imec

GPULAB INTRODUCTION

TABLE OF CONTENTS

WHY USE GPULAB?

Access to a lot of GPU's:

44x GTX 1080 Ti

32x Tesla V100 (DGX-2, HGX-2)

…

Your pip/conda packages are installed and ready to use!

Choose any Docker image with your packages pre-installed

Isolated Storage

Separated storage per project

Automatic Job Scheduling

Jobs are started in FIFO order

WHAT IS GPULAB?

Thin wrapper around GPU-enabled Docker containers:

- Hides complexities of mounting storage, CPU/GPU isolation, etc.
- No need to install CUDA, Tensorflow, etc. on the machine yourself

Job Scheduler:

- Over multiple machines
- With 1 or more GPU's

Authentication:

- Via IDLab iLab.t or Fed4FIRE accounts
- Concept of 'projects' for sharing of resources

JUPYTERHUB
ON GPULAB

Available on
https://jupyterhub.ilabt.imec.be

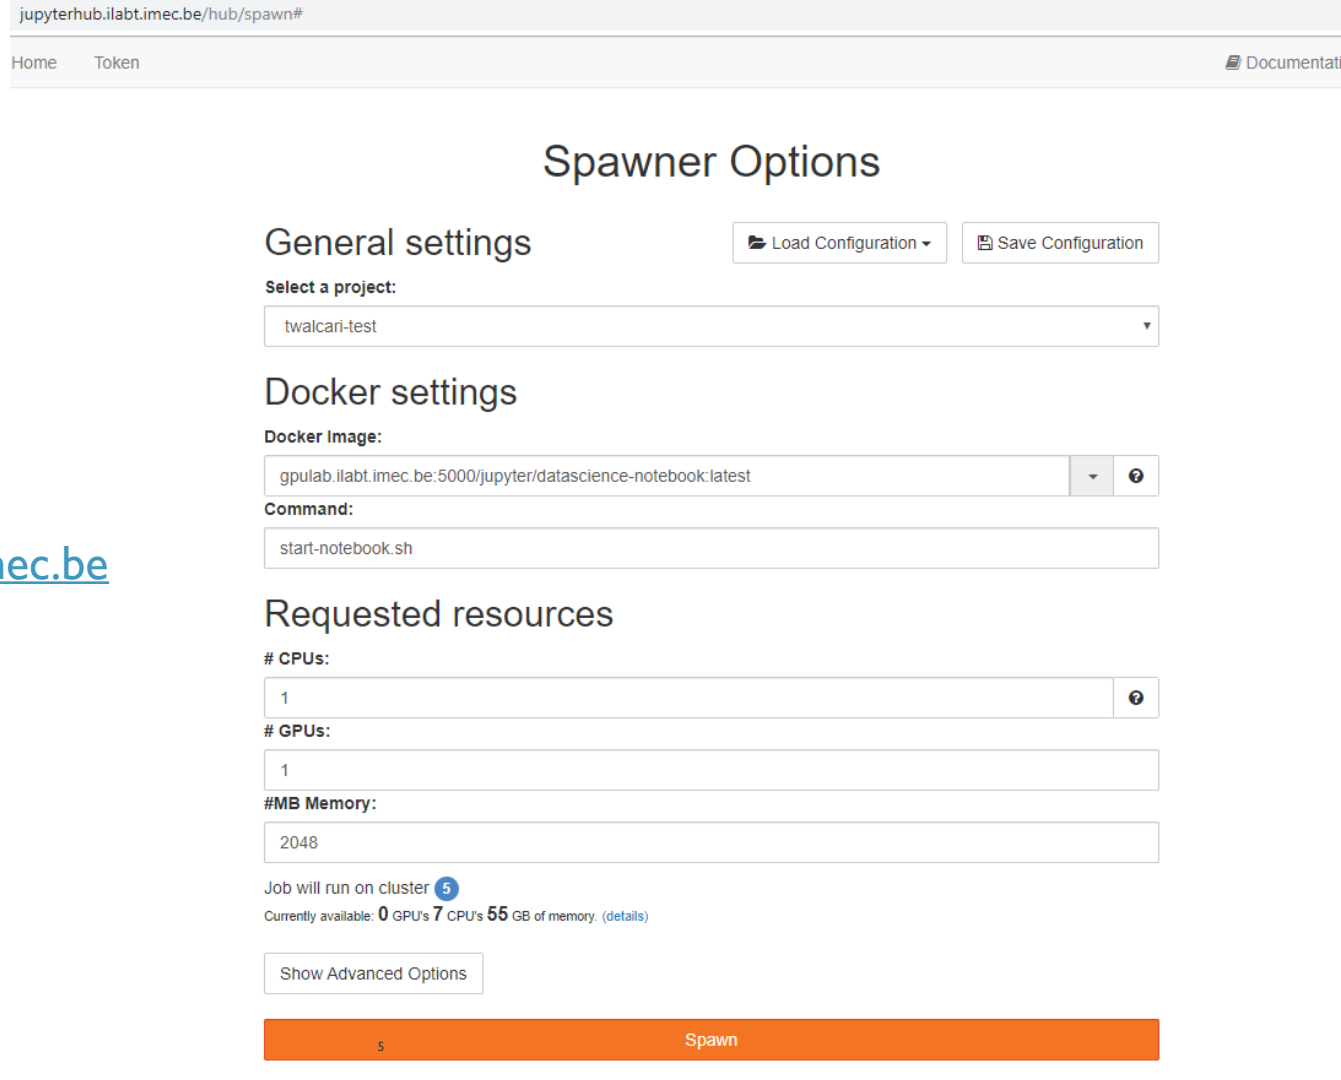

imec IDLab INTERNET & DATA LAB UNIVERSITEIT GENT Universiteit Antwerpen

JUPYTERHUB
ON GPULAB

- Generates and starts a GPULab job for you
- Redirects you to your Jupyter notebook server once started

- **Gotchas:**
  - Server start will timeout after 5 minutes (ex. No GPU's available in chosen cluster)
  - Job will be cancelled after 1 hour of inactivity in the browser, even if a computation is running!
  - Custom docker images must descend from `jupyter/base-notebook`

