GPU MINI-CLUSTER @RGNC

Version 2, 07-05-2021

Miguel Ángel Martínez del Amor,
Research Group on Natural Computing
University of Seville

0. Contents

1. Description

Our mini-cluster is placed at the Department of Computer Science and Artificial Intelligence facilities. It is currently composed of **3** servers, named after the highest peaks in Spain. Their specifications are detailed in www.gcn.us.es/gpucomputing

The configuration is the following:

- **Aneto** server can be openly accessed through a public IP of the university of Seville (we can give you the *IP if you collaborate with us*).
- **Mulhacen** server is hidden behind Aneto server in a private network.
- **Teide** server is hidden behind Aneto server in a private network.

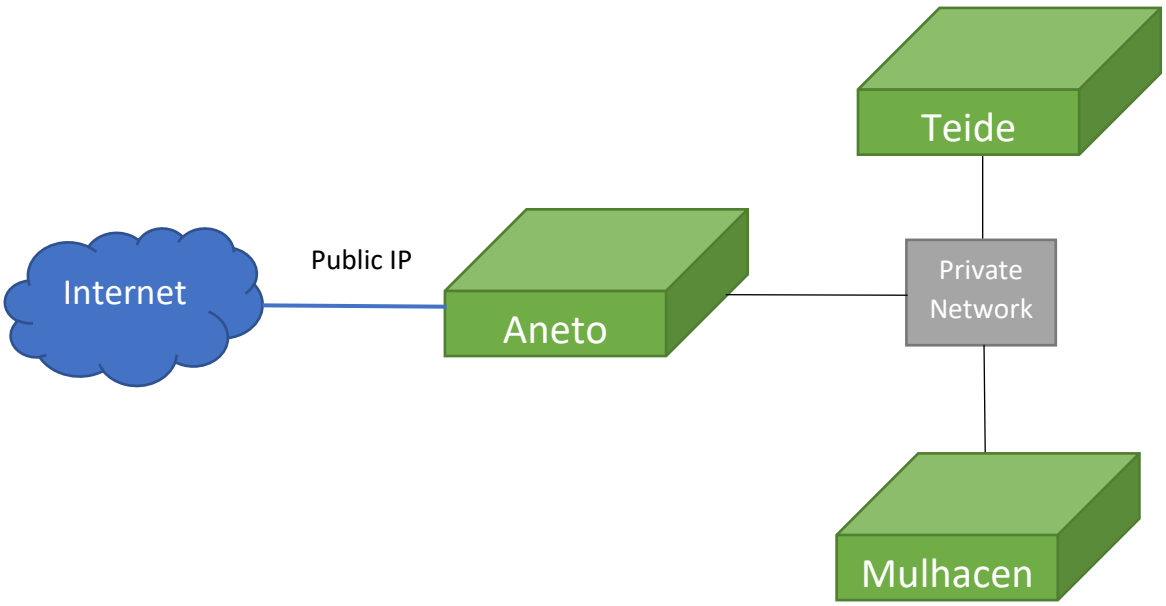

2. Client configuration

If you want to access each node by its name from your machine, do the following. Open the file $HOME/.ssh/config in your home directory in your computer (you should have a folder ".ssh" in your home directory and create the file "config" if it doesn't exist), and paste the following text. Please, replace **GPU-RGNC** by the public IP of our server, and **USER** by the username given to you in the mini-cluster:

```
### First jumphost. Directly reachable
Host aneto
        HostName GPU-RGNC
        User USER
        ProxyCommand none
        ForwardAgent yes
        GSSAPIAuthentication no

### Host to jump to via aneto
Host mulhacen
        HostName mulhacen
        ProxyCommand ssh aneto -W %h:%p
        ForwardAgent yes
        GSSAPIAuthentication no

### Host to jump to via aneto
Host teide
        HostName teide
        ProxyCommand ssh aneto -W %h:%p
        ForwardAgent yes
        GSSAPIAuthentication no
```

If you don't want to enter twice your password when accessing Teide or Mulhacen through Aneto, execute the following in your machine (replace **USER** by your username given in the mini-cluster):

1. `ssh-keygen -t rsa -b 2048` (press enter for all the prompted questions)
2. `ssh-copy-id USER@mulhacen` (insert your password all the required times)
3. `ssh-copy-id USER@teide` (insert your password all the required times)

Try "ssh **USER**@mulhacen", and check that you have to only provide once your password.

