

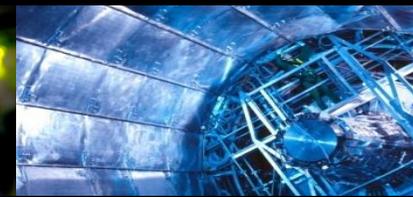
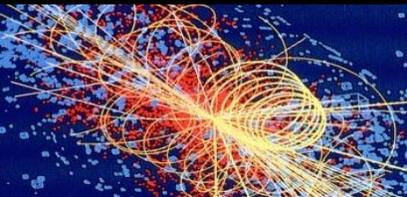
System performance and cost model working group

Josep Flix

Andrea Sciabà

Markus Schulz

WLCG MB, 12-09-2017



Status

- June WLCG workshop had a session on efficiency and cost
- Highlights
 - Need a huge increase in “efficiency” of computing to be able to do computing at HL-LHC
 - How to measure efficiency? What metrics are relevant?
 - Technology will evolve in the next ten years. What will the impact be? Which scenarios should be considered?
 - How to translate the needs of the physics programme into resource requirements?
 - How to quantify the advantage (or disadvantage) of cloud resources vs grid resources (and keep it up to date)?
 - Conclusions
 - A “cost” model, intended as a tool to calculate the impact of current or future computing models on resource needs, would be extremely beneficial for WLCG
 - All experiments have something of the sort to calculate their computing requirements
 - A common approach would be highly desirable
 - Focus not on the money, but on resources and differentiate their needed characteristics (I/O, network, memory etc.)
 - Broad consensus on the proposal to organise a preGDB/GDB and create a working group to discuss and build a system performance model
 - “System” is WLCG in this context
 - Mapping results to cost can be done locally as the last step

Volunteers and Nominated Participants

- In order of appearance:
- Andrea Sciaba, Josep Flix, Markus Schulz
- Helge Meinhard (CERN)
- Gareth Roy (ScotGrid)
- Andrew Sansum (GridPP, AENEAS, SKA)
- Andrey Kiryanov (PIK Computing Centre, Russian Federated Data Storage Project)
- Catherine Biscarat LCG France (+more)
- Domenico Giordano (Benchmarking working group)
- Davide Costanzo and Johannes Elmsheuser (ATLAS)
- Daniele Bonacorsi, David Lange (CMS)
- Concezio Bozzi (LHCb)

Next steps

- Call for volunteers
 - Need to reach a critical mass in terms of interest and availability to contribute
 - Already close
- Created mailing list:
 - wlcg-SystemPerformanceModeling@cern.ch
- Will contact current volunteers to organise a first video meeting after HEPIX
 - Work on a well defined mandate
 - To define objectives and timescales for the first six months
- For feedback, write to
 - jflix@pic.es
 - Andrea.Sciaba@cern.ch
 - Markus.Schulz@cern.ch
 - wlcg-SystemPerformanceModeling@cern.ch

Draft of a potential Mandate

- WLCG Resource and Cost Model Working Group
- WLCG needs metrics that allow us to characterise the resource usages of HEP workloads in sufficient detail so that the impact of changes in the infrastructure or the workload implementations can be quantified with a precision sufficient to guide design decisions towards higher efficiencies. This model has to express the resource utilisation of the HEP workloads in terms of fundamental capabilities and characteristics that computing systems provide, such as storage, memory, network, computational operations, latency, bandwidths etc. . To allow sites and user communities to use the model to improve also their cost efficiency an approach to map these capabilities to local costs is highly desirable. This can't be achieved at a global level, since the conditions at different sites are too different, but the model should be constructed in such a way that this mapping on a local level can be done fairly easily, following given examples.

Mandate

- Bring together workload and infrastructure experts (sites and experiments) to agree on common suitable metrics.
- Identify a set of reference workloads and meter them in different environments as input data for the model.
- Build the model and verify it by predicting resource usage of the reference workloads with respect to changes in the execution environment.
- Provide a few examples of mapping the model to local cost.
 - Covering a small and a medium sized site
- These steps are not sequential and the metrics will have to be modified to be usable.
- Due to the partial overlap with the Benchmarking WG close collaboration is needed.
- Timeline: within 3-6 months a first version of metrics and model should be presented to the community. Further progress should be tracked by the GDB.