USING GPULAB

- Website https://gpulab.ilabt.imec.be for monitoring/submitting simple jobs

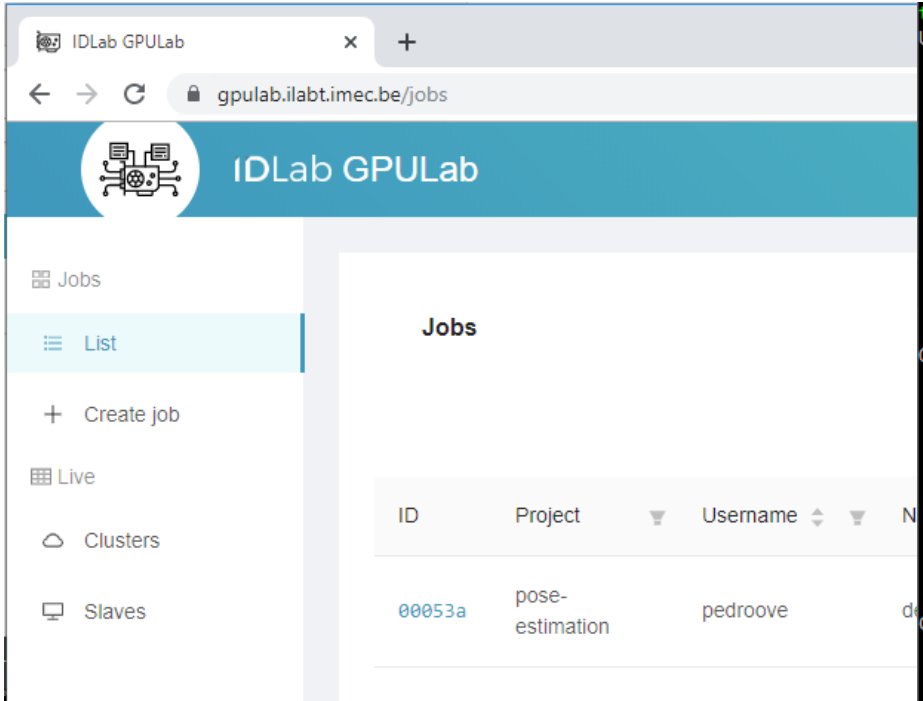

- `gpulab-cli` for submitting jobs from the command line

```
thijs@ibcn055:~$ gpulab-cli
Usage: gpulab-cli [OPTIONS] COMMAND [ARGS]...

  GPULab client version 1.9

  This is the general help. For help with a specific command, try:
     gpulab-cli <command> --help

  Send bugreports, questions and feedback to: jfedbugreports@ilabt.imec.be

  Documentation: https://doc.ilabt.imec.be/ilabt/gpulab/
  Overview page: https://gpulab.ilabt.imec.be/
  Overview page (for --dev): https://dev.gpulab.ilabt.imec.be/

Options:
  --cert PATH          Login certificate  [required]
  -p, --password TEXT  Password associated with the login certificate
  --dev                Use the GPULab development environment
  --stable             Use the GPULab stable environment (default)
  --servercert PATH    The file containing the servers (self-signed)
                       certificate. Only required when the server uses a self
                       signed certificate.
  --version            Show the version and exit.
  -h, --help           Show this message and exit.

Commands:
  cancel       Cancel running job
  clusters     Retrieve info about the available clusters. If a cluster_id is
               specified, detailed info about the slaves of that cluster is
```

USING GPULAB

SUBMITTING A JOB VIA THE CLI

```
thijs@ibcn055:~$ gpulab-cli submit --project twalcari-test < jupyter-scipy.json
87914bc6-10ec-11ea-93a1-d7177117bc9b
```

DEFINING A GPULAB JOB

GPULab-specific bookkeeping

```
{
  "jobDefinition": {
    "name": "helloworld",
    "description": "Hello world!",

    "clusterId": 1,
    "resources": { "gpus": 1, "systemMemory": 2000, "cpuCores": 2 },

    "dockerImage": "nvidia/cuda:10.2-cudnn7-devel",
    "command": "",

    "jobDataLocations": [ ],
    "portMappings": [ ],

    "environment": { },
  }
}
```

Passed to Docker for starting the container

DEFINING A GPULAB JOB

- Modify an example
  - Consult documentation on https://doc.ilabt.imec.be/ilabt/gpulab/jobdefinition.html
- Use the 'Create job' function on the website

