

# Minutes of the HEPiX Benchmarking Working Group Meeting 2016-04-21

Participants:

- Local: Manfred Alef, Martin Bly, Chris Hollowell, Liviu Valsan, Peter Wegner
- Remote: Jerome Pansanel, David Abdurachnarov, Aritz Brosa, Emmanouil Vamvakopoulos, Domenico Giordano, Marco Guerri (and DESY staff)

## Welcome to the HEPiX Benchmarking Working Group!

(Note: This first face-to-face meeting is just an informal one. The official kickoff meeting will be announced asap when some remaining administrative issues (mailing list, wiki, Indico, Vidy, ...) have been fixed.)

## First tasks of the working group:

### 1. Fast benchmarks

There are several candidates of fast benchmarks:

- a) Dirac fast benchmark (LHCb)
- b) Geant4 / single muon events (default workload of Atlas KitValidation benchmark)
- c) ROOT stress test (Alice)
- d) Whetstone
- e) Dhystone
- f) CMS fast benchmark?

It has been suggested to run also the Geant4 full chain benchmark beside of b). D.G. suggests to ignore the first event to remove setup overhead from the benchmark score.

M.A. has shown results from running these 5 benchmarks as single-core jobs in a batch farm [1]. It is important to check cloud environments and multicore jobs as well:

- Site representatives please run the fast benchmarks together with the HS06.
- Experiment representatives please report the performance of their applications (events/s) as well as the fast benchmark scores.

**Benchmark harness:** M.A. proposes to use D.G.'s KV tool to run the benchmarks and to collect the results at a single place. D.G. will present the tool in the next meeting. Further suggestions: report the walltime of all benchmarks, not only the performance scores, and also the memory layout of the system under test.

### 2. Scaling of HEP applications with HS06:

Alice and CMS are invited to do similar analysis as was done by Philippe Charpentier (LHCb), see GDB's 2015-09-09 and 2015-12-09, and by M.A. on Atlas jobs (see WLCG Workshop). If Alice and CMS provide internal accounting details (events/s, WN hostname@GridKa) of representative jobs that run at KIT then M.A. can provide the corresponding HS06 scores.

Knowledge of the correlation of WLCG applications with HS06 will be very helpful for deployment of the next generation HEP-SPEC benchmark.

**3. Deeper investigations of the magic boost** of LHCb jobs, Atlas simulation jobs, and Dirac fast benchmark on Intel Haswell platform compared with the former Sandy Bridge generation, possibly caused by either AVX2 or fast hardware random number generator.

**Miscellaneous:** Liviu has announced that he moved away from benchmarking because of re-organisations at CERN.

[1] <https://indico.cern.ch/event/466991/contributions/1143617/attachments/1260502/1862650/CPU-Benchmarking-2016-04-20.pdf>