

The
HEPiX

Virtualisation
Working Group

Towards a Grid of Clouds

Tony Cass
CHEP 2012
May 24th 2012


Background

- ◆ The HEPiX virtualisation working group was formed to facilitate the instantiation of user-generated virtual machine images at HEPiX (and WLCG) sites.
- ◆ Users were expressing such a wish in 2008/9, but sites were worried about issues such as uncontrolled root access and the maintenance of the traceability logs required by Grid security policies.



Belle Monte-Carlo production on the Amazon EC2 cloud

Martin Sevier, Tom Fifield (University of Melbourne)
Nobuhiko Katayama (KEK)



17th International Conference on Computing in High Energy and Nuclear Physics
21 - 27 March 2009 Prague, Czech Republic

Particularly useful for Peak Demand

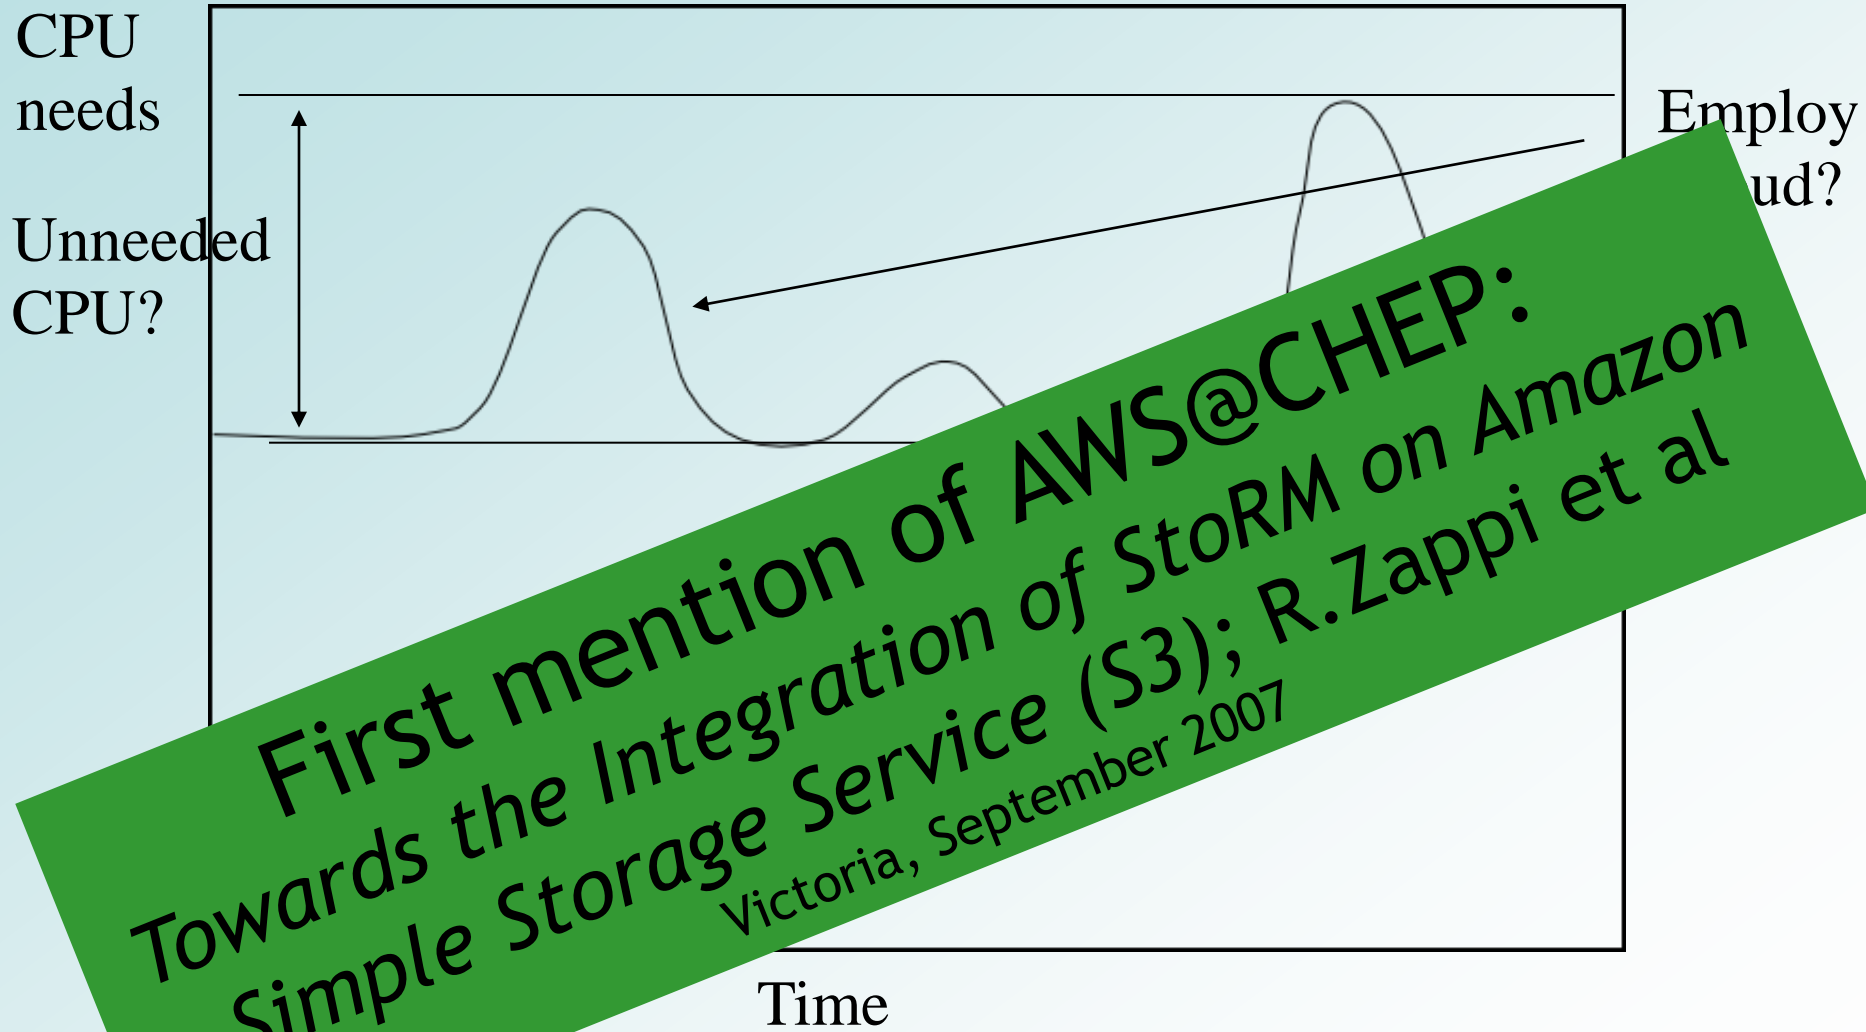


Image endorsement

- ◆ The HEPiX VWG developed a policy that introduced the concept of **image endorsers**: people who would **guarantee** that generated **images could be used safely** at sites.
- ◆ Amongst other things, such images would
 - have no embedded user credentials, and
 - enable sites to contextualise the images to enable the required logging and make other necessary customisations.
 - › Sites agree, however, not to modify the software environment of the image.
- ◆ Sites are free to trust (or not) specific image endorsers but, if they do trust someone in this role, it is expected that any images endorsed by this person can be used at that site **without the need for inspection or manual approval**.
- ◆ The HEPiX VWG policy became the basis of an approved JSPG policy document, “Policy Trusted Virtual Machines”.

Current Status

- ◆ The **endorsement policy is agreed**.
