

Pre-GDB - Clouds

# Testing the CERN Agile infrastructure

Fernando Barreiro  
Mattia Cinquilli  
Marek Denis  
Katarzyna Kucharczyk  
Ramón Medrano

- Common cloud testing
  - AgileInfrastructure@CERN
- Testing process and architecture
- First results
- Summary

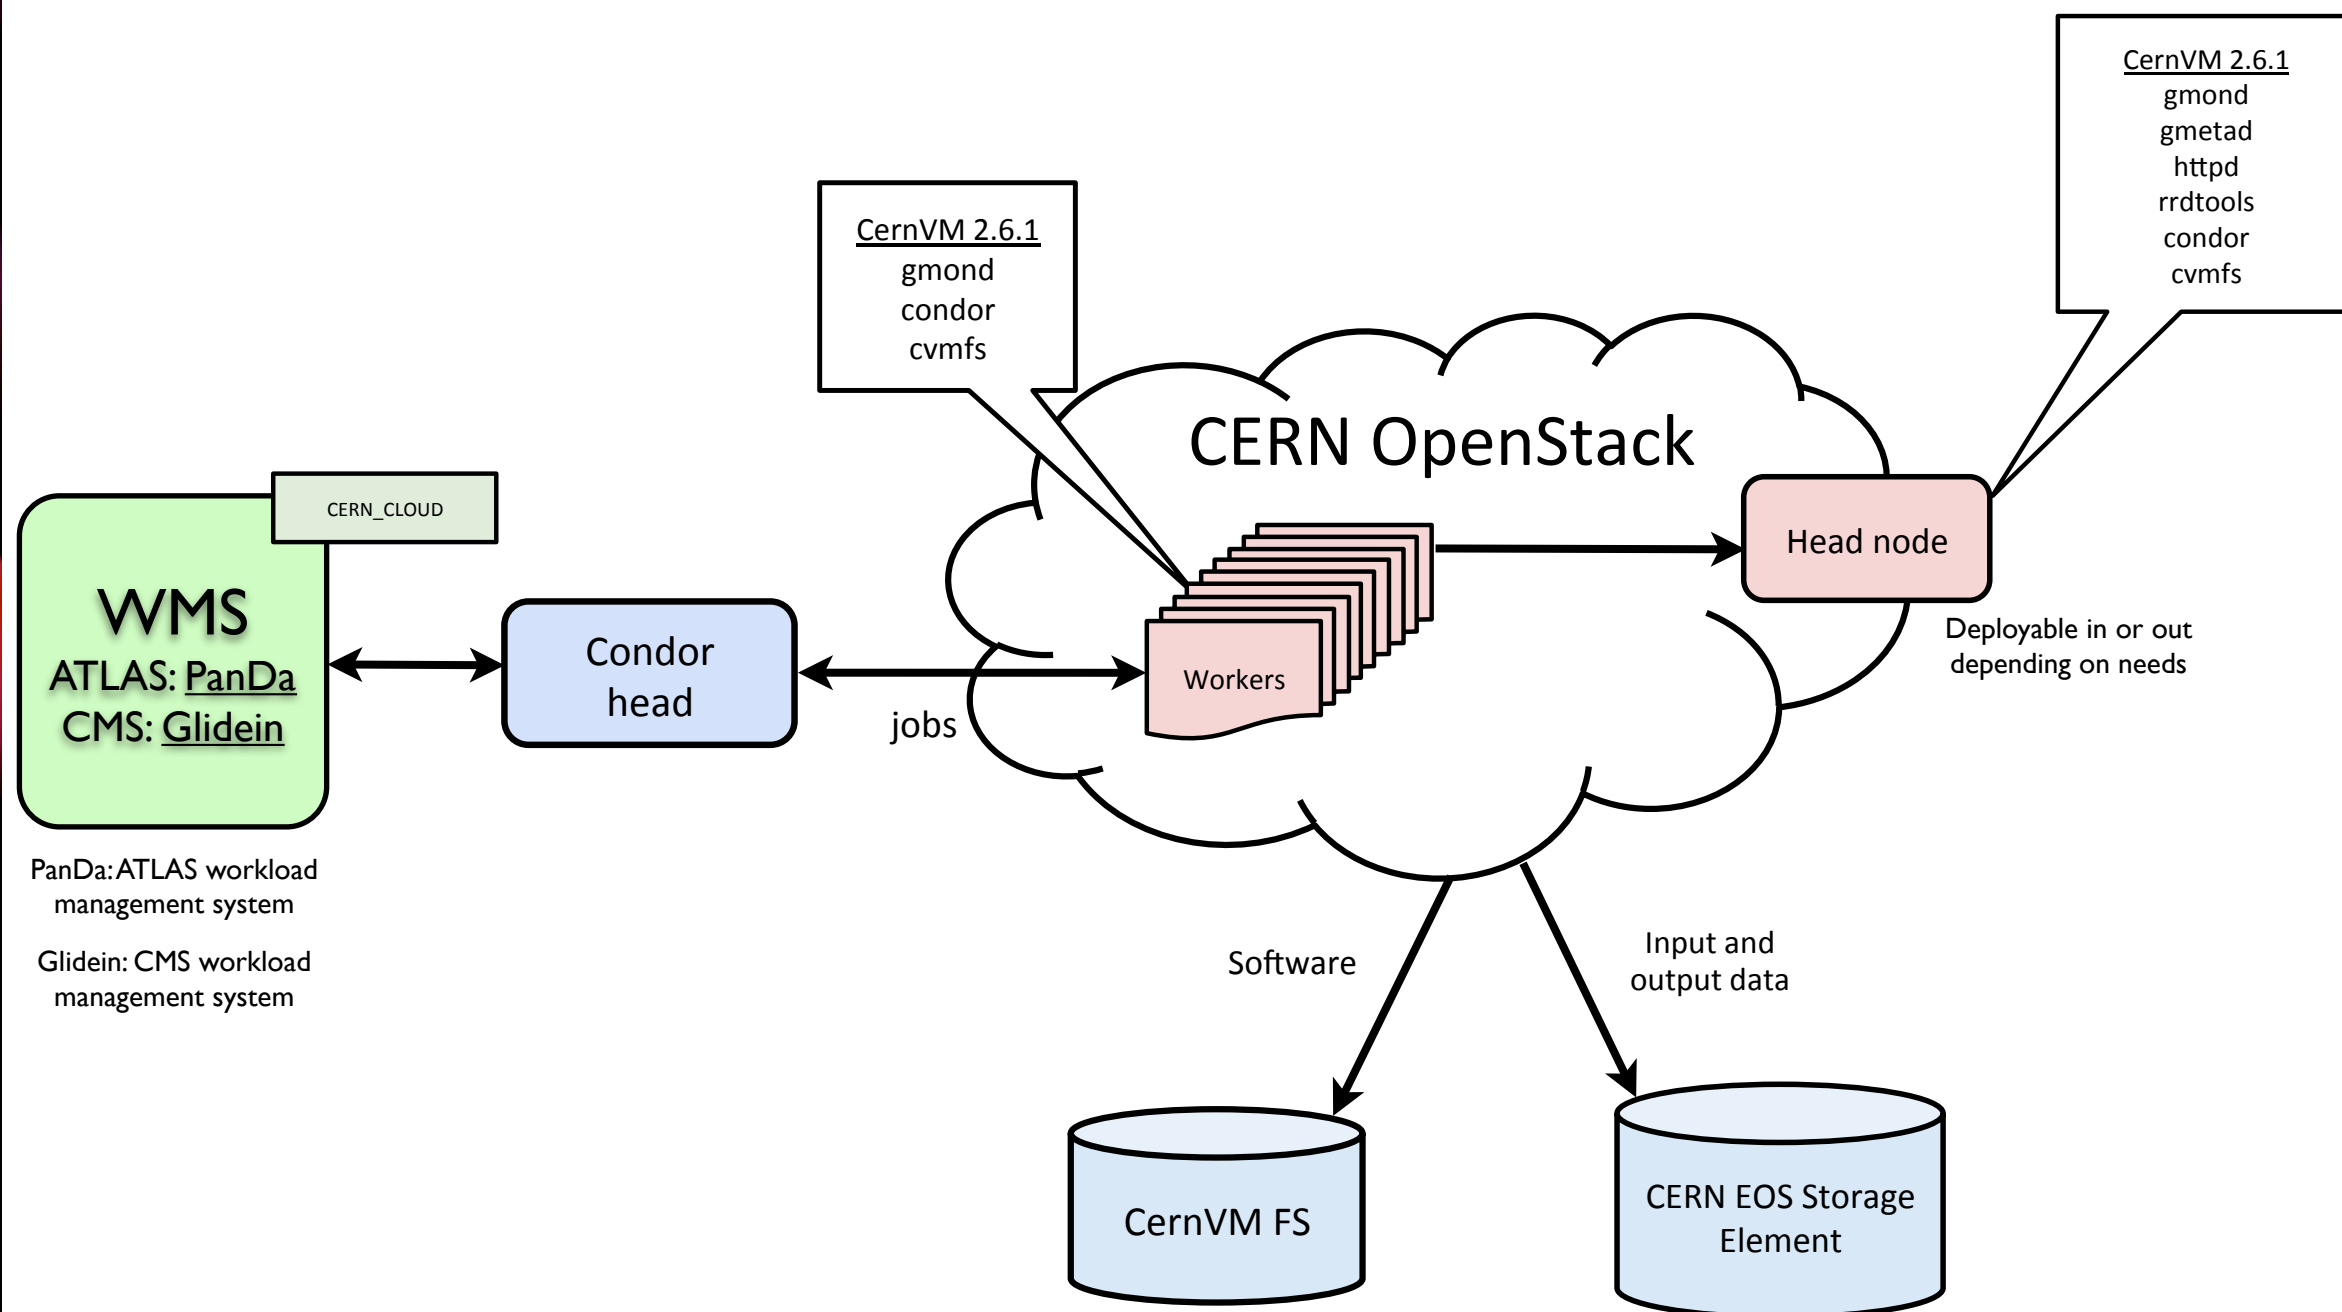
- ✓ Functional
  - ✓ ability to transparently interface with different cloud providers
- ✓ Performance
  - ✓ understand performance and reliability compared to a Grid site
- ✓ Extra-resources
  - ✓ transparently offload experiment workflows to clouds
- Possible future ambition
  - ➔ use commercial services as a significant fraction of overall resources available?