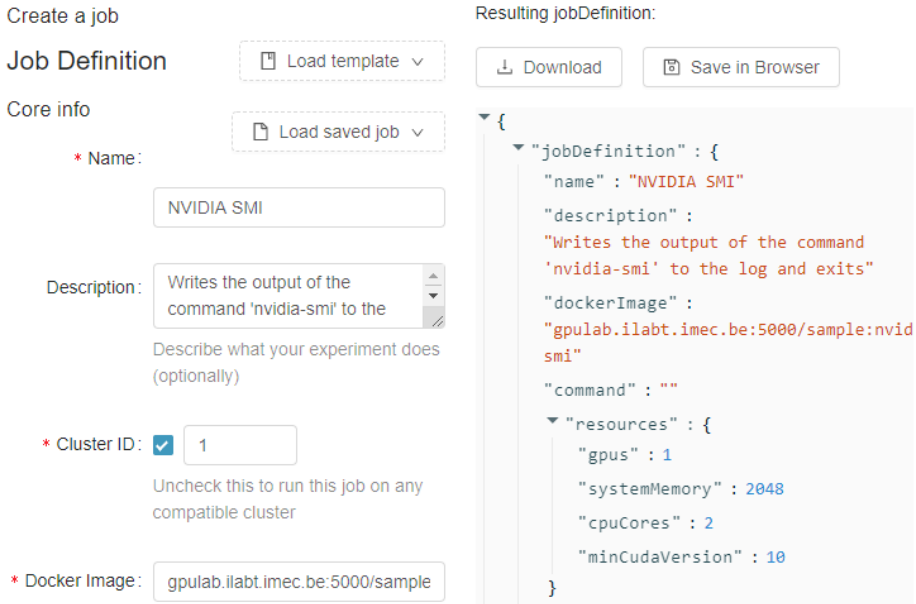

DEFINING A GPULAB JOB

DETERMINING THE CLUSTER AND RESOURCES TO REQUEST: VIA THE WEBSITE

gpulab.ilabt.imec.be/live/cluster

IDLab GPULab

Clusters

GPULab Version	Cluster ID	Comment	GPU model	GPU's	CPU's	Memory (GB)	Slaves	Running Jobs
stable	1	1x 2x GF GTX 1080 Ti	GeForce GTX 1080 Ti	2/ 2	16/ 16	31.47/ 31.47	1	0
stable	2	1x Tesla V100	Tesla V100-PCIE-32GB	0/ 1	10/ 12	23.62/ 31.62	1	1
stable	3	1x RTX2080	GeForce RTX 2080 Ti	1/ 1	12/ 12	31.66/ 31.66	1	0
stable	4	3x 11x GF GTX 1080 Ti	GeForce GTX 1080 Ti	5/ 39	28/ 98	244.93/ 989.58	4	27

DEFINING A GPULAB JOB

DETERMINING THE CLUSTER AND RESOURCES TO REQUEST:VIA THE CLI

```
~$ gpulab-cli clusters
----------+----------------------+----------------------------------------+--------+-------+------
ID        | GPU Model            | Comment                                | Slaves | GPUs  | CPUs
----------+----------------------+----------------------------------------+--------+-------+------
1 stable  |                      | 1x 2x GF GTX 1080 Ti                   | 0      | 0/0   | 0/0
----------+----------------------+----------------------------------------+--------+-------+------
2 stable  |                      | 1x Tesla V100                          | 0      | 0/0   | 0/0
----------+----------------------+----------------------------------------+--------+-------+------
3 stable  | GeForce RTX 2080 Ti  | 1x RTX2080                             | 1      | 1/1   | 12/12
----------+----------------------+----------------------------------------+--------+-------+------
4 stable  | GeForce GTX 1080 Ti  | 4x 11x GF GTX 1080 Ti                  | 7      | 19/30 | 61/93
----------+----------------------+----------------------------------------+--------+-------+------
5 stable  |                      | No GPU, shared CPUs                    | 0      | 0/0   | 0/0
----------+----------------------+----------------------------------------+--------+-------+------
6 stable  | Tesla V100-SXM3-32GB | 1x HGX-2 - 16xTesla V100               | 2      | 3/16  | 38/96
----------+----------------------+----------------------------------------+--------+-------+------
7 stable  | Tesla V100-SXM3-32GB | UAntwerp: 1x DGX-2 - 16xTesla V100     | 1      | 5/16  | 39/96
----------+----------------------+----------------------------------------+--------+-------+------
10 stable |                      | Development only - do not use          | 0      | 0/0   | 0/0
----------+----------------------+----------------------------------------+--------+-------+------
```

STORAGE ON GPULAB

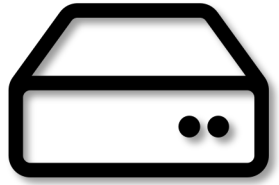

- Use jobDataLocations to set storage mount points
- `/project` is shared NFS storage per project between imec Virtual Wall 2 and Ghent-based GPULab-slaves
  - **Gotcha**: this storage is not available in Antwerp, separate storage on UAntwerp DGX machine!

```
"jobDataLocations": [
  {
    "mountPoint": "/project"
  },
  {
    "mountPoint": "/work",
    "sharePath": "/project/work"
  }
]
```

STORAGE ON GPULAB

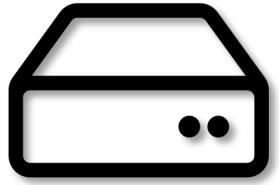

- `/project_scratch`
  - Project storage on a specific slave
  - Permanent, fast SSD (and large), but not accessible from any other slave
  - Only available on specific slaves!

