

# **Virtualisation Working Group Report**

**Tony Cass  
HEPiX Spring 2011  
May 6<sup>th</sup> 2010**

# Organisation

- ◆ Slow start, mostly due to re-org @ CERN.
- ◆ 66 people on mailing list; core of ~10 regular participants, but many more join meetings on occasion.
  - Many thanks to Ian Gable for taking notes during meetings.
- ◆ Group identified 5 work areas
  - Image Generation Policy
    - » Dave Kelsey & Keith Chadwick
  - Image Exchange
    - » Owen Synge
  - Image Expiry/Revocation
    - » Later agreed to be part of policy & exchange area
  - Image Contextualisation
    - » Sebastien Goasguen
  - Multiple Hypervisor Support
    - » Andrea Chierici

# Policy for Trusted Image Generation

- ◆ You recognise that VM base images, VO environments and VM complete images, must be generated according to current best practice, the details of which may be documented elsewhere by the Grid. These include but are not limited to:
  - any image generation tool used must be fully patched and up to date;
  - all operating system security patches must be applied to all images and be up to date;
  - images are assumed to be world-readable and as such must not contain any confidential information;
  - there should be no installed accounts, host/service certificates, ssh keys or user credentials of any form in an image;
  - images must be configured such that they do not prevent Sites from meeting the fine-grained monitoring and control requirements defined in the Grid Security Traceability and Logging policy to allow for security incident response;
  - the image must not prevent Sites from implementing local authorisation and/or policy decisions, e.g. blocking the running of Grid work for a particular user.
- ◆ [http://www.jspg.org/wiki/Policy\\_Trusted\\_Virtual\\_Machines](http://www.jspg.org/wiki/Policy_Trusted_Virtual_Machines)

# Image Cataloguing and Exchange

Change Virtual Machine Image | Django site admin

cern.ch https://vmrepo.cern.ch/vmic/admin/

## Change Virtual Machine Image

History

**VMI endorsement**

Endorser: Romain Wartel +

**VMI download location**

VMI filename: Amstrad\_OS3.tar.gz

**Status of the VMI**

This VMI is APPROVED to be run locally     This VMI can be shared with other sites

**Metadata about the VMI**

VMI UUID:	Amstrad_OS_1234	Production date:	Date: 2010-08-16 Today   📅
			Time: 14:17:34 Now   🕒
Endorsement date:	Date: 2010-08-16 Today   📅	VMI checksum:	13242345
	Time: 14:17:37 Now   🕒	Hypervisor:	Xen

**Metadata about the VM**

OS version:	Amstrad OS
Architecture:	ARM

Done

MS

)



# Image Contextualisation

- ◆ Contextualisation is needed so that sites can configure images to interface to local infrastructure
  - e.g. for syslog, monitoring & batch scheduler.
- ◆ Contextualisation is limited to these needs! Sites may not alter the image contents in any way.
  - Any site are concerned about security aspects of an image should refuse to instantiate it and notify the endorser.
- ◆ Contextualisation mechanism
  - Images should attempt to mount a CDRROM image provided by the sites and, if successful, invoke two scripts from the CDRROM image:
    - › prolog.sh before network initialisation
    - › epilog.sh after network initialisation

# Multiple Hypervisor Support

## ◆ Andrea

- Surveyed sites; results show that kvm and Xen dominate as hypervisors, especially in batch virtualisation area.
- Documented method to produce VM image that can be used with both kvm and Xen
  - » Method tested by Sebastien Goasguen and Abdeslem Djaoui (RAL)

# Current Status @ HEPiX in Cornell

## ◆ Generation policy

- Clear, but probably needs to be formally approved by JSPG?      **DK: Wait until we have real experience**

