

## Survey & Discussion

- Alessandra Forti (Atlas experiment and site repres.)
- Andrew McNab (LHCb and site repres.)
- Manfred Aef (WG chair and site repres.)
- Pepe Flix (CMS experiment and site repres.)
- Latchezar Betev & Costin Grigoras (ALICE experiment repres.)

# Survey

- Collects a number of topics discussed during the WG meetings
  - Agreed with the panelists
  - Experiment Representatives collected answers
  - In addition, point of view of few site representatives

# About HS06...

# (1/13) HS06 score Vs HEP simulation workloads

Is the issue of HS06 score Vs time Simulation workload confirmed, within the accuracy you need?

ALICE	<i>HS06 does not scale with our simulation workload, DB12 does</i>
LHCb	<i>Not really. We could increase the safety margins in the events estimation.</i>
ATLAS	<i>There has been an increase of jobs killed by the batch system, but it is not clear if the problem lies with hs06 use in the brokering. Needs more digging</i>
CMS	<i>For the moment, the results obtained in these studies are pretty compatible (to some extend) to a linear correlation.</i>

# (2/13) HS06 score Vs HEP reconstruction workloads

Is the correlation still good for reconstruction jobs?

ALICE	<i>HS06 does not scale with reco jobs either. Since CPU intensive jobs are the majority of what we do, the fast benchmark is adequate to gauge the machine performance.</i>
LHCb	<i>Better than for MC yes. We don't rely on it there though.</i>
ATLAS	<i>Needs more digging</i>
CMS	<i>We don't have made detailed studies in this yet.</i>

## (3/13) HS06 score Vs HEP workloads

How was it studied in the recent years?

ALICE	<i>Fast benchmarks: studied and validated on the Grid and on isolated machines with the same result. We have now covered ~all CPU types in any configuration available. HS06 have been taken from machines where they were available.</i>
LHCb	<i>Job Slots</i>
CMS	<i>Since March 2017, studies made in PIC Tier-1, with isolated machines</i>
ATLAS	<i>A comprehensive study hasn't been done.</i>

## (4/13) HS06 expectation if compiled at 64 bit

Belief/Evidence that HS06 would scale better if compiled at 64 bits?

ALICE	<i>We don't know and what would be the difference?</i>
ATLAS	<i>It would be closer to the application setup. However to do such a change it would have to last for a reasonable period. We shouldn't change for the sake of change.</i>
CMS	<i>We didn't test this yet.</i>
LHCb	<i>Not sure. In principle yes. Disruptive to do though.</i>

# About Fast Benchmarks...



## (5/13) Benchmarking information in job slots

Does the experiment need to access benchmarking information in the job slot? For which purpose?

LHCb	<p>Yes.</p> <p><i>It's used to calculate how many events to try to simulate in MC jobs.</i></p>
ALICE	<p><i>Not a 'must have' to have in MJF. If there is an approved benchmark (for example DB12) and its results are published on the WNs, then it dispenses with the need to run fast benchmark with every job.</i></p>
ATLAS	<p><i>While short benchmarks have marginally better results than HS06, for most applications this is not enough to warrant running them in each pilot.</i></p>
CMS	<p><i>We don't have at the moment "job masonry" as part of our model. [...] Jobs are of pre-defined length, job to slot matching right now based on such pre-defined job length</i></p>

# (6/13) Adoption of fast benchmarks in pilot framework

What is the state of the art?

LHCb	<i>Fully adopted. Every MC job relies on it. Every pilot job calculates and records it</i>
ALICE	<i>DB12 is fully integrated in ALICE JobAgent, ready for production. We can switch to other benchmarking as needed</i>
ATLAS	<i>Currently the pilot runs DB12 and whetstone in the pilot and the numbers are then written in a ES instance in UC.</i>
CMS	<i>We are running DB12 already, other fast benchmarks could be easily added as well. Results are injected into the CERN ELK instance</i>

## (7/13) Preferred fast benchmarks

What are the preferred fast benchmarks from the experiment point of view?

