

WLCG

Performance Discussion

Slides to provoke a discussion

Markus Schulz (slides presented by Andrea Valassi)

GDB, 9 November 2016

Partitioning???

- **WLCG Deployment Area**
 - Workflows
 - Sites managers + experiments operations people
 - Data Management
 - Sites/data management systems, experiments operations
 - Networking
 - Infrastructures
- **WLCG Application Area/HSF**
 - Software
 - Experiments, ROOT, GEANT,.....
- How to link these activities?
 - Or keep them apart?

Do we know what to optimise for?

- Efficiency and performance are very fluid concepts
- Everyone agrees that these are less than perfect quantities:
 - CPU/Wall
 - Time to completion
 -
- Can we get meaningful quantities derived from workloads and the costs of resources?
 - ALICE did a good job for the HLT deciding between:
 - m_1 nodes with CPU_x at n_1 chf /node
 - m_2 nodes with CPU_y + m_3 GPUs at n_2 chf/node

Cost Model Working Group?

- Can we identify a reasonably small (<10) number of “typical” workflows?
 - Standard Candles
 - That can be run by site managers independently from the experiment frameworks
 - Requires common taxonomy of jobs
- With different I/O patterns
 - Staging to WN
 - Local storage system read
 - Remote read
- Single threaded vs. MP/MT
- This exists already in most experiments
 - With different degrees of “ability to be run independently”
 - Currently difficult to compare
- Plus a model for the infrastructure costs
 - Required to compare with commercial cloud resources ;-)
 - Some interesting presentations at CHEP
- Add everything up (weighted by the workflows frequency) → event/chf

Starting with something (relatively) simple:

- Shared taxonomy for workloads
 - mapped to groups of similar experiment specific detailed job types
- Should be made accessible to the site managers directly
 - Not via harvesting workload management system log files.
- Will make it easier for sites to understand what kind of job type a running job is representing
 - Allows sites to compare behaviour
 - Allows to spot mismatches between expected and observed behaviour
 - Allows sites to give concrete feedback on changes, drifts, impact on the infrastructure etc.

Is there a need for a forum to discuss performance related activities at regular intervals?

- Software, Infrastructure and Workflow related
 - Tools, observations, plans, open issues
 - Bringing software and operations people together
- Could be linked to the WLCG-Workshop, Hepix or the HSF Workshop
 - With focus on only a few topics each time
 - Covering over time developments in all areas
 - Not longer than ½ -1 day with opportunities for people who work on related subjects to discuss in small groups issues in depth
 - Not a day of presentations

Executive Summary

- The LHCC and the Funding Agencies asked to start a process to define and address the cost of computing for HL-LHC
- We (WLCG) started this process and what I just summarized are the first steps forward
- New Alice and LHCb computing TDRs target Run-3
- The plan is to produce ATLAS and CMS TDRs on the timescale of 2020
- MHO: innovation toward HL-LHC needs to be an adiabatic process involving SW, computing and infrastructure. And needs to start now

Any other ideas?