

# **Virtualisation Working Group Report**

**Tony Cass  
HEPiX Spring 2012  
April 25<sup>th</sup> 2012**

# Agenda

- ◆ Image Generation Policy
- ◆ Image Contextualisation
- ◆ Image Exchange
- ◆ Summary

# Agenda

- ◆ ~~Image Generation Policy~~
- ◆ Image Contextualisation
- ◆ Image Exchange
- ◆ Summary

No changes...

# Agenda

- ◆ Image Generation Policy
- ◆ **Image Contextualisation**
- ◆ Image Exchange
- ◆ Summary

# Contextualisation update

- ◆ Previously:
  - Attempt to mount CDROM image
  - If successful, invoke scripts pre- and post- network initialisation
- ◆ Problem: Not compatible with EC2-style services
- ◆ New—and agreed—proposal
  - also attempt to download contextualisation scripts from <http://169.254.169.154>; site scripts at `/context/default.tgz`
    - » user contextualisation scripts may also be present, but site scripts are run after these!
- ◆ This is compatible with the previous scheme and works with multiple cloud orchestrators (such as OpenStack and OpenNebula).

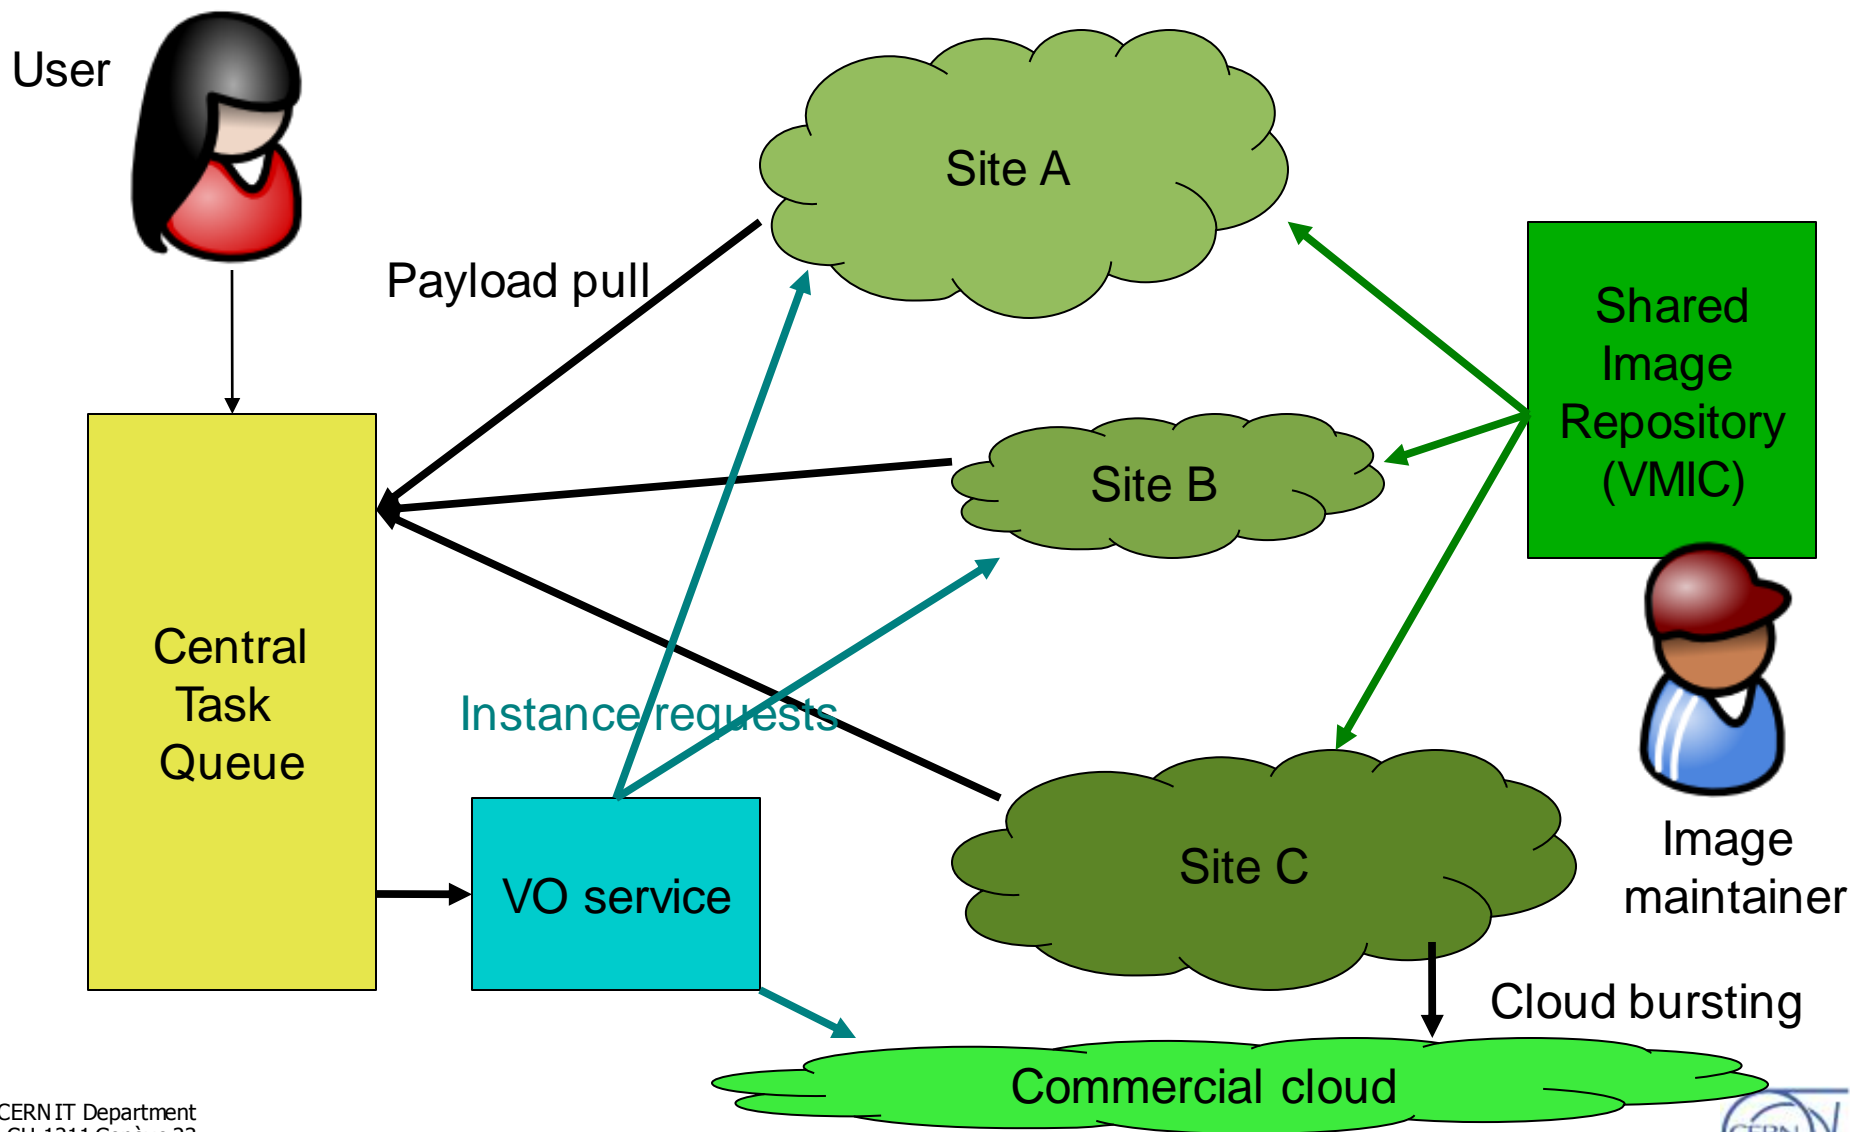
# PS: Machine Features

- ◆ Virtualisation (and batch) systems are supposed to make available
  - \$MACHINEFEATURES pointing to a directory which contains:
    - » hs06 : the hs06 rating of the worker node
    - » shutdowntime (if present) the epoch of the next shutdown of the box (i.e. if the box is draining)
  - \$JOBFEATURES: pointing to a directory containing details for current job:
    - » jobstart : the epoch of when your job started
    - » mem\_limit the current memory limit (if any)
    - » run\_limit : the max number of minutes the job may run
    - » cpu\_limit: the max number of CPU minutes the job may run
- ◆ Implemented for LSF@CERN; NiKHEF working on PBS implementation; discussing with IN2P3 for GridEngine. Anyone interested in implementing for another scheduler?

