


# Likwid tool

HPC User Support @ CINECA  
January, 14th 2021

# New SLURM plugin on Marconi cluster



- ❑ Prerequisite to run **Likwid** tool on Marconi cluster.
- ❑ Installed during latest scheduled maintenance operations on Marconi cluster on January 12<sup>th</sup>, 2021.
- ❑ This plugin allows to apply a set of configurations to the node and to ensure that these are restored at the end of the job. It is configured at present to lower the value of `perf_event_paranoid` from 1 to 0.
- ❑ It can be activated by specifying the following directives in the job script:

```
#SBATCH --gres=sysfs  
#SBATCH --exclusive
```

We have performed preliminary tests by asking for one node without the new gres and without specifying the “—exclusive” flag:


```
$ srun --partition=skl_usr_prod --nodes=1 --ntasks=48 --time=24:00:00 --account cin_staff --pty /bin/bash
```

This way we still obtain `perf_event_paranoid` related errors.

No errors are found when specifying all the required flags, and the `perf_event_paranoid` is correctly set to 0 for the current job:

```
$ srun --partition=skl_usr_prod --gres=sysfs --exclusive --nodes=1 --ntasks=48 --time=24:00:00 --account cin_staff --pty /bin/bash
```

# New SLURM plugin on Marconi-SKL



## Known issues that might be observed

*INFO: You are running LIKWID in a cpuset with 1 CPUs. Taking given IDs as logical ID in cpuset*

This is a known likwid bug that was already fixed by tool developers (a patch must be applied).

We will apply the patch and proceed with more extensive tests for this plugin in the forthcoming days.