WLCG Trusted Virtual Images

Michel Jouvin

EGI Virtualisation Workshop

Amsterdam, May 12th 2011

Disclaimer - Credits

- This is not WLCG official views on virtualisation
  - There are several views, use cases and experiences in WLCG
- This work comes from sites
  - HEPiX is a site admin forum (http://www.hepix.org)
- Goal: enable a virtualised environment for user jobs in the grid world
  - Virtual image produced/maintained by VO, not by site
  - Transparent to existing job submission frameworks/paradygms
  - Not a cloud approach… but doesn't prevent it
- Most slides stolen from Owen Synge and Tony Cass presentation at last HEPiX (GSI, May 6th)

Challenges for Trusted Images

- Image Generation Policy: policy defining basic rules to build an image eligible to be trusted
  - Mostly based on the idea of widely best practices
  - Defines a few restrictions, eg. no root access
  - Discussed in JSPG framework: http://www.jspg.org/wiki/Policy_Trusted_Virtual_Machines
- Secured Image Transfer and Exchange
- Image Contextualisation: well-defined mechanism to allow sites to customize the image configuration at VM startup without modifying it
  - Cannot be used to update the system (eg. kernel, libs)
- Multiple Hypervisor Support, mainly Xen/KVM
  - Xen interest waning

Trusted Virtual Image

- A standard virtual image + some metadata + 1 signature
  - Signature of metada + a SHA512 hash of the image
- Image is signed by an *endorser*
  - Endorser role is to certify (not validate) the image has been built according to the policy…
  - … and to revoke it in case an image is considered no longer appropriate for use (eg. if a security vulnerability affects the image)
  - An endorser is typically a well-known person from a VO
- Endorsed images must be approved by a site admin before instanciation at a given site
  - Site retains control on what it accepts on a per-image basis
- A site is not allowed to patch an image
  - Must refuse to instantiate it if it has a concern