LHCb	<i>DB12 Python is what we have in production and is sufficient for our current needs.</i>
ALICE	<i>We have evaluated all current contenders: DB12 python, DB12 C++, HepSPEC, GeekBench and others. Only DB12 python shows correlation with our MC workload.</i>
ATLAS	<i>Currently still DB12 python and whetstone. DB12 cpp may be used in the future.</i>
CMS	<i>So far, DB12 (python), KV compared with HS06 and a special CMS ttbar simulation suite and the results obtained are not yet conclusive.</i>

# Preparation for the next long-running benchmark

## (8/13) Suite of HEP workloads - candles

- What are the suggested workloads to use?

ALICE	<i>Only for MC (2/3 of the Grid activity, CPU-bound jobs).</i>
LHCb	<i>MC as this is vast majority of our CPU use.</i>
ATLAS	<i>It is important to measure the correlation with all the applications. (Full simulation is ~37% of the ATLAS , evgen is 22%)</i>
CMS	<i>TTbar simulation multi-threaded suite, [...] pretty stable when running &gt;500 events (and it typically takes 1h of execution time). We should discuss how to provide a RECO workload for testing.</i>

## (9/13) Suite of HEP workloads - candles

Action that the experiments can take to make available such workloads

ATLAS	<i>Work already ongoing, workloads exposed</i>
CMS	<i>No problem!</i>
ALICE	<i>The code itself is available in CVMFS, together with the scripts to run it, however there is some work to be done on making it ALICE environment agnostic. We will work on this.</i>
LHCb	<i>We expect to be able to containerise this yes. We don't have an estimate for when it will be available.</i>

# (10/13) Suite of HEP workloads: multi-threading?

What is the status of adoption of multi-threading

ALICE	<i>No immediate plans for multithreaded jobs.</i>
LHCb	<i>This is an active area of development and imminent deployment in preparation for the LHCb upgrade. It is definitely something we will rely on.</i>
ATLAS	<i>Under validation to move G4 to multi-threaded running by 2018. For reconstruction this is targeted for Run 3, unlikely to swing into action before 2020. IMO now is the right time to add a multi-threaded benchmark for sites - we could even offer ATLAS G4 MT if that's interesting.</i>
CMS	<i>This is particularly important for CMS, since all main CMS applications became multi-threaded in 2016. [...] benchmarks will need to reflect that.</i>

# (11/13) New CPU architectures under study

what is the set of new architectures where the WLCG workloads will run (that then need to be benchmarked?)

LHCb	We expect to continue purely with x86 (unless sites go in a different direction) but SSE4 and AVX2 optimised code is already being studied in detail for the performance benefits so it would be useful to be able to reflect that in the benchmarking in the future
ALICE	<i>largely x86</i>
ATLAS	Aarch64, Power8,9
CMS	



# (12/13) Status of adoption of GPUs

what is the set of new architectures where the WLCG workloads will run (that then need to be benchmarked?)

ATLAS	<i>No GPU for now</i>
LHCb	<i>Some institutes have developed analyses which make heavy use of GPUs (e.g. for multidimensional fitting). However they're not used by the centrally managed offline system and we don't have plans to use them outside user jobs.</i>
ALICE	<i>GPUs are used in specific environment and for specific code: HLT and big role in the O2 system. The code executed currently and in the future on the Grid will not depend on GPUs</i>
CMS	<i>The project of GPUs for pixel tracks at HLT is growing fast.</i>

# (13/13) Suite of HEP workloads – Collection of results

How shall we collect results? Is there a need of a common DB for hardware models?  
N.B. This is not for the accounting use case, but to consolidate the approaches to run, collect and share the benchmark results.

ALICE	<i>We are ready to provide the DB12 score and the observed simulation performance for all CPUs and nodes on the Grid where we have run the benchmark. ALICE will use the gathered results internally and for accounting purposes. <u>No plans beyond the current exercise.</u></i>
LHCb	<i>We don't feel a central DB is needed by LHCb operationally. It would be interesting for performance studies though.</i>
ATLAS	<i>From the WLCG point of view it would be ideal to have all the results together. However when we tried to write in their infrastructure from the pilot we couldn't do it because it is not certificate aware.</i>
CMS	<i>Results could be published to CERN ELK [...] A how-to-run page for running and publishing results could be welcome.</i>