3. Access

If you have configured your ssh client as in previous section, you can access to each node independently. So, depending on the server (replace **USER** by your username in the mini-cluster):

- Access to Aneto: `ssh USER@aneto`

- Access to Teide: `ssh USER@teide`
- Access to Mulhacen: `ssh USER@mulhacen`

Once you have accessed to the desired server, you can run your programs remotely through ssh. If you need to copy files to and from the servers, do the following (replace **USER** by your username and **NODE** by your desired node, either aneto, mulhacen or teide):

- Copy to a NODE: `scp YOUR_FILE USER@NODE:`
- Copy from a NODE: `scp USER@NODE:YOUR_FILE .`
- Copy a folder to a NODE: `scp -r YOUR_FOLDER USER@NODE:`
- Copy a folder from a NODE: `scp -r USER@NODE:YOUR_FOLDER .`

4. Working in a node

Once you are logged in a node, you can check which are the GPUs available and their status typing: `nvidia-smi`. By default you are given GPU number 0. If it is busy, or you want to use GPU number 1, type: `export CUDA_VISIBLE_DEVICES=1`. If you want to use both GPUs, type: `export CUDA_VISIBLE_DEVICES=0,1`. If you want to use GPU number 0 again, type: `export CUDA_VISIBLE_DEVICES=0`

We use **cvmfs** in Mulhacen and Teide for accessing different compilers and libraries versions. In order to use it:

1. `ls /cvmfs/sft.cern.ch/lcg/releases` (this will mount the remote file system)
2. If you want to use another version of *gcc* than the default one, select your version (see `ls /cvmfs/sft.cern.ch/lcg/releases/gcc`), and then (for example, for GCC 8.3.0), type:
   `source /cvmfs/sft.cern.ch/lcg/releases/gcc/8.3.0/x86_64-centos7/setup.sh`
3. Some other libraries instead of requiring a source, just need to adjust your $PATH or your $LD_LIBRARY_PATH to that folder (e.g. for cmake)

If you need to execute a remote Jupyter notebook environment, then do the following (replace **USER** by your username and **NODE** by your desired node, either aneto, mulhacen or teide):

1. Type in your machine: `ssh -L 8888:localhost:8888 USER@NODE`
2. In the node you chose, type:
   `jupyter notebook --port=8888 --no-browser`
   Copy the URL with the token shown at the end.
3. Now in your browser, paste the URL you copied before.

You can use the remote application development in NVIDIA Nsight since version 6.5. More info: http://devblogs.nvidia.com/parallelforall/remote-application-development-nvidia-nsight-eclipse-edition/

There is no problem if you want to use our servers for remote application development. However, for Mulhacen and Teide you have to do some extra work, since it cannot be seen by Nsight. A workaround for this is to do the following (replace USER by your username and NODE by your desired node mulhacen or teide):

- In Nsight, go to properties of the project, Build → Target Systems.
- Select Manage..., and Add a new server
- Configure the new server as follows:
    - Host name: localhost
    - User name: *USER*
    - Label: *NODE*
    - System type: ssh, port: 9999
- Once you have that, just do the following in a Terminal before accessing the server via Nsight:
    - To run and profile an application, type:
      `ssh -L 9999:localhost:22 USER@NODE`
    - To debug an application (if you have configured port 2345 for the debugger), type:
      `ssh -L 2345:localhost:2345 USER@NODE`

Have fun, and wish you an efficient code!