

Instructions to use Cheaha GPU Server for Deep Learning

Sai Sumana Narva, Da Yan (yanda@uab.edu)

UAB Department of Computer Science

Step 1: Login to Cheaha

- For **Mac/Linux** users:
`ssh [your_Blazer_ID]@cheaha.rc.uab.edu`
- For **Windows** users:
Using **putty** to log in

After typing your Blazer password, you will be on the Cheaha login server node `login001`

Step 2: Now we need to get another machine to run our deep learning job (**do not** use `login001` to run or you will get warned by IT people). Use the following command to ask for a GPU server.

- `srun --ntasks=4 --cpus-per-task=4 --mem-per-cpu=32000 --time=08:00:00 --partition=pascalnodes --job-name=JOB_NAME --gres=gpu:4 --pty /bin/bash`

You may change the green parts according to your need. You must ask for `--mem-per-cpu=32000`, as if it is small, you are very likely to get assigned somewhere with many users running their jobs.

Now, you should be assigned to a GPU node `c0xxx`

Note that you have a time limit of `08:00:00` hours here and you can request more; be sure to turn on TensorFlow's checkpointing so you can recover the weight parameters from the latest checkpointed parameters next time to continue training.

You can bypass the time limit by directly ssh to `c0xxx` once you know such a host exists, but please do not do that as you will get warned by IT people.

Step 3: If this is your first time using Cheaha GPU server (skip this step if you have done it before), please follow the instructions below:

- `echo ". /share/apps/rc/software/Anaconda3/5.2.0/etc/profile.d/conda.sh"`
`>> ~/.bashrc`
- `echo "conda activate" >> ~/.bashrc`
- `source ~/.bashrc`

This will set up your `conda` so that you can enter an `conda environment` using "conda activate" in the future.

Now we need to create a conda environment, let's call it "`deep`":

- `git clone https://gitlab.rc.uab.edu/rc-data-science/horovod-environment.git`
`/data/user/$USER/nbotw` (--- download environment "nbotw" meaning 'notebook of the time' maintained by UAB IT)
- `module load cuda10.0/toolkit` (for using GPU)
- `module load Anaconda3/5.2.0`
- `cd /data/user/$USER/nbotw`
- `conda env create -f nbotw.yml --name deep` (--- this command copies "nbotw" into "`deep`", it will take a while as it needs to install all packages into "`deep`")

Step 4: Please make sure you have done Step 3. Now if you have set up your conda environment, say "`deep`", as in Step 3, you can now enter the environment (Note: if you just installed, we find that you may not be able to see GPUs unless you exit and restart from Step 1):

- `module load Anaconda3/5.2.0`
- `conda activate deep`
- `module load cuda10.0/toolkit`

Now you entered your previously created environment "`deep`".

Step 5: To confirm if GPU is being used user can [start python](#) and run the following commands:

- `from tensorflow.python.client import device_lib`
- `device_lib.list_local_devices()`

Make sure you see the GPUs, not just CPU.

Step 6: You may start Jupyter Notebook server on `c0xxx` by running the following commands

- `unset XDG_RUNTIME_DIR`
- `jupyter notebook --no-browser --ip=$host`

You should see Jupyter Notebook running and provide you an URL like:

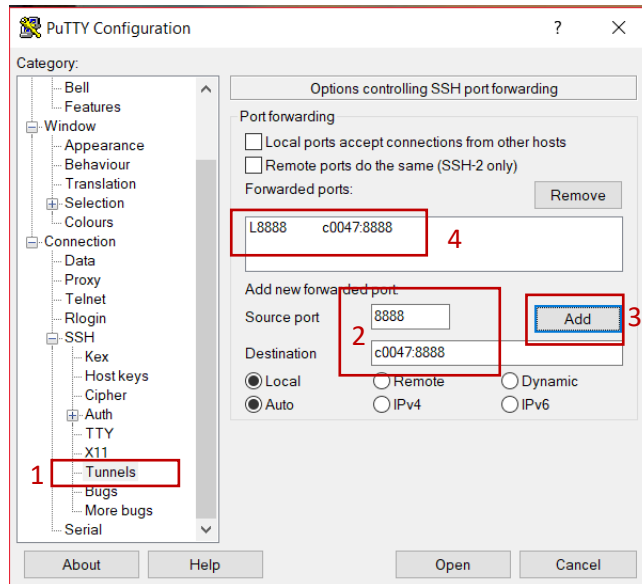
- `http://c0047:8888/?token=73da89e0eabdeb9d6dc1241a55754634d4e169357f60626c&token=73da89e0eabdeb7d6dc1241a55754634d4e169357f60626c`

Next, you want to use a browser in your local machine to connect to the Jupyter Notebook server. Run the following command [on your local machine](#):

- For Mac/Linux users:
 - `ssh -L 88XX:c0XXX:88XX [your_Blazer_ID]@cheaha.rc.uab.edu`

`c00xx` is the hostname of your assigned GPU server where you started Jupyter Notebook.

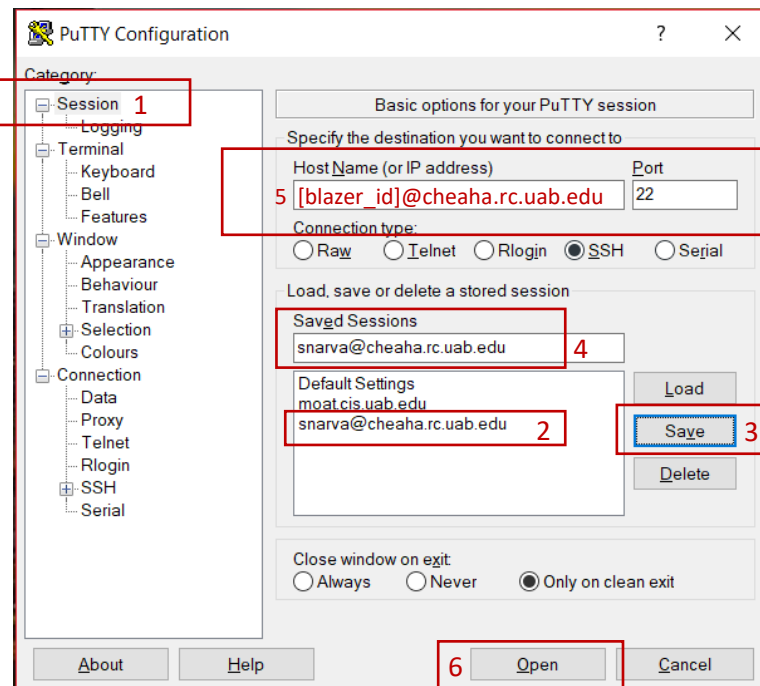
- For Windows users:



Open **Tunnels** on putty, enter the source port and destination hostname:port. Be sure to **click the ADD** button.

Now, go back to **Session**, click session `[your_Blazer_ID]@cheaha.rc.uab.edu`, and then click the **SAVE** button.

Then log in by typing the hostname `[your_Blazer_ID]@cheaha.rc.uab.edu`.



Now you can open browser on local system and paste the URL returned by the Jupyter Notebook server by replacing the `hostname of the server` with `localhost` (i.e. `localhost:8888` instead of `c0047:8888`).

Step 7 (If you want to monitor your GPU usage): open a new console and log in to Cheaha following Step 1. Then directly ssh to your previously assigned node `c0xxx`.

- `ssh c0xxx`

Now you are on `c0xxx`, run the following to monitor GPU usage

- `module load cuda10.0/toolkit`
- `watch nvidia-smi`

Run command `exit` to end the ssh session when done.