

Benchmarking worker nodes using LHCb simulation productions and comparing with HEP-Spec06

Thursday, October 13, 2016 3:15 PM (15 minutes)

In order to estimate the capabilities of a Computing slot with limited processing time, it is necessary to know with a rather good precision its “power”. This allows for example pilot job to match a task for which the required CPU work is known, or to define the number of events to be processed knowing the CPU work per event. Otherwise one always has the risk that the task is aborted because it exceeds the CPU capabilities of the resource. It also allows a better accounting of the consumed resources.

The traditional way the CPU power is estimated in WLCG since 2007 is using the HEP-Spec06 benchmark (HS06) suite that was verified at the time to scale properly with a set of typical HEP applications. However the hardware architecture of processors has evolved, all WLCG experiments moved to using 64-bit applications and use different compilation flags from those advertised for running HS06. It is therefore interesting to check the scaling of HS06 with the HEP applications.

To this purpose, we have been using CPU intensive massive simulation productions from the LHCb experiment and compared their event throughput to the HS06 rating of the worker nodes. We also compared it with a much faster benchmark script that is used by the DIRAC framework used by LHCb for evaluating at run time the performance of the worker nodes.

This contribution reports on the finding of these comparisons: the main observation is that the scaling with HS06 is no longer fulfilled, while the fast benchmarks have a better scaling but are less precise. One can also clearly see that some hardware or software features when enabled on the worker nodes may enhance their performance beyond expectation from either benchmark, depending on external factors.

Tertiary Keyword (Optional)

Secondary Keyword (Optional)

Accounting and information

Primary Keyword (Mandatory)

Data processing workflows and frameworks/pipelines

Primary author: CHARPENTIER, Philippe (CERN)

Presenter: CHARPENTIER, Philippe (CERN)

Session Classification: Track 6: Infrastructures

Track Classification: Track 6: Infrastructures