Minutes of the HEPiX Benchmarking Working Group Meeting 2016-07-01

Participants: Manfred Alef (minutes), Omar Awile, Martin Bly, Cristovao Cordeiro, Domenico Giordano, Costin Grigoraş, Chris Hollowell, Michele Michelotto, Emmanouil Vamvakopoulos

Presentation by Manfred Alef: Next Steps

In the meeting 4 weeks ago Domenico Giordano had presented the CERN Cloud Benchmark Suite. First tests have shown that this suite is a helpful tool to run several fast benchmark candidates and to collect the results at a single place for later analysis. On the other hand there are some weak points which prevent running form traditional batch jobs:

Dependencies: batch systems don't grant root permission to user jobs, therefore the tool must run in user space, and required packages should be added to the HEP_OSLIBS meta package to make them available on all WLCG WNs, so that benchmark jobs can run on a wide range of differently configured hosts (HT, turbo boost, number of job slots, ...).

Licensing: all proposed fast benchmarks are coming with free licenses, except some Atlas extensions of the KV tool. Therefore KV will not ship with the CERN Cloud Benchmark Suite but read KV from CVMFS. Open question is how to run it at sites which are not supporting Atlas.

Parallel runs: the suite can run once per user account. It creates a lockfile which prevents to start a second benchmark instance by the same user while the first instance is still running. This will be fixed.

Users will need special credentials for sending results to the ES data collector. Please contact Cristovao Cordeiro or Domenico Giordano.

Regarding the talk by Manfred Alef in the last meeting 2 weeks ago Emmanouil Vamvakopoulos asks the question about the wide spread of benchmark scores on Intel hosts while the AMD numbers are quite flat. Answer: the HS06 measures the system performance under worst-case conditions while as many benchmark copies are running in parallel as there are job slots configured for the system under test. When the WN is only partially loaded then it can run at higher clock speed and provides some amount of bonus performance for free. The flat AMD results are caused by the fact that clock speed is fixed at AMD hosts at GridKa because of some critical issues in the past.

Indico link:

https://indico.cern.ch/event/544935/

Next meeting:

Fri June 15th 14:00 CEST (12:00 UTC): https://indico.cern.ch/event/549169/