

Benchmarking – The GDB Discussion and the Way Forward

Helge Meinhard, CERN-IT
WLCG Management Board
15-Sep-2015

Outline

- GDB discussion – summary and personal impressions
- Four issues to address
 - Predict power of batch slot
 - Benchmark accurately full machine
 - Check accounting figures
 - Decide on mechanisms for storing and distributing benchmarking information
- How to move forward – a proposal

GDB Discussion 09-Sep-2015

- https://twiki.cern.ch/twiki/bin/view/LCG/GDBMeetingNotes20150909#CPU_Benchmarking
- Machine-job features
- Benchmarking: experiment contributions
 - Widely varying level of detail
 - Long discussions highlighting uncomfortable feelings
 - A number of areas to be improved
 - IMO a little confused
- Attempt to structure... (Manfred Alef, Ian Bird, Michel Jouvin, HM, ...)

Issues to Address (1)

- Predict power of compute slot (batch, cloud) for the running job
 - Often needed for job matching and masonry
 - Two approaches:
 - Use HS06 (via MJF) – possibly underestimate because of advanced CPU features
 - Determine on the running job – possibly unprecise if workload changes
 - Needs to be fast and to require access to job slot only
 - HS06 clearly inappropriate – takes hours, requires licence, expects access to full machine
 - Known candidates: LHCb Python script, ROOTmark, Drystone/Whetstone, KitValidation, HTCondor benchmarks, ...
 - Some work done, but not conclusive yet
 - Ideally one single choice for all experiments and application types

Issues to Address (2)

- Benchmark precisely a whole farm
 - Needed for procurements, pledges, installed capacity, CPU accounting, ...
 - Requires (possibly long-running) benchmark programs controlling the full machine
 - HEP-SPEC 06 emerged from common WLCG/HEPiX activity back in 2007/2008
 - No known issue with HS06 per se
 - No evidence of scaling issues beyond 10% - the initial objective
 - Choice of boundary conditions for running HS06 has served community well
 - Applications, machines and industry-standard benchmarks have evolved since
 - Should move forward to a new benchmark soon, following proper verification against typical experiment applications
 - Candidates: (Subset of?) new SPEC CPU benchmark suite (expected to be released soon), Geant4 benchmark, ...

Issues to Address (3)

- Accounting: Acute or latent suspicion about accounting numbers being inaccurate
 - CPU time used times slot power in HS06
 - HS06 of machine divided by number of slots, or
 - Average HS06 per slot of a whole compute farm or CE
 - Increasingly inaccurate due to increased machine sophistication – factors including
 - Symmetric multi-threading / hyper-threading
 - Turbo-boost
 - Virtualisation
 - ...
 - Exactly the same reasons why (1) and (2) are potentially very different!
- Check whether this is the only source of discrepancy (and unhappiness)

Issues to Address (4)

- Storage and transport of information
 - Current attempt: Machine-job features
 - Definitely the right direction
 - Deployment is easy (still risks to take long due to chicken-egg situation)
 - Needed at least for precise estimate of lower limit of job slot performance, and for proper per-job accounting

Proposal

- Quite some expertise around
 - ... and even way more diverging views ;-)
- Co-ordination and planning needed
- Establish a task force mandated to plan concrete steps to tackle issues (1) to (4)
 - Include all LHC experiments, selected site reps and benchmarking and accounting experts
 - Report back to MB (and GDB)
 - Subject to MB approval, kick off activities around issues (1) to (4) with clearly defined objectives and target dates
 - In particular for (1) and (2), collaborate with HEPiX and their benchmarking experts

Questions? Decision?