Trusted Virtual Machine Images

a step towards Cloud Computing for HEP?

Tony Cass
on behalf of the HEPiX Virtualisation Working Group

October 19th 2010

Motivation for the Working Group

- Virtualisation is attractive!
  - For sites: ease operational procedures
  - For VOs: guarantee execution environment
    - » Although VOs remain worried about virtualisation penalty
- The EC2 model not acceptable for many sites
  - if infrastructure is not able to provide adequate isolation for images where user has root access
    - » e.g. where NFS (or other protection systems based on uid/gid identification) is used to provide access to shared resources.
- HEPiX working group established to enable exchange of trusted virtual machine images between sites
  - Assumes that not all users will need to generate VM images. This is reasonable as not all users need to install software at sites in today's Grid environment.

Approach

- Establish a policy for generation of trusted images
- Document method for sites to contextualise images to meet local policies
- Enable exchange of trusted images between sites.

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

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 CDROM image provided by the sites and, if successful, invoke two scripts from the CDROM image:
    - prolog.sh before network initialisation
    - epilog.sh after network initialisation

Image Cataloguing and Exchange

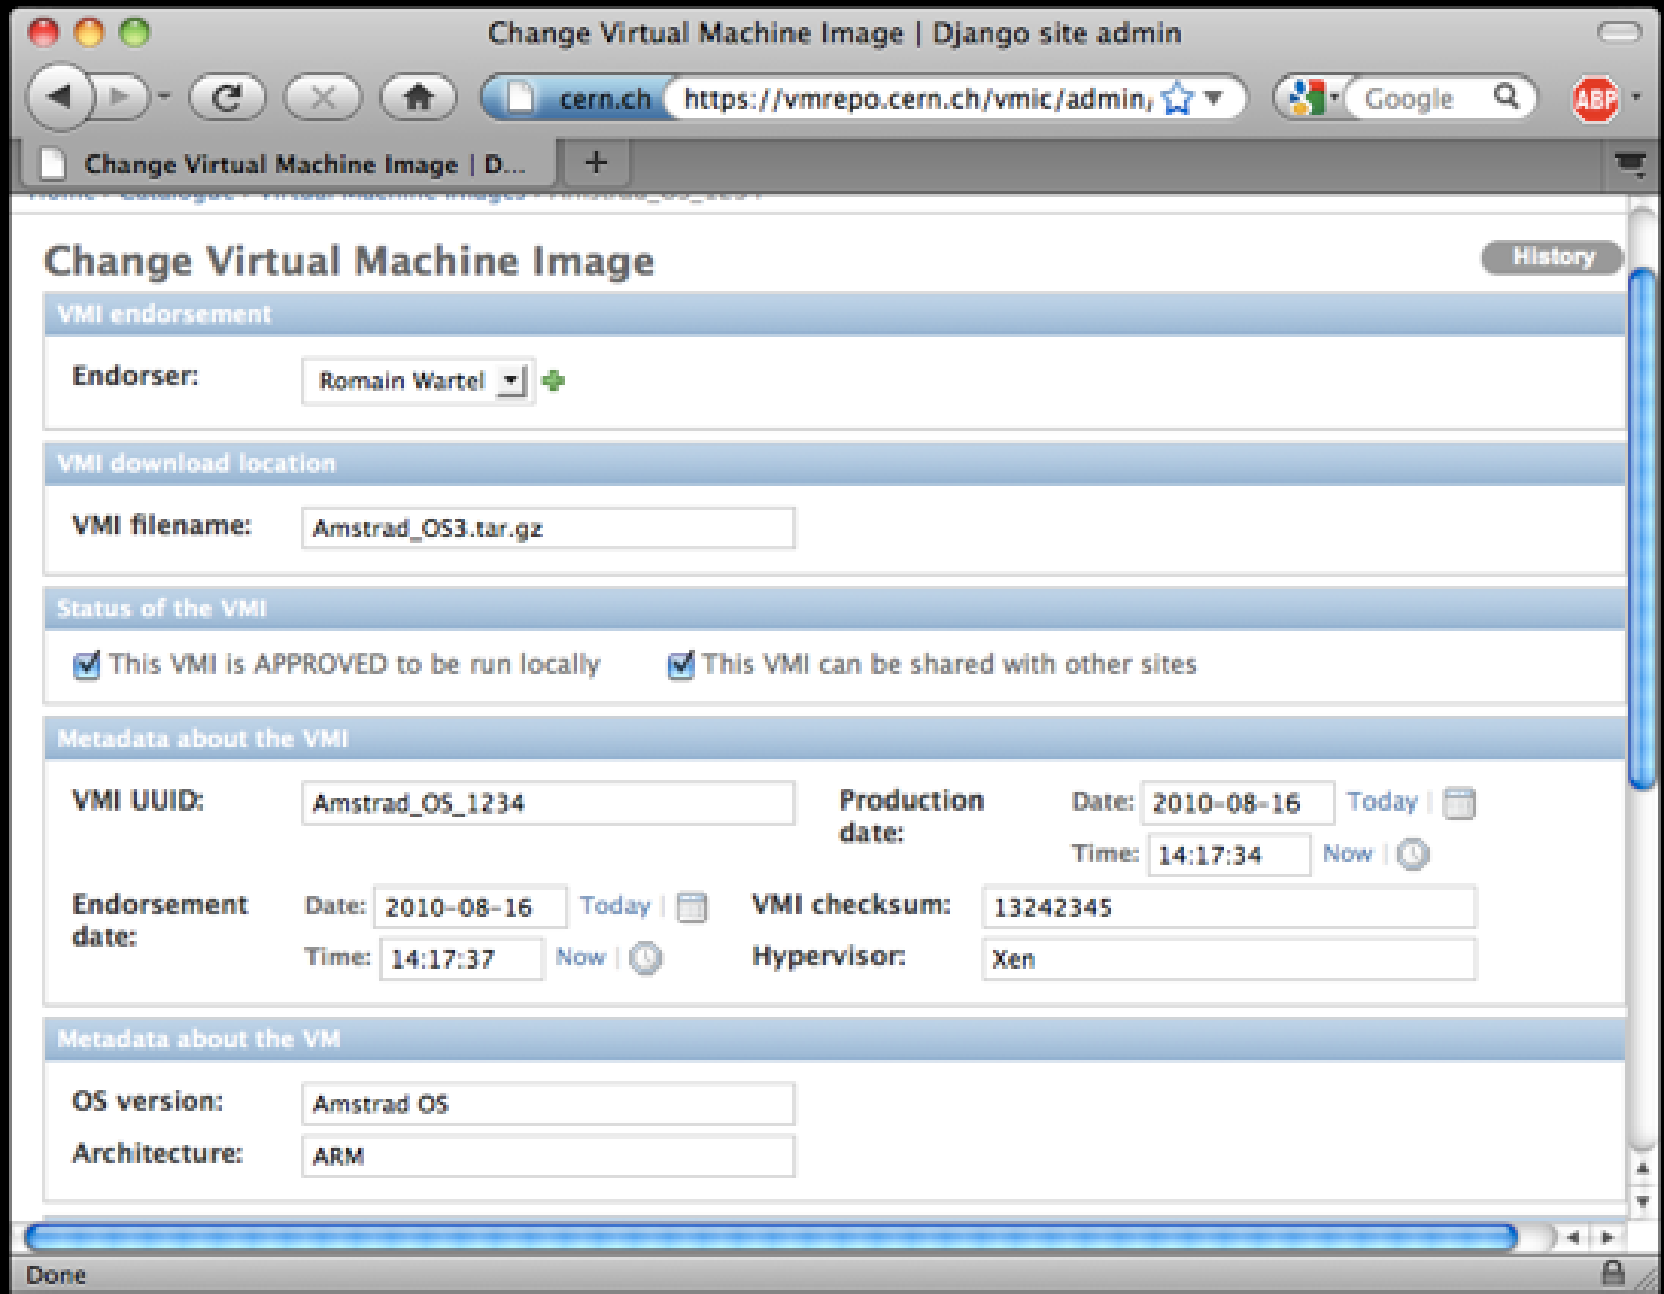

A step towards Cloud Computing?

- Most (all?) current implementations of virtualised cpu servers integrate virtual machines with local workload management systems
  - requests to the local WMS lead to instantiation of virtual machines
  - locally provided machines have WMS client installed; contextualisation allows WMS client installation for remotely provided images.
- There is an alternative!
  - Sites can instantiate VM images that connect directly to a VOs pilot job infrastructure.
  - This decouples completely
    - » Site role to allocate resources according to funding
    - » VO interest in managing priorities between different users and tasks.
  - Dynamic instantiation of VM images according to demand patterns leads to changing “virtual batch system” for the VOs.

Summary

- The HEPiX Virtualisation Working Group has established a framework which should enable the free interchange of virtual machine images between HEP sites
  - a policy governing how images should be generated, to enable trust between sites and endorsers, and
  - a mechanism for contextualising images to be compliant with local policies at the instantiating site, and
  - a mechanism for cataloguing endorsed images and to track images approved for instantiation at a given site
- VO supplied images could connect to pilot job infrastructure directly, not the local workload management system
  - allowing sites to instantiate images dynamically according to workload patterns, a step towards Cloud Computing for HEP.
- Interestingly, the StratusLab project has very similar ideas (see PS50 @ 16:30). Collaboration looks to be a possibility, especially in the Image Catalogue area.