Currently available on:

- HGX-2 at UGent:
  - 100 TB NVMe storage (`/project_scratch`)
- DGX-2 at UAntwerp:
  - 28 TB NVMe storage (`/project_scratch`)
  - **Gotcha:** `/project` is an alias for `/project_scratch`

```
"jobDataLocations": [
  {
    "mountPoint":
        "/project_scratch"
  },
]
```

STORAGE ON GPULAB

IMPORTING YOUR DATA INTO GPULAB

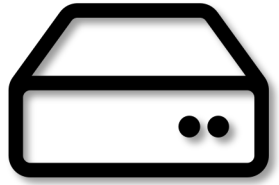

- Use Jupyter notebooks or SSH-access to your jobs to explore your storage
  - **Gotcha**: No SCP available
- Start job with SFTP-server

- **`/project:`** Swap in a Virtual Wall 2 machine to prepare your data in `/groups/ilabt-imec-be/<projectname>`

- **`/project_scratch:`** Copy from `/project` to `/project_scratch` for fast storage on HGX-2 machine

Only applicable for UGent NFS storage

Documentation: https://doc.ilabt.imec.be/ilabt/gpulab/storage.html

EXPOSING PORTS OF YOUR CONTAINER

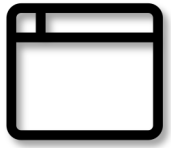

- You can define ports to be exposed in `portMappings`
- Use `containerPort` to specify which port of you container you want to access
- Host address/port is determined during job scheduling
- Hosts have no public IPv4 address!
  - Only public IPv6
  - Or private IPv4 (via UGent idlab-vpn)
- **Gotcha:** exposed ports on UAntwerp DGX are only available within UAntwerp IDLab VPN

```
"portMappings": [
  {
    "containerPort": 5000
  },
  {
    "containerPort": 5001,
    "hostPort": 5001
  }
]
```

WARNING: Job will fail if another container is already mapped to that port!

EXPOSING PORTS OF YOUR CONTAINER

FINDING THE HOST/PORT VIA THE WEBSITE OR CLI

```
thijs@ibcn055:~$ gpulab-cli jobs bdad9168
         Job ID: bdad9168-1058-11ea-93a1-1f5b50ba1116
           Name: Jupyter SciPy
    Description: Scipy jupyter notebook server
        Project: F000884-practica
       Username: dlpract
   Docker image: jupyter/scipy-notebook:latest
        Command:
         Status: RUNNING
        Created: 2019-11-26T15:26:31+01:00
  State Updated: 2019-11-26T15:26:40+01:00
         Queued: 2019-11-26T15:26:31+01:00
     Cluster ID: 4
      Worker ID: 7
    Worker Name: n053-02
  Port Mappings: 8888/tcp -> 33216
    Worker Host: n053-02.wall2.ilabt.iminds.be
        Started: 2019-11-26T15:26:38+01:00
       Duration: 6 minutes, 44 seconds
       Finished: -
       Deadline: 2019-11-27T01:26:38+01:00
```

Job bdad9168-1058-11ea-93a1-1f5b50ba1116

RUNNING

Refresh Cancel

Project: F000884-practica

Description: Scipy jupyter notebook server

General Info | Logs | Debugging Logs | Raw Job JSON

Job Description

Job Status

Created on November 26, 2019 3:26 PM

Queued on November 26, 2019 3:26 PM

Started on November 26, 2019 3:26 PM

Running ...

