

Using SURFsara -> Cartesius (BETA)

To get the stuff on cartesius, login into the machine and copy the directory with the exercises:

```
cp -r /home/analucia/HPC_GPU_2k17 .
```

To be able to use the GPUs on cartesius, first you will need to load the CUDA module. Use:

```
module load cuda/8.0.44
```

Next, go into the folder HPC_GPU_2k17/vector-add.
You should be able to successfully run:

```
make
```

... which will produce a "vector-add" executable.

Finally, for execution, you have the following options:

1. Use a job script like "myjob.gpu" and use the instructions from here:
<https://userinfo.surfsara.nl/systems/cartesius/usage/batch-usage#heading14>

(effectively, use sbatch myjob.gpu => the results will be written upon completion into slurm-<jobid>.out; use squeue -u <username> to check whether your job is still in the queue)

2. ssh to gcn1 - a single test node for GPU development; from there, after loading the module like above, you can simply run the code:

```
./vector-add
```

3. run a job in interactive mode, which effectively means you call srun to get a "bash" on one of the nodes, from where you use ./vector-add directly :

```
srun -n 1 -t 10 -p gpu_short --pty bash -il
```

... wait ... then you get access to a gcnXX node, like this:

```
[analucia@gcn17 vector-add]$
```

... which you have for 10 minutes. Once there, you can run the code from the prompt and get the results.