

## Working with Piz Daint supercomputer

You are granted access to the main supercomputer in the CSCS (Swiss Supercomputer Center in Lugano): Piz Daint. The credentials (a username and a password) were distributed to you by e-mail.

Piz Daint consists of two parts: a GPU-enabled cluster Cray XC50 and a multicore cluster Cray XC40 (former Piz Dora). You will only have to use the GPU part, that offers 5'320 nodes with an Intel<sup>®</sup> Xeon<sup>®</sup> E5-2690 v3 CPU and an NVIDIA<sup>®</sup> Tesla<sup>®</sup> P100 GPU.

Detailed information about the machines and how to use them can be found on CSCS web sites: `cscs.ch` and `user.cscs.ch`

**Logging in** Computational resources of CSCS are only accessible from the internal network, so you will have to first login in frontend cluster Ela:

```
$ ssh your_username@ela.cscs.ch
```

Now you can access a login node of Piz Daint:

```
your_username@ela1:~> ssh daint
```

Once on Piz Daint, you can run `hostname` to verify that you are indeed on the desired machine.

**Compilation** Piz Daint uses `module` system to control the environment and to simplify usage of different tools and libraries. In order to compile CUDA code, you will have to load several modules:

```
$ module load daint-gpu
$ module swap PrgEnv-cray PrgEnv-gnu
$ module load cudatoolkit
```

Now you can use the `nvcc` compiler and `gcc` of version 5.3.

**Job submission** All CSCS systems use the SLURM batch system for the submission, control and management of user jobs. SLURM provides a rich set of features for organizing your workload and provides an extensive array of tools for managing your resource usage, however in normal interaction with the batch system you only need to know a few basic commands:

- `srun` - submit an executable or run an executable from within the sbatch script
- `sbatch` - submit a batch script

- `squeue` - check the status of jobs on the system
- `scancel` - delete one of your jobs from the queue

From the Piz Daint login node you can submit a hostname task to the jobs queue: `srun -C gpu hostname`. Use `squeue -u your_username` to check the status of your job, or `squeue` to see your job relative to all the other jobs in the queue. After the job is executed, you will see the output of the command on the screen. Observe that the hostname is now different: you see the name of the compute node which has executed your task.

In order to have a more precise control over the parameters of your job (i.e. runtime, number of nodes or tasks submitted, output files, etc.) you can prepare a SLURM script and submit it with `sbatch` command. A common script looks like this:

```
#!/bin/bash -l
#SBATCH --job-name=job_name
#SBATCH --time=01:00:00
#SBATCH --nodes=1
#SBATCH --ntasks-per-node=1
#SBATCH --constraint=gpu
```

```
srun ./your_executable
```

See more details about job submission here:  
[https://user.cscs.ch/getting\\_started/running\\_jobs/](https://user.cscs.ch/getting_started/running_jobs/)

**Interactive mode** You can also run your jobs on compute nodes interactively. To do so, run `salloc -C gpu` and wait (typically a few minutes) until you get a shell prompt. Run `srun hostname` once again to make sure that you are on a compute node. Now you can use `srun your_application` command without having to submit it in the queue.

**Important** Note that the total node-hour budget for the students is 1600 node-hours per month, which is about **20 node-hours per month per person**. The limits are not hard, i.e. you can exceed them, but the priority of your jobs will decrease and the queue waiting time will increase significantly. Be considerate about resource usage, especially don't overuse interactive mode!

If you have any questions or issues, please **DO NOT contact the CSCS help desk**, always address to the Teaching Assistants first!