# Agenda

- ◆ Image Generation Policy
- ◆ Image Contextualisation
- ◆ **Image Exchange**
- ◆ Summary

# What We Want



Slide courtesy of Ulrich Schwickerath



# So, where are we?

- ◆ Where we should have been in Vancouver!
- ◆ F2F meeting finally arranged at CERN two weeks ago. Successful:
  - Image generated at DESY and published then
  - imported to CERN, LAL and DESY, and
  - instantiated and contextualised.
- ◆ Useful discussions on features of the vmlisub tool and integration with OpenStack/Glance and StratusLab.
  - See Owen's talk...

# Agenda

- ◆ Image Generation Policy
- ◆ Image Contextualisation
- ◆ Image Exchange
- ◆ **Summary**

# Summary

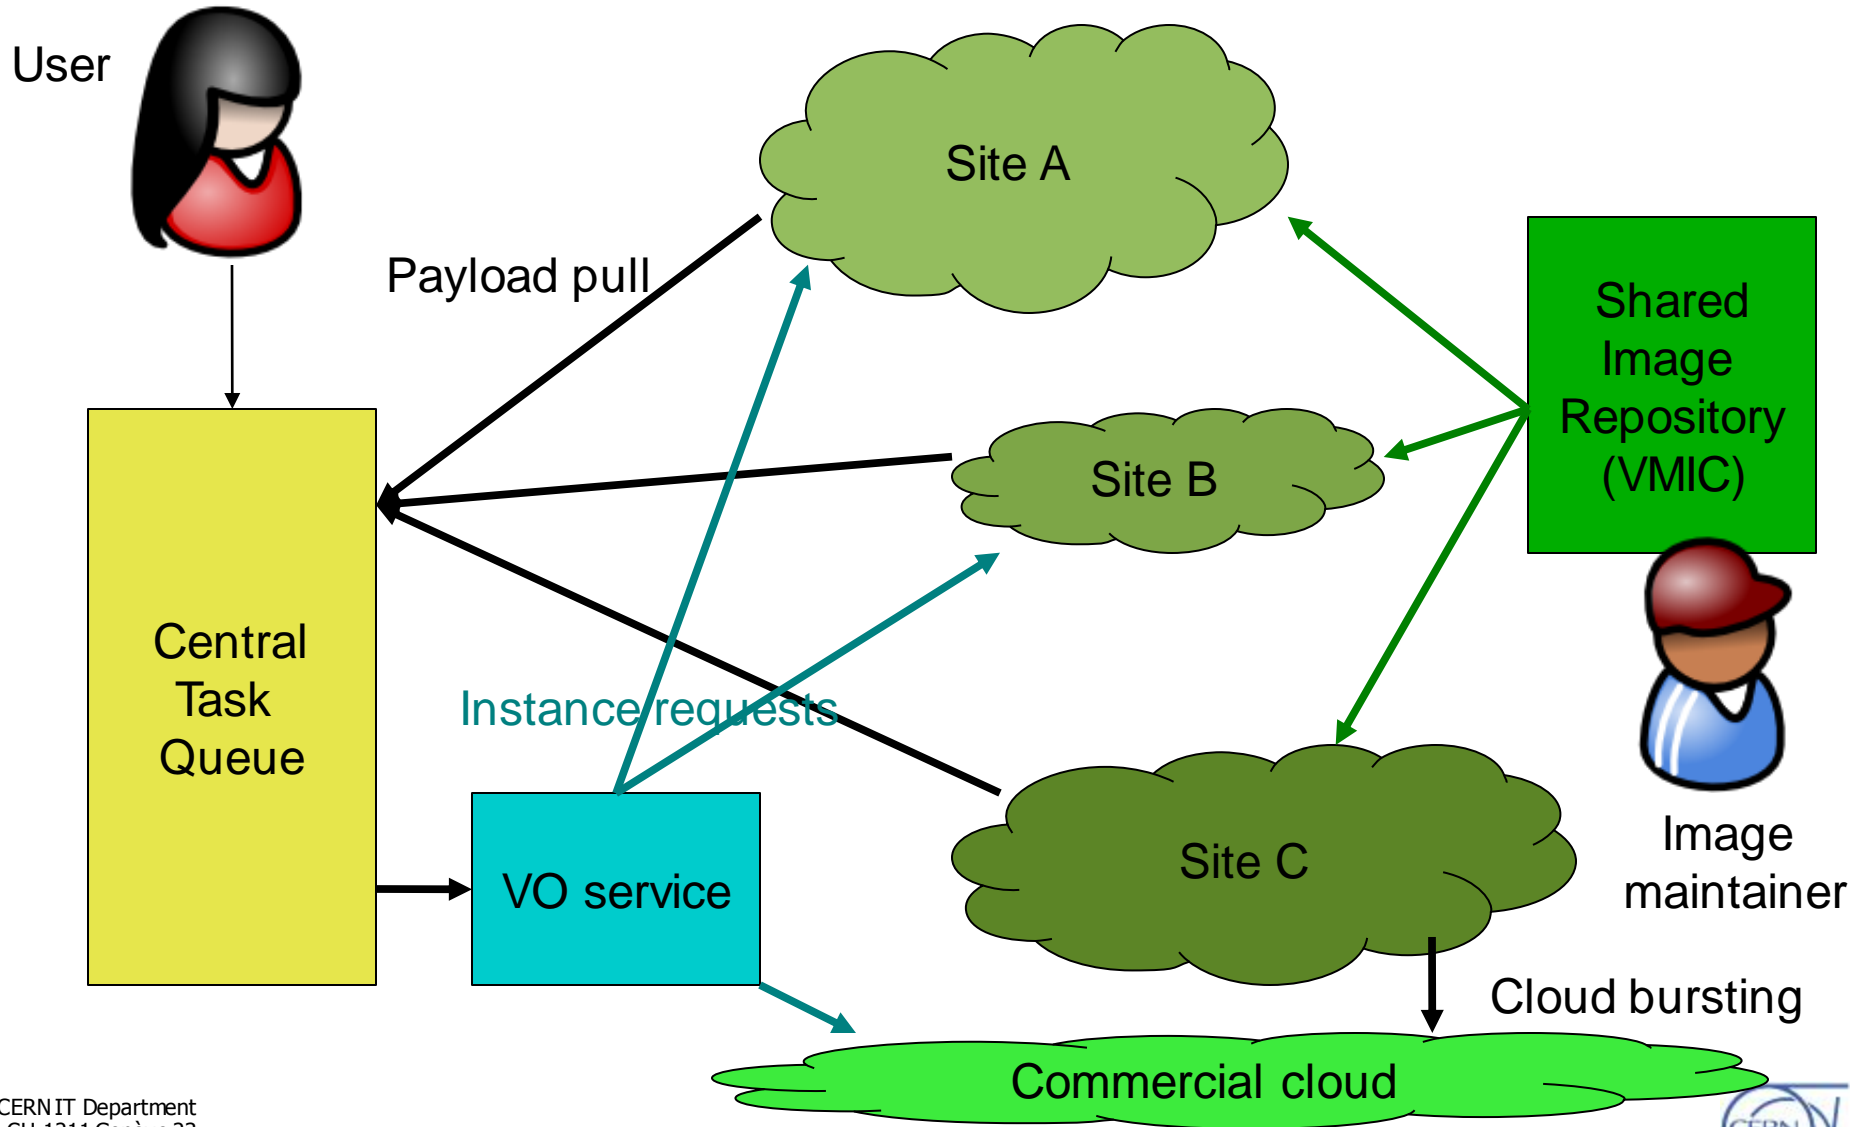
- ◆ The working group has made good progress in establishing policies to allow the exchange of VM images... ✓
- ◆ ... but not such good progress in delivering a distributed catalogue of endorsed images.

# Summary

- ◆ The working group has made good progress in establishing policies to allow the exchange of VM images... ✓
- ◆ ... and also good progress in delivering a distributed catalogue of endorsed images. ✓
- ◆ CVMFS is probably the neatest solution to the problem of VO software distribution... ✓
- ◆ ... but VM exchange remains interesting
  - as an option for sites to run hypervisors not used and automatically migrate to latest patched systems as images instantiate, and
  - if the VM images can contact pilot job frameworks directly, simplifying the scheduling problems at sites.

This is!! be feasible by HEPIX in Prague!  
More thought needed in this area... WLCG TEG?  
Followup at GDB and WLCG in NYC

# What We Want



Slide courtesy of Ulrich Schwickerath