Endorsement / Approval Workflow

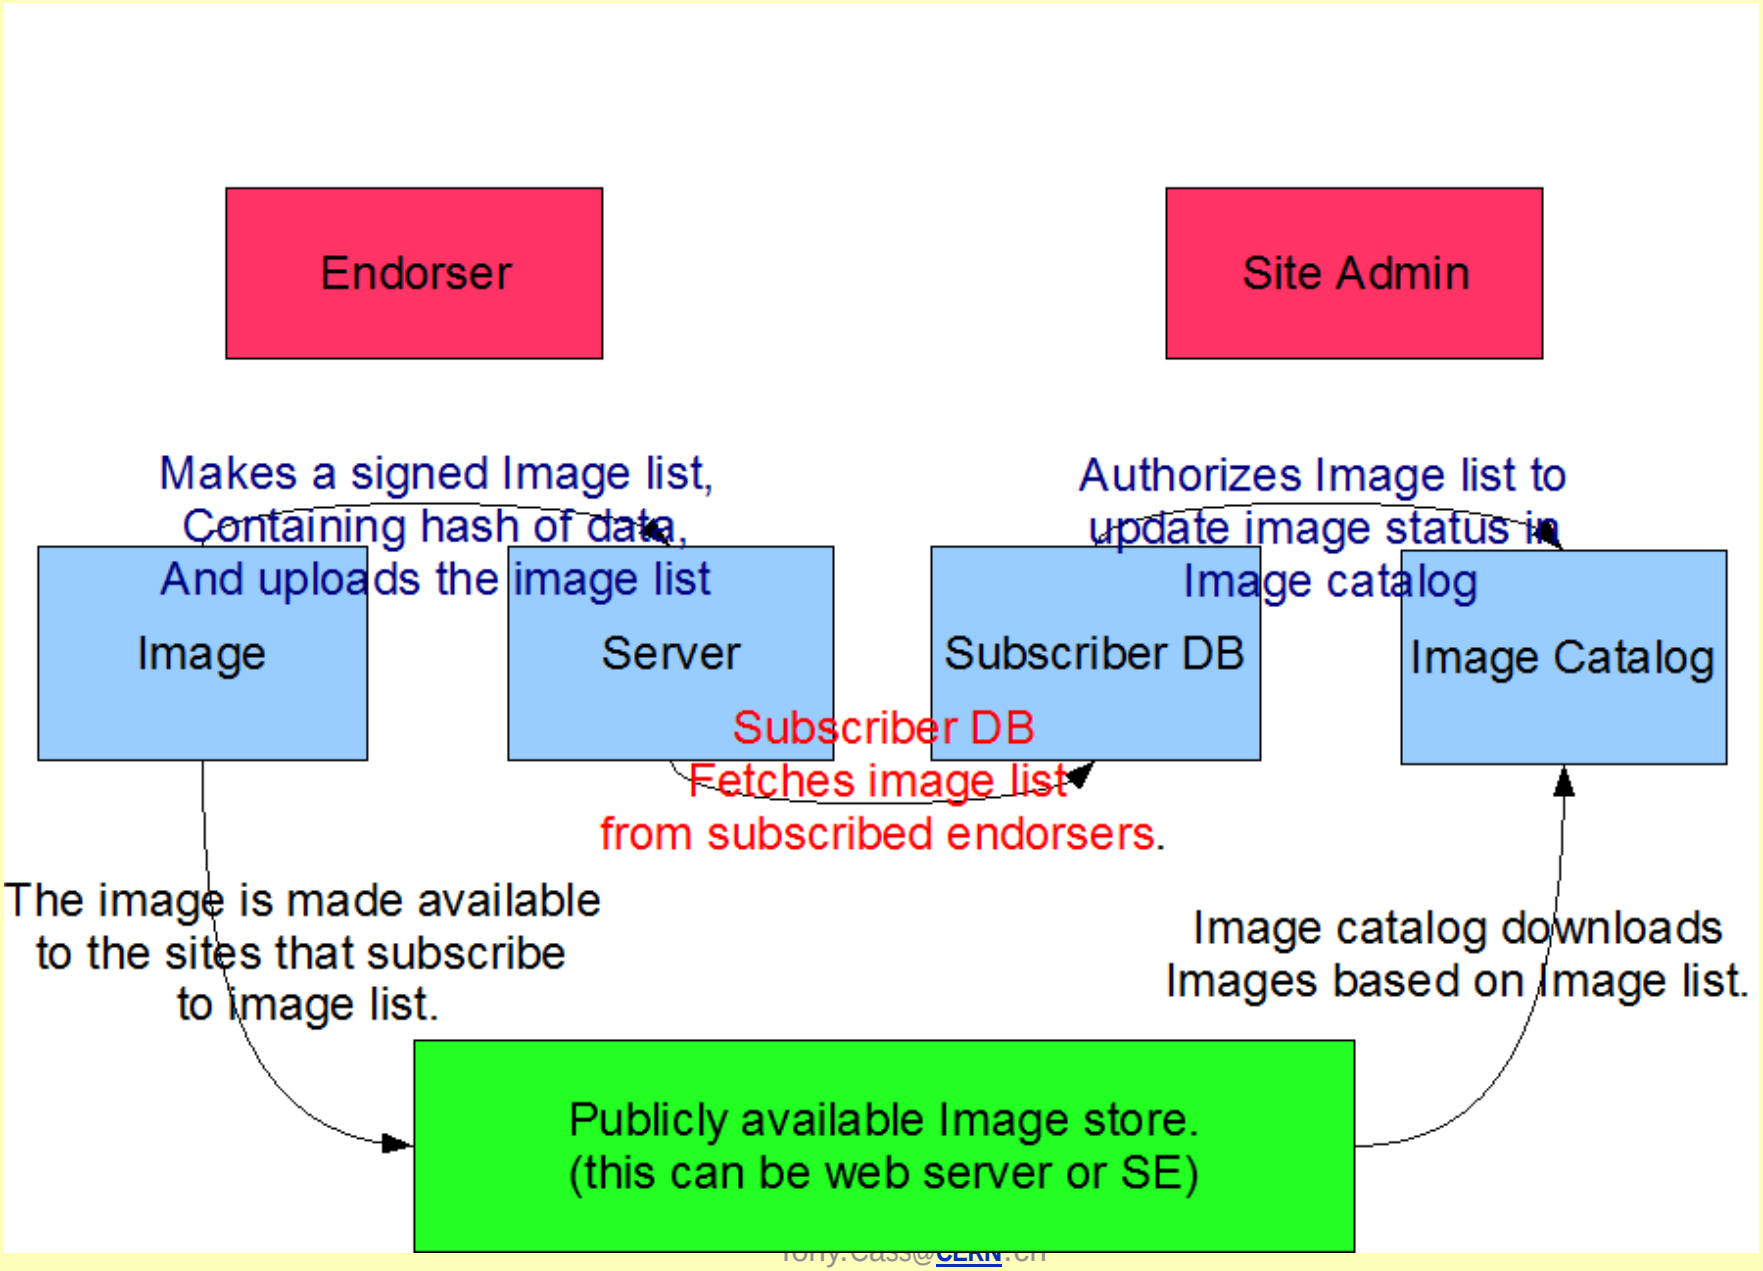

StratusLab Marketplace Workflows

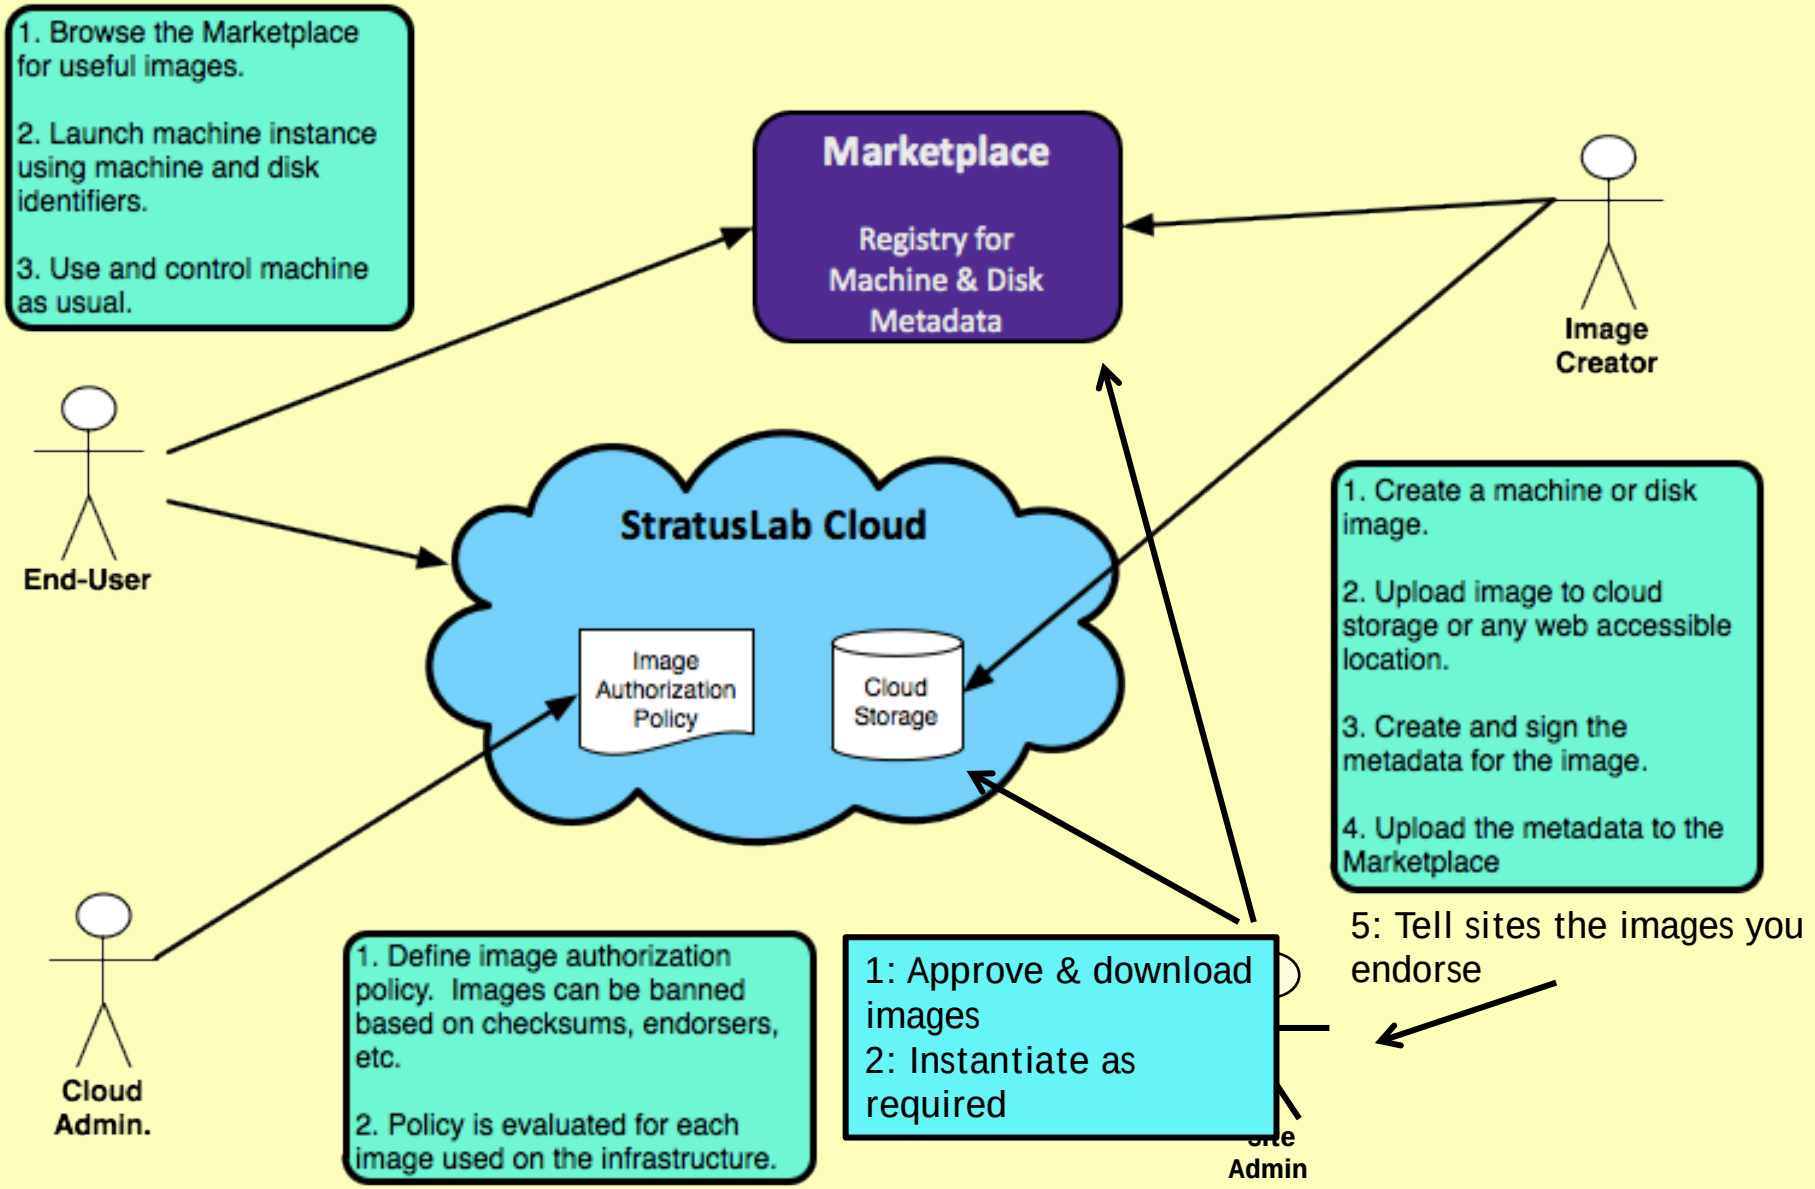

Other thoughts

- Virtualisation is an area with much scope for communication failures!
- We must be clear that "image endorsement" is a very rapid process
  - As soon as it is endorsed, an image is then immediately available for instantiation by all sites who trust the endorser; no need for a lengthy process of verification at sites.
- Some sites talk about restricting instantiated VM images but the actual impact for end-users is likely small
  - e.g. VM images would have no need to connect to a NFS-based shared storage area, in particular if using CVMFS for SW distribution
- If the VM images can contact pilot job frameworks directly, may simplify the scheduling problems at sites.

Summary

- Virtualised environment for job execution may improve flexibility and usability of grids by existing and new users
- Trusted Virtual Images are a key enabler for acceptance by sites of images produced by VOs
  - Different from arbitrary user images
- Image endorser role is at the heart of the trust but sites retain capability to approve/revoke endorsed images
  - Revocation by site is preventing that image maintenance is supported by site
- StratusLab Marketplace may offer the basis for a long-term supported trusted VM infrastructure

Backup Slides

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