- ◆ Technical arrangements have been defined for
 - **image contextualisation**
 - » these are compatible with EC2/OpenNebula/OpenStack
 - exchange of information between the site infrastructure and a running virtual machine
 - » e.g. remaining lifetime, that the virtual machine can be terminated, ...
- ◆ A framework for image endorsers to publish and distribute images has been developed.
 - This has been integrated with StratusLab's marketplace at LAL and is being integrated with OpenStack Glance at CERN.
- ◆ **CERNVM images are compatible** with the HEPiX VWG policies
 - and there has been a security review of the underlying technology.

Job done then. What now?

◆ Why virtualise?

- Give user control over the execution environment
 - » An important topic in 2008/9, but addressed now by CERNVMFS, even for real machines.
- Improve overall resource management at sites.
 - » Still relevant but doesn't need user-supplied images.

Conclusions

- Grid works (for ~1% of all scientists)
 - Allows LHC to achieve scientific results almost instantaneously
- Operations are complex and costly
 - Still immature middleware and faulty hardware
 - Highly customized very different application frameworks
 - Different resource ownership and service levels
- **Clouds will not make Grids cheaper or redundant** ✓
 - **But surely will add extra complexity** ✗
- Grid is here to stay
 - Scientific data will always be distributed
 - Global science is a collaborative effort, and so is Grid
- Standardization and convergence to common approaches is badly needed
 - Otherwise Grid efficiency will remain relevant only to few selected applications, like HEP
- Something totally different will certainly come

Job done then. What now?