Job Execution Environment

Summary: 4B

Cluster: 4

Slave: n053-02

Port mappings

Port 8888/tcp → n053-02.wall2.ilabt.iminds.be:33216

CHECKING THE LOGS OF YOUR CONTAINER
ON THE CLI

```
thijs@ibcn055:~$ gpulab-cli log bdad9168
2019-11-26T14:26:38.818588342Z Executing the command: jupyter notebook
2019-11-26T14:26:39.233037365Z [I 14:26:39.232 NotebookApp] Writing notebook server cookie secret to /home/jovyan/.local
/share/jupyter/runtime/notebook_cookie_secret
2019-11-26T14:26:41.245393769Z [I 14:26:41.244 NotebookApp] JupyterLab extension loaded from /opt/conda/lib/python3.7/si
te-packages/jupyterlab
2019-11-26T14:26:41.245434141Z [I 14:26:41.245 NotebookApp] JupyterLab application directory is /opt/conda/share/jupyter
/lab
2019-11-26T14:26:42.955435040Z [I 14:26:42.955 NotebookApp] Serving notebooks from local directory: /home/jovyan
2019-11-26T14:26:42.955473156Z [I 14:26:42.955 NotebookApp] The Jupyter Notebook is running at:
2019-11-26T14:26:42.955494464Z [I 14:26:42.955 NotebookApp] http://e1e266b4d721:8888/?token=0650d6f66e2f98fb4eb513f4c993
9c063b85a0d3c4541531
2019-11-26T14:26:42.955546481Z [I 14:26:42.955 NotebookApp]  or http://127.0.0.1:8888/?token=0650d6f66e2f98fb4eb513f4c99
39c063b85a0d3c4541531
2019-11-26T14:26:42.955586375Z [I 14:26:42.955 NotebookApp] Use Control-C to stop this server and shut down all kernels
(twice to skip confirmation).
2019-11-26T14:26:42.960878486Z [C 14:26:42.960 NotebookApp]
2019-11-26T14:26:42.960896909Z
2019-11-26T14:26:42.960902168Z     To access the notebook, open this file in a browser:
2019-11-26T14:26:42.960906520Z         file:///home/jovyan/.local/share/jupyter/runtime/nbserver-6-open.html
2019-11-26T14:26:42.960911507Z     Or copy and paste one of these URLs:
2019-11-26T14:26:42.960915909Z         http://e1e266b4d721:8888/?token=0650d6f66e2f98fb4eb513f4c9939c063b85a0d3c4541531
2019-11-26T14:26:42.960920383Z      or http://127.0.0.1:8888/?token=0650d6f66e2f98fb4eb513f4c9939c063b85a0d3c4541531
2019-11-26T14:26:46.694747826Z [I 14:26:46.694 NotebookApp] 302 GET / (192.168.124.38) 0.82ms
2019-11-26T14:26:46.740732997Z [I 14:26:46.740 NotebookApp] 302 GET /tree? (192.168.124.38) 0.95ms
```

