

Rocky 8 Transition Guide.

On the occasion of the arrival of the GPUs at the TeideHPC supercomputing center, a new version of the operating system used by both the computing nodes, GPU and login nodes has been introduced. The entire cluster related to Centos 6 and Centos 7 will be at the end of life in a few months.

Moreover, a new cluster called AnagaGPU has been created and some changes have been made to the TeideHPC cluster as well as the software, so if you have already run jobs in TeideHPC you will need to make adjustments to your workflow.

In summary, these are the most significant changes at the level of the Operating System, access, software, slurm:

Operating system and login nodes

- Each cluster (TeideHPC and AnagaGPU) has its own access IP.
- The operating system of both clusters and the new nodes is Rocky 8.
- There are **4 new login nodes** arranged in high availability (HA) through 2 access IPs.
- The assignment of login node during access is random and depends on the number of users.

Software

- The change in the operating system means that most user software based on Centos 6 or CentOS 7 will not work and has to be recompiled.
- The *TCL* modules tool (Centos 6) is deprecated under **Lmod**
- The installed software becomes organized using a **flat nomenclature**
- Each type of node has the specific software installed and compiled for each node architecture. This means:

The software installed depends on the architecture of the nodes

- Basically there are **2** architectures: **icelake** (nodes with GPUs) and **sandybrige** (CPU nodes).
- Look at the **cluster description** on the main page as well as the page "**How to request GPU and compute resources**".

Each cluster has its own software

To see the software available in each cluster you must enter through the access IP of each cluster.

There are modules that do not depend on the architecture

```
----- /share/easybuild/software/common/modules/all -----  
EasyBuild/4.7.0 Go/1.18.3 Miniconda3/22.11.1-1 Singularity/3.11.0 slurm/teide  
EasyBuild/4.8.2 (L,D) Mamba/4.14.0-0 Miniconda3/23.5.2-0 (D) Squashfs/4.3
```

Slurm

- **Node allocation changes from *NON-Shared Mode* to *Shared Nodes*.**

This means that by simply requesting 1 compute node, the full node is not requested for the user, thus forcing the user to make a full reservation of resources.

- The default parameters assigned by slurm are:

```
#SBATCH --node=1  
#SBATCH --ntask=1  
#SBATCH --ntask-per-node=1  
#SBATCH --cpu-per-task=1  
#SBATCH --mem=2GB
```

- A new partition has been created to **request the use of GPUs**.
- We strongly recommend using the `srun your_application` command to run applications. **In this link** You can see a simple explanation of what implications using or not using it may have.
- You can study the **efficiency of your completed jobs** with a simple command.

Public repository with examples

To facilitate the start and access on HPC computing at TeideHPC, we have created a public repository on github where we will publish examples of use of applications.

We encourage you to collaborate on it. https://github.com/hpciter/user_codes