

# Future Work on Virtualisation and Clouds: Introduction to the Discussion

Michel Jouvin  
LAL, Orsay  
[jouvin@lal.in2p3.fr](mailto:jouvin@lal.in2p3.fr)

GDB, January 2013

# From December GDB

- Report by Tony Cass about the legacy from HEPiX WG on Virtualization
  - > Demonstrated feasibility of using images maintained by “users” (experiments) on virtualized resources operated by sites
  - > Tackled the different issues: see Tony’s talk
- Proposal to have this work taken over by WLCG to enable a wider use of the cloud resources set up at many sites
  - > Focus on the interaction between pilot frameworks and cloud resources
  - > Experiment actual issues like passing credentials to VMs
  - > Evaluate the future need for batch systems in this model or the ability to do batch-less VM instantiation
- Today’s discussion: assess interest for this proposal and if there is one, start to discuss a work plan

# From Yesterday's pre-GDB...

- ◎ Session on cloud tests and plans during Operations Coordination pre-GDB but with a different focus
  - > Issues with accounting
  - > ATLAS and CMS tests of CERN Agile infrastructure
  - > Plans to use HLT farms as standard computing resources during LS1
    - Most experiments plan to run a cloud to achieve this
- ◎ Accounting : main issue is accounting of public cloud resources used as there is no option to access the accounting services
  - > ATLAS exposed possible plans based on using PanDA job statistics
  - > No full agreement that public resource usage should go into WLCG accounting
  - > For private clouds, APEL has demonstrated its ability to collect resources from clouds as it does with other non grid resources

# ... From Yesterday's pre-GDB

- ◎ CERN Agile infrastructure tests
  - > Joint activity from ATLAS and CMS, **driven by IT/ES**
  - > Done with a mix of test jobs and real production/analysis jobs
  - > ~200 VMs of 4 cores/8 GB for each experiment
  - > Condor used to instantiate VMs (VM provisioning)
    - Not used exactly the same way by both experiments as ATLAS is using PanDA
  - > Successful tests, despite some unexpected CPU efficiency problems in ATLAS that are still to be understood
- ◎ HLT farms during LS1: all experiments except LHCb planning to build a cloud from their HLT clusters for running offline jobs
  - > ATLAS/CMS using OpenStack but distinct from CERN Agile infra
  - > ALICE will use CernVM Cloud
  - > In every case, coexistence between HLT usage and this one

# Today's Discussion

- ◎ Short introduction by T0 and a few experiments
  - > Yesterday's presentations illustrated that there was different approaches for VM submissions, in particular regarding the role of batch systems
- ◎ Really a discussion on the main topics experiments (and sites!) are interested to work on with respect to cloud usage
  - > Is batch-less interaction with cloud of interest to several experiments? Is there some work in progress?
    - LHCb and ALICE at least seem interested
  - > What are the main issues to address based on tests done so far?
    - Ex: Image management/endorsement, credential forwarding to VM...
  - > How to organize the future work, if any is agreed?