

Hepix Benchmark WG
Draft minutes on 17 June 2016
(e.v. took notes)

- **indico angeda** : <https://indico.cern.ch/event/540544/>

Costin Grigoras presented the « ALICE results on fast CPU benchmarks »

- **70%** of ALICE calculations correspond to Monte Carlo Simulation Jobs. The fast benchmark should scale with this activity.
- First tests with RootStress and Condor_kflops, both of benchmarks meet the fast execution requirement (e.g. 1-2 minutes)
- Results demonstrate that the RootStress scale rather well with Alice's MC workload. In contrast, there is not any correlation with Condor_kflops results and Alice's MC workload.
- Next tests took place with SysBench. SysBench did not exhibit better scaling performance with respect to RootStress results.
- Alice tried a commercial benchmark : CeekBench
 - get evaluation licenses to run the benchmark on the grid on 64b.
 - The CeekBench gives good estimation of MC performance job
- LHCb python benchmark is running for long time in production and is giving well correlation with MC simulation job's performance of ALICE experiment.
- HS06 vs MJF exhibits large variation but this correspond a different value of HS06 on MJF for HS06 values (RAL Case, reported by P.C.)
- **The CPU performance and the correlation with the enabled or disabled state of Hyper-threading capability of the CPUs , it is not obvious from those results.**
- Plan to run the fast benchmark with the pilot and keep history for each WN, if is possible.

Manfred Alef presented the **Fast Benchmark Candidates**:

- Use cases, requirements and the running condition of the fast benchmark were presented at KIT.
- All the fast benchmarks exhibit rather well scaling with experiment workload but also exhibit difference with HS06 (up to 40% in some cases).
- Comment and Proposition from M.A.:
 - **Make standard the name of the fast-benchmark candidates** (e.g. Dirac fast benchmark, KV single muon simulation, root stress test, Drystone and Whetstone).
 - **Can we agree on frozen benchmark version and Compiler flags ?**
 - **Dirac Fast Benchmark**: run with default Python
 - **KV Single Muon Simulation**: need more transparent compiler version and flags
 - **ROOT Stress Test**: very tricky to compile; use binary which is used at GridKa?
 - **Drystone, Whetstone**: use default Makefile
 - **Attention**, the calibration factor in Dirac benchmark on single core version and multi-core version is not the same (250 vs 360).
 - (P.C.) The reason for new value (360) is that describe better the AMD CPUs, it is not related with running mode (serial vs parallel).
- **Benchmark harness**
 - CERN Cloud Benchmark Suite (see talk by Domenico Giordano, see at: https://indico.cern.ch/event/535458/contributions/2176092/attachments/1284582/1909948/CERNCloudBenchmarkSuite_HEPIXBmkWG_giordano.pdf, <http://bmkwg.web.cern.ch>)
 - CERN Cloud Benchmark Suite can runs several types of benchmarks (**KV , Dirac and**

Whetstone)

- Collect results at single place for final analysis:
 - (D.G.) We need more feedback from benchmark users in order to check the option of global collection of benchmarks results from different sites)
- **CERN Cloud Benchmark Suite have strong binding from cvmfs** (KV release is stored on CVMFS). Therefore in a fresh installed machine (vm and or physical) need root access to setup properly the fuse module (plus some rpms).
- Create a repository with all available benchmarks (?)
- **Next meeting on 1 July 2016**