CHECKING THE LOGS OF YOUR CONTAINER
ON THE WEBSITE

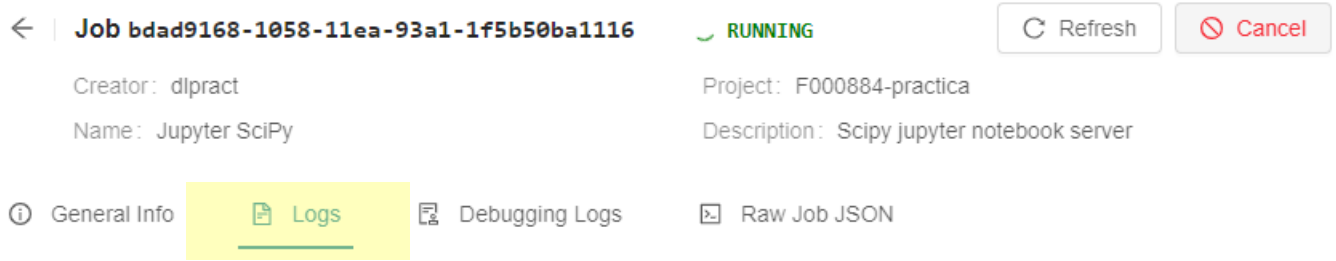

GETTING SSH-ACCESS TO YOUR CONTAINER

Use gpulab-cli ssh <job-id>

```
thijs@ibcn055:~$ gpulab-cli ssh 87914bc6-10ec-11ea-93a1-d7177117bc9b
Warning: Permanently added the ECDSA host key for IP address '2001:6a8:1d80:27::242' to the list of known hosts.
setsockopt IPV6_TCLASS 16: Operation not permitted:
The authenticity of host 'n051-02.wall2.ilabt.iminds.be (<no hostip for proxy command>)' can't be established.
ECDSA key fingerprint is SHA256:Q88ZLhdKW8JNXc/Mf5HARu78RhS+99UzxNzuCTa5DCE.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'n051-02.wall2.ilabt.iminds.be' (ECDSA) to the list of known hosts.
jovyan@441af988cc0d:~$ w
 08:08:27 up 13 days, 15:55,  0 users,  load average: 16.86, 16.88, 17.19
USER     TTY      FROM             LOGIN@   IDLE   JCPU   PCPU WHAT
jovyan@441af988cc0d:~$
```

BEST PRACTICES

1. Develop and test your code locally or in a Jupyter notebook
    - Create your own Docker image with custom software if necessary
2. Scale up to the full dataset once your code is ready
3. Add error-catching and retry mechanisms to your I/O operations
4. Add checkpointing for your intermediate results

GET STARTED!

- **Documentation**: http://doc.ilabt.imec.be

- **GPULab**: https://gpulab.ilabt.imec.be
- **JupyterHub**: https://jupyterhub.ilabt.imec.be

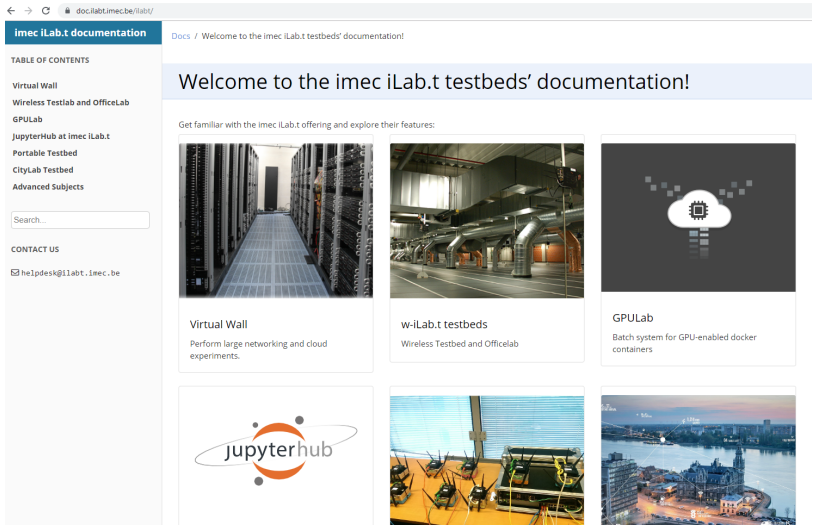

SUPPORT helpdesk@ilabt.imec.be

Mattermost https://mattermost.ilabt.imec.be

Channels:
- GPULab Support
- JupyterHub support

GET STARTED!
GET AN ACCOUNT

- GPULab uses accounts/projects from the IDLab Testbeds Portal:
  https://account.ilabt.imec.be

