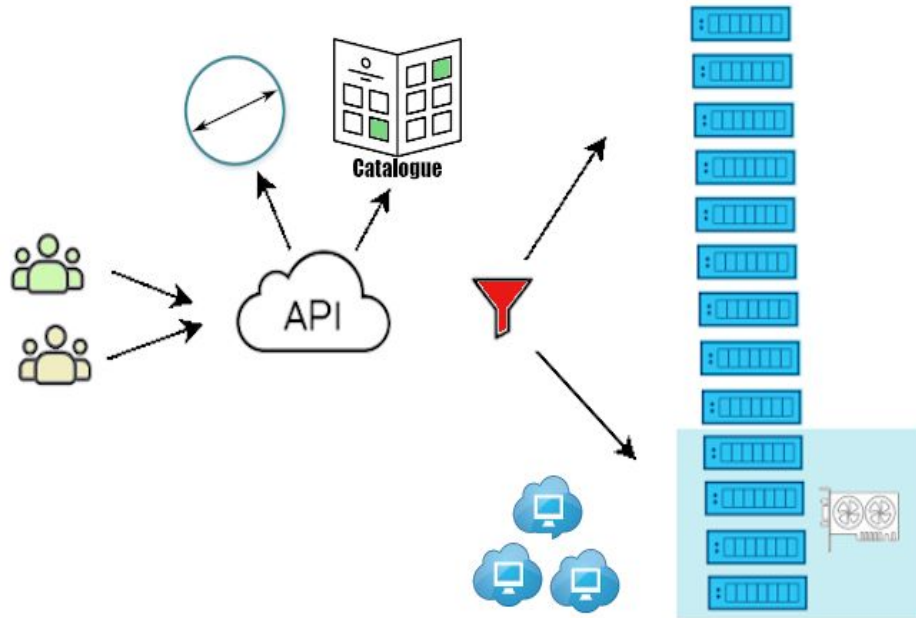
- 3 x Serveur Dell PowerEdge R740
 - 1 x server - 2 x Xeon(R) Silver 4216 CPU @ 2.10GHz 64 cores + Raiser gpu (3 pci)
 - 1 carte V100 32 GB
 - margpu01
 - 2 x server - 2 x Xeon(R) Silver 4210 CPU @ 2.20GHz 40 cores + Raiser gpu (3 pci)
 - 2 x 3 x GEFORCE RTC 2080 Ti - 11GB
 - margpu02 and margpu03
- 2 types of platform
 - 2 servers « standalone », margpu01 and margpu02
 - 1 server « cloud », margpu03

Status: Configuration « standalone »

- Cuda + drivers nvidia
- Software Anaconda +
- Monitoring (nagios)
- Share de ~ 50 Tb
- Auth via ssh (only key pair, no password)
- ~ 10 users (LHCb, Atlas, DarkSide,)

Status: Configuration « cloud »

- margpu03 was integrated in the cloud infra (IIAS) (test platform)
- Pass-through (instance have direct access the the gpu card)
- Preconfigured images
 - OS CentOS7 or CentOS8
 - Drivers nvidia
 - Software : cuda, anaconda, etc ...
- Usage of the API
 - Choose an image in the catalogue
 - Choose a flavor (ram, cores, disk +)
 - Create instance
 - Access through ssh
 - Use the machine
 - Free the resources when the job is done



Documentation and support

- Administration
 - We are tracking everything we are doing with elog
- User Documentation (work in progress)
 - store the documentation source on gitlab.in2p3.fr
 - the documentation will be published on gitlab pages
- User Support (work in progress)
 - use gitlab issue tracker for support and discussion follow up
 - use the gitlab wiki to share our experiences (successful or not)
- FAQ meeting on regular basis, every 3 months (or less if needed)
 - first was this morning

Preliminary feedbacks

- RTX2080TI allows real FP32 acceleration
This is code and model dependent but some tests shown RTX2080TI could be more efficient than the V100 in FP32
- V100 is about 10 times more expensive than the RTX2080TI, if you want V100, better to be sure you absolutely need double precision!
- Basic tests (small model, only use GPU) shown equivalent performances when using GPU with « standalone » server and « cloud » instances thanks to PCI Pass-through

On going discussion

- Migrate margpu01 and/or margpu02 to the cloud ?
 - pass margpu03 to prod infra
 - user test and feedback
 - report and address the question to next FAQ
- Pass-through VS virtual GPU (vgpu)
 - Pass-through: allocate the all GPU card
 - And when allocated to an instance, the GPU card is locked
 - vgpu may allow a more flexibility, to be explored
 - RTX2080TI can not be used with vgpu, only the V100 can