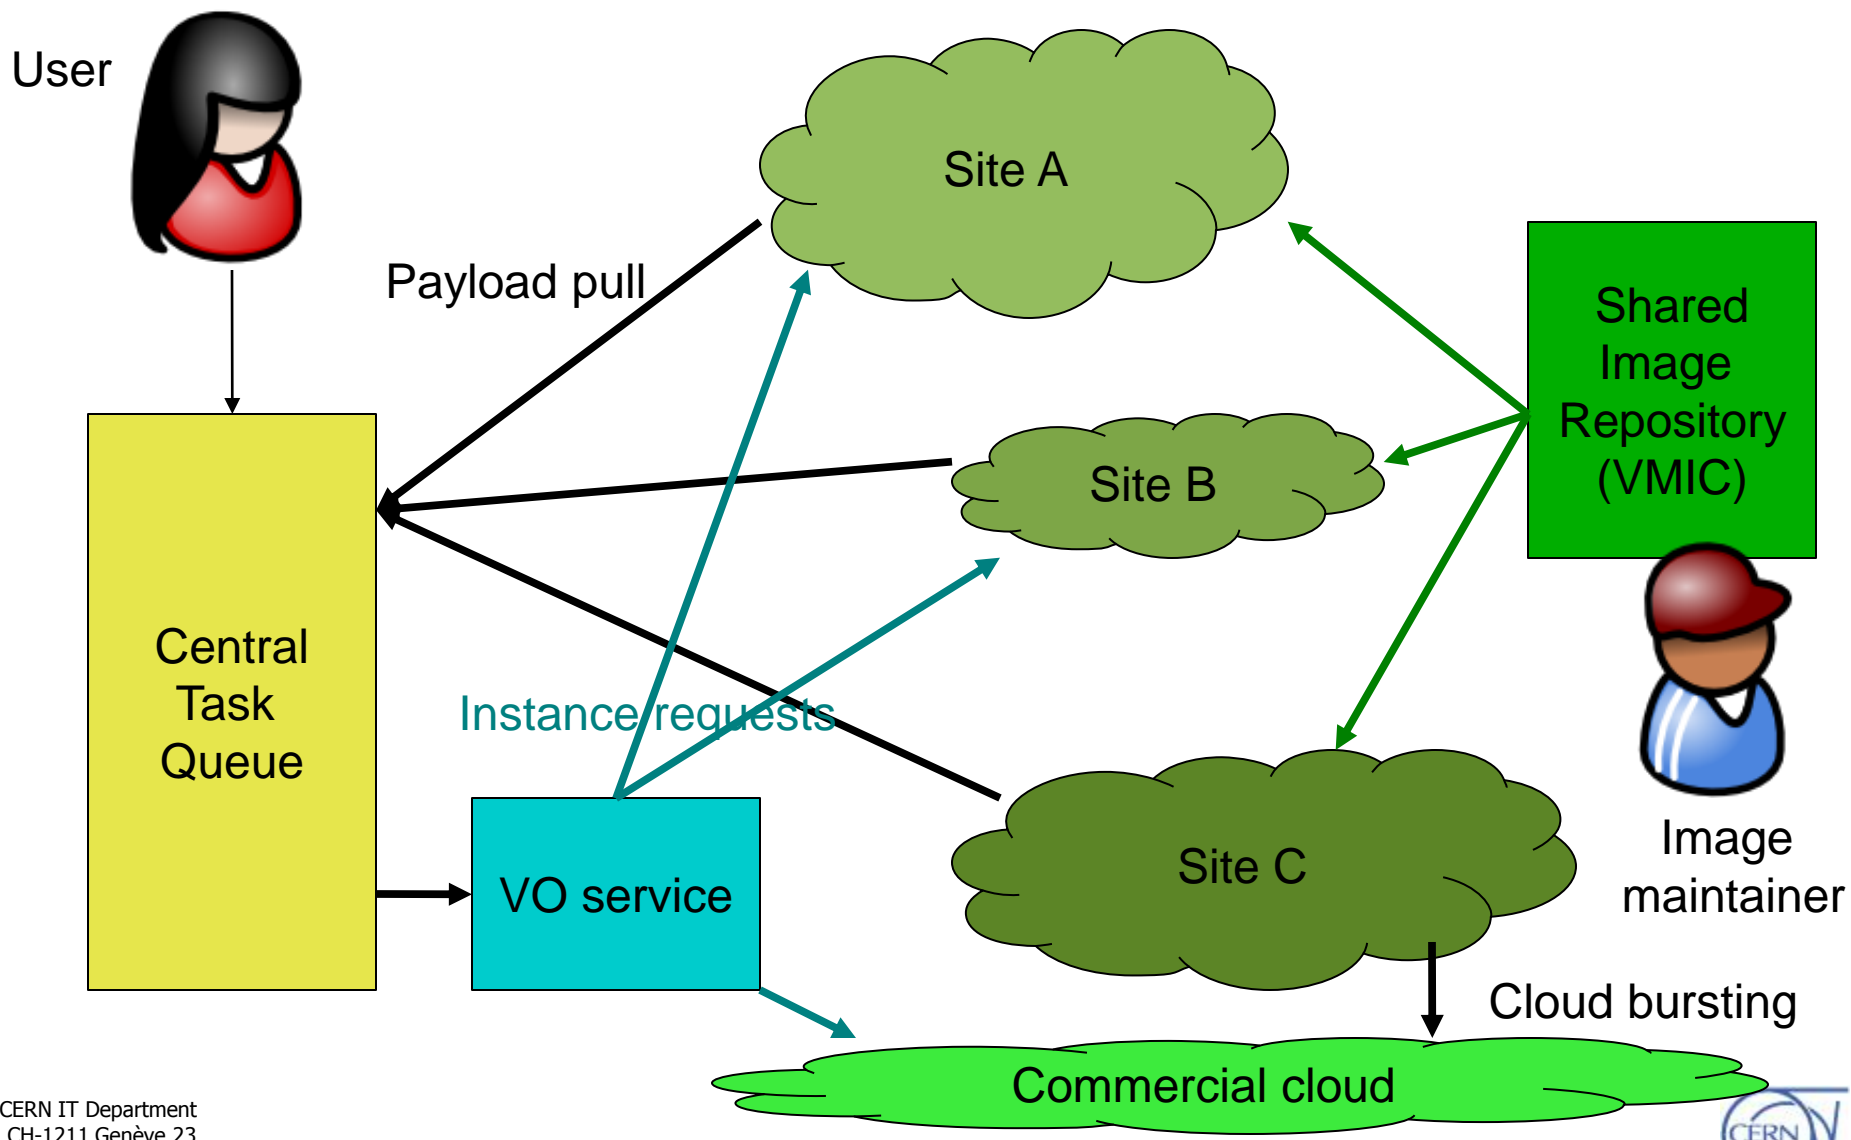
◆ Why virtualise?