- Common strategy between ATLAS and CMS
  - LHCb starting
- Already tested other cloud infrastructures
  - StratusLab, LxCloud
- Currently focusing on AgileInfrastructure@CERN
  - IT point of view:
    - productive system testing
  - Experiments point of view:
    - first time running jobs in the new CERN IaaS
    - possibility to use extra resources

- Project reviewing the CERN computer-centre management toolset
  - improve infrastructure efficiency and management
  - increase functionalities and QoS
- Horizontal IaaS: OpenStack as cloud software for virtual machines and image management
  - Two service models:
    - Pets: unique, long lived, with given names
    - Cattle: identical to other cattle, random names



- (1) Upload image to the cloud provider repository
- (2) Boot VM with uploaded image
- (3) Contextualization of the VM
- (4) Interface VM with experiment Workload Management System
- (5) Test - Monitor - Improve cycle



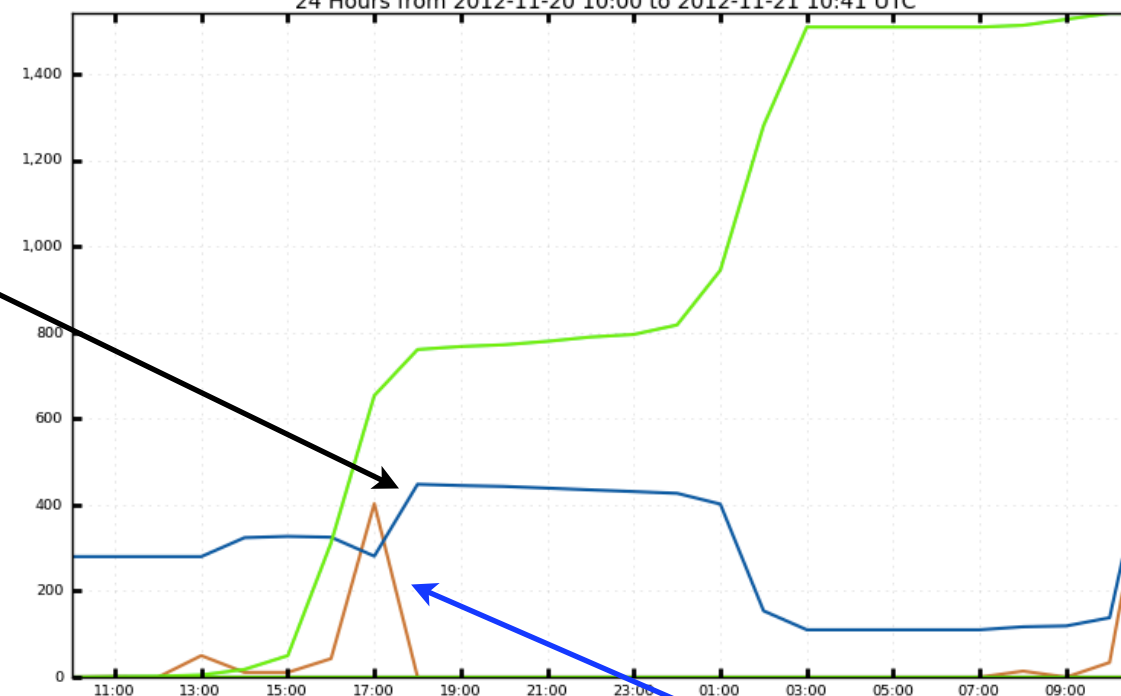
- ATLAS
  - Job type:
    - HC tests and real production (evgen, merge, pile, reco, simul, validation)
  - constantly 200 VMs (4cores, 8GB RAM, 80GB disk)
- CMS:
  - Job type:
    - MonteCarlo: pythia simulation tests running ~1h/job
    - analysis: ~8h test jobs reading /DoubleElectron/Run2012B-PromptReco-v1/AOD from EOS
  - peak of 200 VMs (4cores, 8GB RAM, 80GB disk)




dashboard

OpenStack CMS testing

24 Hours from 2012-11-20 10:00 to 2012-11-21 10:41 UTC

~450 jobs running  
in parallel~700 jobs running  
in parallel



OPENSTACK\_CLOUD - week

pending jobs  
correctly absorbedPlots taken from PanDA  
monitoring and Dashboard

ATLAS

CERN - 15 January 2013

Jobs correctly load  
Virtual Machines  
resources



Plots taken from VM  
resource monitoring  
system configured during  
contextualization

- Job efficiency is very high (1% failures for the ATLAS use case)
  - CERN-PROD failure rate 0.2-05%
- Apparently network latency not an issue
  - deeper investigation being done
- First results on CPU efficiency:
  - ATLAS tests with HammerCloud to evaluate CPU/network performances
    - overhead for the CPU that sits somewhere in between 7% and 21%
    - virtualized environment? under investigation
  - No HammerCloud tests for CMS yet



- CERN IT Agile Infrastructure provides a total of ~1.6K cores for our testing (800 for each ATLAS and CMS)
  - 200 VMs: 4-cores, 8GB RAM, 40-80GB local disk
- Manually handling VM life cycle
  - CLI/Web tools (e.g.: euca-tools, nova client)
  - customized scripts or tools

- BNL development of automatic provisioning in PanDA pilot factory on going
- GlideinWMS has automatic provisioning through Condor
  - VMs requested/stopped depending on job queue
  - already tested on AI with sleeping jobs
  - need to send real jobs and to increase the scale
  - need to evaluate potential load issues



- Work conceptually very close between the experiments
  - sharing existing efforts, experiences and tools
- Demonstrated the possibility to run real jobs on the AgileInfrastructure@CERN
- Evaluate puppet for machine configuration
- Investigate job efficiency metrics
- Test automatic provisioning systems