

Minutes of the ABP Computing Working Group meeting 1

29th September 2016

Participants: X. Buffat, J.M. Coello de Portugal, R. Costa, L. Deniau, R. De Maria, G. Iadarola, J. Iven, B. Jones, K. Li, E. McIntosh, L. Mether, K.S. Sjobaek.

R. De Maria presented a summary of the use cases of AFS within the ABP group, along with the typical work flow within which AFS is used. He highlighted several issues that he will be encountered in case no replacement fulfilling the same requirements is not provided. He insisted on the required manpower which seems to exceed the resources currently allocated, since the alternatives seem to require a major refactoring for many of the codes that are currently run on lxplus/lxbatch as well as a complete change of work flow for people heavily working remotely on lxplus.

He mentioned that the current documentation is very poor and several technologies are suggested for different applications. This results in significant cost in assesment for the user. Since the roadmap seems uncertain, it is unclear whether the efforts invested will pay off.

Without attempting to use HTCondor, it seems that this technology could fulfil the same needs as LSF.

J. Iven showed how they identified a decline of the OpenAFS community. Most institutes are moving towards other technologies. The workload is on purpose put on the users, which will have to adapt to the new technology choices. This was agreed at the sector and department levels, it is unclear why the information did not reach the group. OpenAFS is not an optimal filesystem for several application for which it is used, this major refactoring is hoped to bring optimised solutions for the different use cases. Fortunately the migration may take a few years, as the OpenAFS project does not have a fixed end date. IT will support AFS as long as possible for dedicated critical applications only.

The typical workflow presented by R. De Maria should not be affected by a migration to EOS, yet it would not be efficient. The proposed strategy is to move locally data required at run time.

J. Iven mentioned that AFS has been kept on the technical network for convenience, however it represents a security threat that should have been removed long ago.

Running shared applications remotely (e.g. MAD-X, Mathematica), as currently available on LXPLUS, would require the setting up of a dedicated cluster with its own shared file system. Nevertheless, the executables could still be available on CVMFS or CERNBOX, but will need to be executed locally.

H. Burkhardt is the contact person for important requests to IT. For support or give feedback, one may use tickets.

LX BATCH is running currently as its maximum capacity in terms of number of jobs, it does not fulfil the needs for the treatment of the LHC data. B. Jones presented the migration to its replacement : HTCondor. IT will end support for LSF at the end of run II. While based on slightly different concepts (no queues, no normalised minutes), HTCondor has all the features provided by LSF and possibly more. It profits from a larger community.

A pilot version is already running, it is possible to use it. B. Jones will provide the working group with the

documentation.

The possibility to operate a dedicated cluster for MPI type of jobs is being investigated. If realised, it will be accessible through HTCondor.