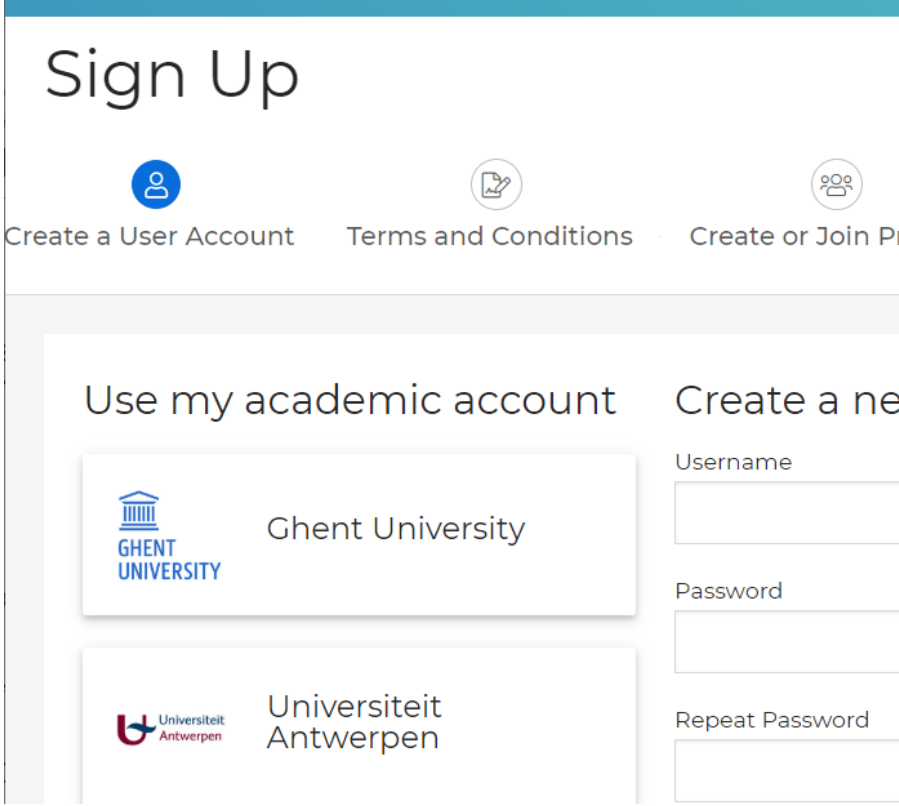

- Signup using your University Login

- Request a new project to isolate your file storage

- or -

- Join the 'GPULab UA Tutorial' project for testing:

https://account.ilabt.imec.be/invite/gpulab_ua_tut?key=V6CjOYsbnozUIEhc

GET STARTED!
SETTING UP THE CLI

- Download `gpulab-client-2.0.tar.gz` from https://doc.ilabt.imec.be/ilabt/gpulab/
- Install using `sudo pip3 install gpulab-client-2.0.tar.gz`
- Get your Login Certificate (PEM) from https://account.ilabt.imec.be/

```
thijs@ibcn055:~$ gpulab-cli --cert login_ilabt_imec_be_twalcari@ugent.be.pem clusters
+---------+----------------------+-----------------------------------------+--------+-------+------
| ID      | GPU Model            | Comment                                 | Slaves | GPUs  | CPUs
+---------+----------------------+-----------------------------------------+--------+-------+------
| 1 stable |                     | 1x 2x GF GTX 1080 Ti                    | 0      | 0/0   | 0/0
+---------+----------------------+-----------------------------------------+--------+-------+------
| 2 stable |                     | 1x Tesla V100                           | 0      | 0/0   | 0/0
+---------+----------------------+-----------------------------------------+--------+-------+------
| 3 stable | GeForce RTX 2080 Ti | 1x RTX2080                              | 1      | 0/1   | 10/12
+---------+----------------------+-----------------------------------------+--------+-------+------
| 4 stable | GeForce GTX 1080 Ti | 4x 11x GF GTX 1080 Ti                   | 7      | 21/30 | 31/93
+---------+----------------------+-----------------------------------------+--------+-------+------
| 5 stable |                     | No GPU, shared CPUs                     | 0      | 0/0   | 0/0
```

embracing a better life