

Minute of the meeting

Date: July 4th, 2022

Participants

EUROfusion: France Boillod-Cerneux, David Coster, Gilles Fourestey, Roman Hatsky, Serhiy Mochalskyy, Michal Konrad Owskiak

CINECA: Susana Bueno Minguez, Alessandro Marani, Nitin Shukla

Subjects discussed

We have reviewed events announced to all users via hpc-newsletter that have affected production from June 9 until June 30.

Ticket revision

We have reviewed tickets escalated to NVIDIA support. Other tickets were briefly reviewed and discussed, with no particular observations raised. It was suggested to close ticket #24227 due to lack of user feedback, and it was remarked the importance of having resolved ticket #25494 about upgrading the software stack on Marconi100.

Ticket statistics

We have reported and reviewed information collected for all tickets received, on both the first and second level support queues, from June 9 until June 29.

Module usage on Marconi cluster

We couldn't provide data for the month of June, but it will be forwarded to all participants within a week; it will refer to the module usage of those batch jobs submitted by EUROfusion users/accounts to Marconi-SKL partition.

Sanity checks on Marconi and Marconi100 clusters

Sanity checks were performed during production and maintenance on MARCONI, and during production on MARCONI100. The results were shown during the Ticket Meeting. All the tests proved that the situation is stable and there are no significant problems in any of the tests performed, apart for some slightly slow performance on some nodes in both clusters, that are kept under check.

Power capping on Marconi

No updates on this topic, but the system administrators of CINECA are gathering additional data to improve their procedures that detect and uncap nodes with power capping on.

Starwall benchmarks

We present the STARWALL benchmark and its findings. The STARWALL benchmark was performed during maintenance on the 14th of June. We compile and run the binary using Intel MPI 2020. The program uses 16 Nodes, we submitted 175 jobs occupying 2800 nodes. We obtained the Min time of 18.31 sec and the Max time was 32.35 sec. Results compelled us to investigate discrepancies in the time of some runs. We have performed further runs with Intel MPI 2018 and 2021, which performed similarly and showed better performance than Intel MPI 2020. We launched 140 runs and obtained min time = 17.3 and max time = 25.27 sec. We concluded that it could be a particular problem with Intel MPI 2020. We need to reach out to the Intel team. We also compared the results of what happened when it performed better and worst. We used COUNTDOWN, "A Run-time Library for Application-agnostic Energy Saving in MPI Communication Primitives" and illustrated the difference between these two runs. EUROfusion asks to perform an additional run on the same switch.

Hpcmd tool

We have given updates about the last actions performed. The preliminary plan agreed (that foresees in a first phase the installation of the hpcmd tool in the 16 nodes available on the skl_fua_dbg partition, and, in a second phase, to extend to a total of 128 nodes the nodes available on this partition for short period of time of about 2 days, so to allow extensive tests on a bigger number of nodes) is ongoing. We have now completed the setup of the tool on one production node on the skl_fua_dbg partition and adjusted the configuration for the database and web site. Nexts step is to install the tool in all skl_fua_dbg nodes, we will inform all members of the TM team as soon as this will be done.

Next ticket meetings

The scheduled dates for next ticket meetings are:

TM-67: Tuesday, September 13th, 2022 at 10:00

TM-68: Monday, October 10th, 2022 at 10:00

TM-69: Monday, November 7th, 2022 at 10:00