



# **Batch Virtualization coordination**

(and cloud computing at CERN)

- iozone tests follow-up:
  - check read ahead
  - check drop caches
  - Follow-up: possible kernel issue: some VMs developed unkillable processes. After a reboot, the issue was gone.
- cvmfs support: idea is to run on hypervisor + export
  - Follow-up : according to Jakob this will not work. He recommended to run cvmfs on each node. Sharing caches wouldn't work either.

## Adjustments needed:

- Availability of people
- New LSF master is late
- No solution for dynamic hosts, may have to go for static configuration

## Description

## Due

- **First virtual machines based on ISF in production system** 26/11/2010 → 9/12/2011
- **ONE production master ready** 6/12/2010
- **LSF master replacement** 19/11/2010 → Jan 2011
- **First dynamic virtual nodes in the production system** 29/11/2010 → Jan 2011
- 96 virtual machines in the production system 23/12/2010
  
- First SLC6 release 17/01/2011 (assumption)
  - Migration to SLC6 on hypervisors 23/02/2011
  - First SLC6 golden node 14/03/2011
  - First SLC6 worker node 15/03/2011

- ▶ ISF status (Belmiro)
  - ▶ Master node status
  - ▶ Production release preparation
- ▶ ONE status (Sebastien, Ulrich)
  - ▶ Packaging
  - ▶ Master node installation status
- ▶ Worker node and Ixcloud status: Ulrich
- ▶ VM worker node benchmarking: Qiulan

## Worker node status: performance tuning

- ▶ `enable_ept=0` configured on lxcloud in quattor
- ▶ Additional LVM areas on lxcloud (per VM slot):
  - ▶ pool
  - ▶ cvmfs

## Worker node status: performance tuning: /pool

- ▶ Vm64slc5 has support for external /pool
  - ▶ Areas already created
  - ▶ Mounted at contextualization time
  - ▶ Required changes in LXBATCH
    - ▶ Patched LSF rpm with additional script (in preprod)
    - ▶ Configuration change on the Grid queues
- ▶ Under test in batch7
  - ▶ vmce002 (CREAM) submission works fine

## Worker node status: performance tuning: /pool

- ▶ Timelines:
  - ▶ Thursday: roll out of new LSF version into preprod
  - ▶ Monday morning:
    - ▶ Roll out into prod
    - ▶ Change of grid queue configuration
  - ▶ Tuesday: start using of new image in prod

## Worker node status: CVMFS

- ▶ 2GB cache externally mounted
- ▶ New LSF attribute “cvmfs”
- ▶ Will use frontier SQUIDs
- ▶ Configuration in sync with plus/batch



## Throughput tests

- ▶ 16 vi\_10\_20 nodes in share\_short downgraded to 8 job slots only
- ▶ 2 more hypervisors moved to production with ONE
  - ▶ issue with lease files in one. Under discussion with them
  - ▶ Need to drain existing leases in order to extend the lease file
- ▶ Idea: use final setup of WN, with /pool outside the image and enable\_ept=0. One more benchmark running now.

- ▶ Iptables setup in Ixcloud:
  - ▶ Requires review
  - ▶ Right now switched off
  - ▶ Updated configuration under test
- ▶ Code management:
  - ▶ Temporary move of VMIC code into SVN at CERN
- ▶ Plans for Xmas:
  - ▶ 96 VM in production, half ISF half ONE
  - ▶ Ideally vm64slc5 setup with CVMFS
  - ▶ Fallback solution vm64slc5kvm (as now)
  - ▶ Allow running also long jobs