- Give user control over the execution environment
 - » An important topic in 2008/9, but addressed now by CERNVMFS, even for real machines.
- Improve overall resource management at sites.
 - » Still relevant but doesn't need user-supplied images.

◆ Exploit user-supplied images and cloud interfaces to simplify meeting the different goals of

- **sites**: use resources efficiently and respect commitments for resource sharing, and
- **experiments**: allocate tasks amongst available resources to meet physics priorities.

How this could be used



Summary

- ◆ The HEPiX Virtualisation Working Group has shown how trusted user generated images can be safely instantiated at Grid sites
 - compatible with site security policies and obligations, and
 - with a guaranteed environment for the experiments.
- ◆ Three options now for following this up
 - Ignore and hope it goes away
 - Integrate trusted virtual images into a traditional batch environment
 - **Exploit the proposal** to deliver a “Grid of Clouds” that enables sites to easily and simply present to experiments a seamless—but dynamically changing—set of resources that the experiments can exploit to schedule work according to physics priorities.

We don't have to embrace it

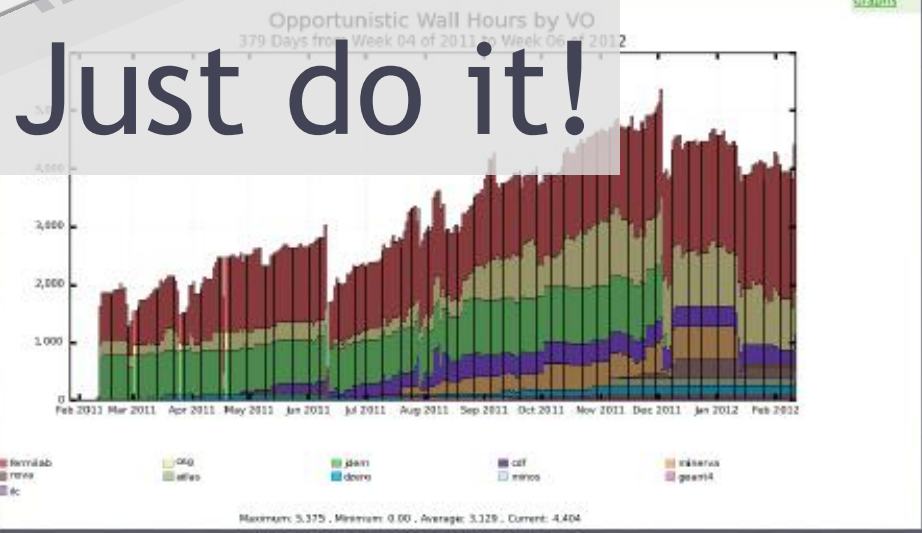
From Ian Gable
"Grid of Clouds" use for ATLAS
but it needs to grow



embrace this by Amsterdam

- Pares in comparison to the scale seen in industry

Just do it!



Thanks to Ulrich Schwickerath and Steve Timm for figures
Disclaimer: Opinion is not theirs :)