## ◆ Contextualisation & kvm/Xen support

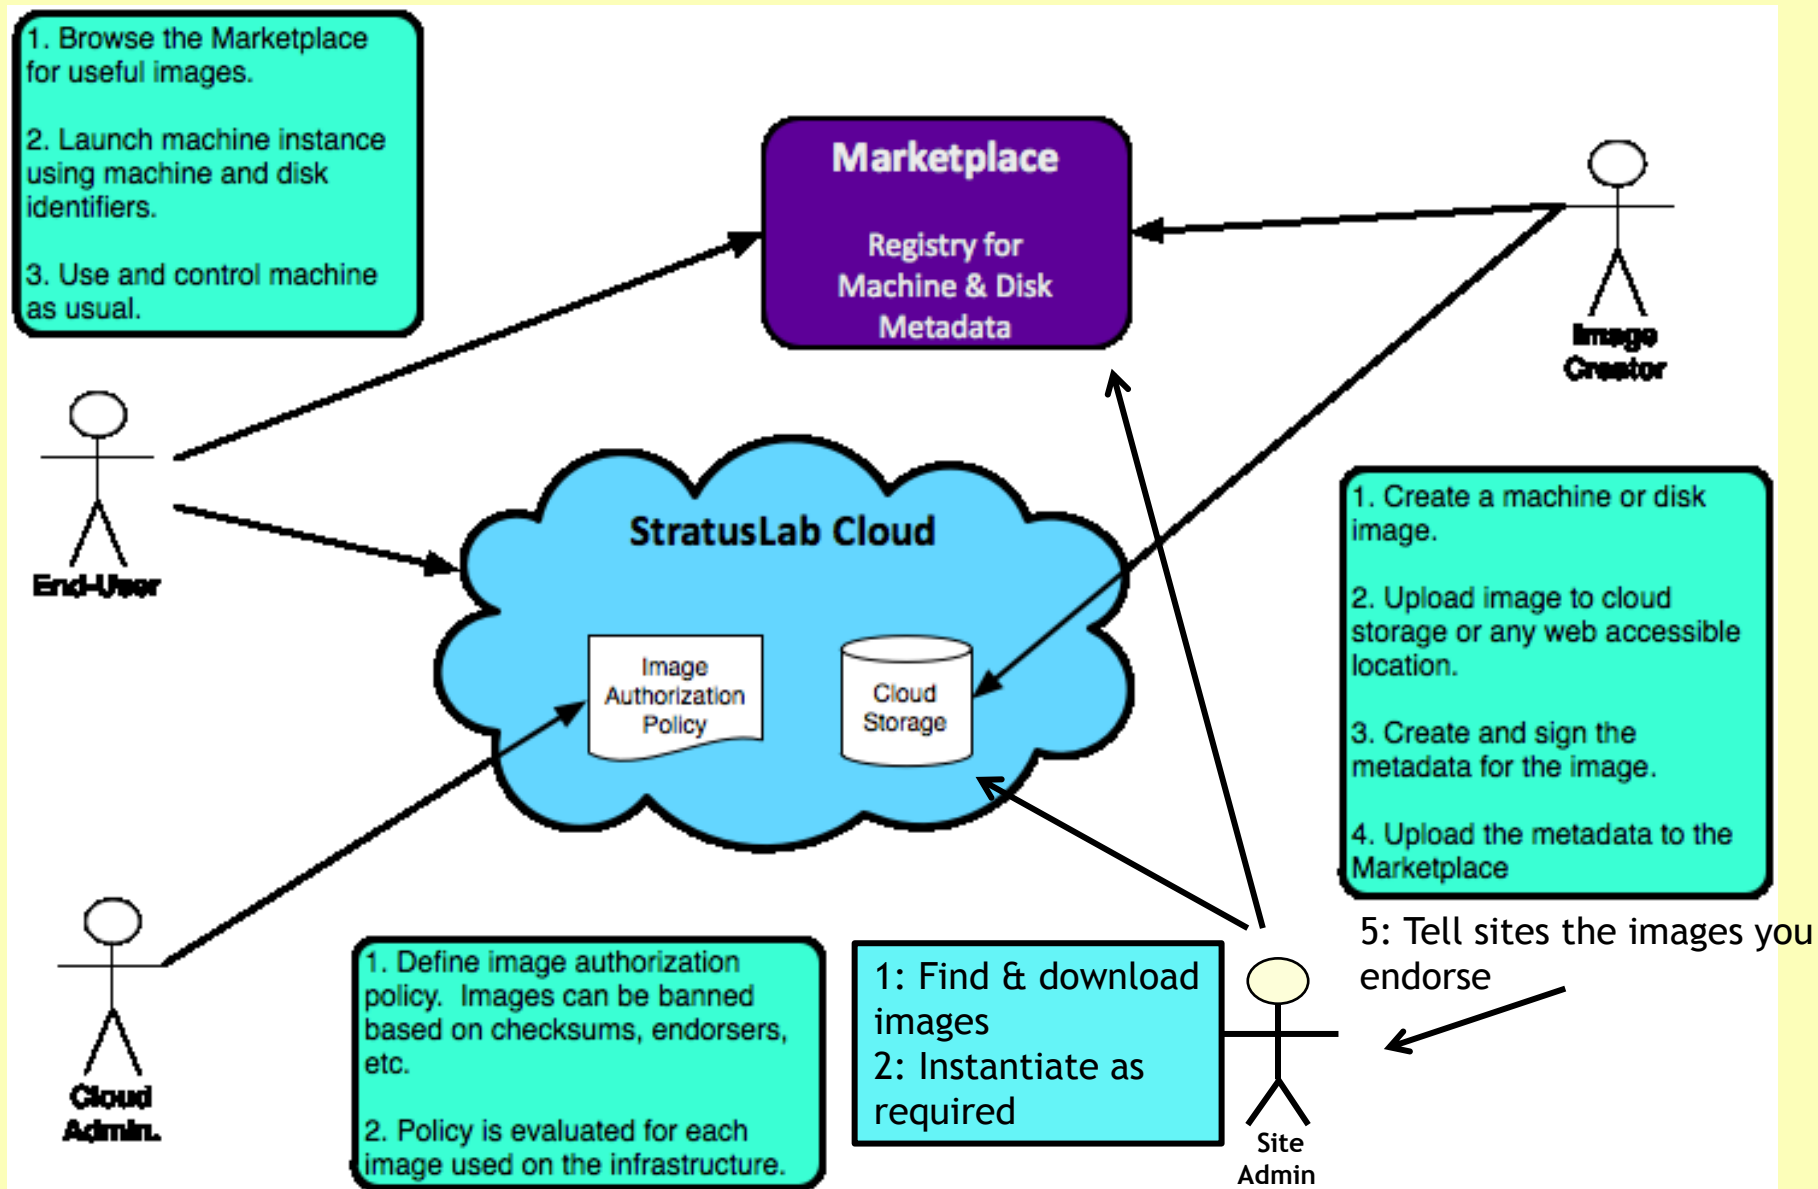
- Also clear.      **Interest in Xen waning**

# Progress in 2011

- ◆ Owen funded to work on virtualisation
- ◆ Stratus Lab discussion



# Stratus Lab Workflows



# Progress in 2011

- ◆ Owen funded to work on virtualisation
- ◆ Stratus Lab discussion
- ◆ Victoria->CERN image distribution

# Other thoughts



# 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 Work distributed catalogue of endorsed images. in Progress!
- ◆ CVMFS is probably the neatest solution to the problem of VO software distribution... ✓  
Also, images created at one site (Victoria) have been
- ◆ ... but VM exchange remains interesting  
instantiated and contextualised  
- as an option for sites  
However, work is  
automatically  
coupled at  
instantiated  
more coherent  
- If the VM  
Video conference “days” planned to achieve this.  
Still trying to bring expt on board. ATLAS?  
works directly,  
systems at sites.

