

Minutes of the ABP Computing Working Group meeting

24 May 2018

Participants: D. Amorim, F. Asvesta, D. Castro, R. Costa, R. Dembria, L. Deniau, J. Evans, G. Iadarola, A. Latina, A. Mereghetti, M. Migliorati, L. Mether, J. Molson, A. Oeftiger, V. Olsen, D. Piparo, G. Rumolo, R. Scrivens, L. Sabato, G. Sterbini, E. Tejedor.

SWAN: service for web based analysis

Our IT colleagues illustrated the capabilities of the SWAN service:

- The idea of SWAN is to provide a tool to perform analysis directly in a web-browser profiting of cloud computing resources.
- The chosen interface is the jupyter notebook. The storage is made on the personal EOS space of the user.
- The software is distributed via CVMFS (CERN distributed system for storing software) and the user is given the possibility of installing additional libraries (in the local storage).
- The available interfaces are Python, Root, R and Octave.
- Recently added features include the possibility of having shared projects among users and the integration with the Spark big-data infrastructure.
- The usage of SWAN requires a CERN account, this makes it non-trivial to use for schools and trainings involving non-cern users.
- An annoying feature is that sessions are killed after 6h of inactivity. The users would very much like the possibility of freezing/reloadings the state. A workaround consist in leaving the browser tab open.
- It is possible to install SWAN locally.
- At the moment IT maintains the infrastructure and EP is making and maintaining the software.
- A pilot program is starting to provide GPU computing via SWAN.
- A SWAN interface to submit jobs to HTCondor would be very useful for many ABP use-cases.
- The development team found extremely useful to expose the project on GitHub, as this attracted many contributions from volunteer developers.