

Minute of the meeting

Date: July 15th, 2021

Participants

EUROfusion: France Boillod-Cerneux, David Coster, Jacques David, Roman Hatzky, Serhiy Mochalskyy

CINECA: Susana Bueno, Alessandro Marani

INTEL: Giacomo Rossi

Subjects discussed

We have reviewed events announced to all users via hpc-newsletter that have affected production since June 16 until July 14.

Ticket revision

We have reviewed the two tickets escalated to SchedMD:

- ticket 16506: reports the issue referring to the allocation of 1 GPU card job to nodes in random way and not packing tightly on Marconi100 cluster. SchedMD support reported that the bugfix requires a notable change in the select plugin, so they expect the fix to be available only in slurm 21.08 (August 2021). We agreed to wait for the slurm upgrade.
- ticket 18895: reports the issue referring to a second mpirun in a job that fails after slurm upgrade; the fix for this bug is expected on 20.11.8 version.

We have also reported and reviewed the status for the two tickets escalated to Intel support:

- ticket 3932: Intel support is working on this bug (CMPLRIL0-33599) and recently provided feedback and possible resolution together with a new test code that we reported to the user.
- ticket 16019: we obtained information about the motherboard configuration (UPI links) for Marconi SKL compute nodes; we also prepared and shown a report containing also results obtained for latest tests performed using a different tool for intra-node bandwidth evaluation as suggested by the Intel MPI team.

We reported other tickets on the second level support queue. We focused mainly on ticket 19157, that show a known issue randomly occurring for jobs running Euterpe code that run up to the time limit with no output generation. Differently from previous known cases where users that were using an old executable could fix the issue by recompiling the code, the wrong behaviour seems to show this time when using an executable two days old. We will discuss internally about a possible strategy so to try to recognize the possible bug involved on these jobs.

Ticket statistics

We have reported and reviewed information collected for all tickets receive, on both the first and second level support queues, since June 14 until July 12.

Module usage on Marconi-SKL

We reported data collected for June 2021 (referring to new 5th Cycle Allocation projects that started on March 1st) for the modules usage of those batch jobs submitted by EUROfusion users/accounts to Marconi-SKL partition, that was integrated with information about module usage by partition/qos.

Sanity checks on Marconi and Marconi100 clusters

Sanity checks were performed between June and July on Marconi and Marconi100, both during maintenance and during production. The results for production checks were incomplete because the data collection is still in progress. The results shown that there are no particular anomalies and both clusters are stable in their performances.

likwid & hpcmd tools

Latest tests performed with likwid tool installed as module on Marconi cluster show a failure in the main memory data collection, this data is missing (it is reported a zero value). We have done additional tests that confirm these observed values. We have contacted the developer of this tool and we have provided him access to both Marconi and Marconi100 clusters where the likwid tool is available. He gave us some relevant feedback, we are doing now further investigations with our system administrators (exposure of offcore registers by the operating system).

We presented progress done on the several steps required to set up the whole pipeline to make properly work the hpcmd tool as indicated in the public documentation (<http://mpcdf.pages.mpcdf.de/hpcmd/>). We have reported and described a new workflow that has been put in place on Marconi cluster, that manages to collect messages generated by the hpcmd daemon on each Marconi compute node (skl_sys_test partition) into a shared GPFS area; the collected data is then made available onto a second area, that is mounted on the virtual machine on our CLOUD infrastructure. We are working at present on the setup of the procedures required for the data ingestion into Elasticsearch database, also in the setup of the required dashboards for data visualization.

Next ticket meetings

The scheduled dates for next ticket meetings are:

TM-56: Tuesday, September 14th, 2021 at 10:00

TM-57: Friday, October 15th, 2021 at 10:00

TM-58: Monday, November 22nd, 2021 at 10:00