HEPiX Virtualisation Working Group

Status, ~~February 10th 2010~~ April 21st 2010

Tony.Cass@cern.ch

Objective

- Enable virtual machine images created at one site to be used at other HEPiX (and WLGC) sites.
- Working assumptions
  - images are generated by some authorised or trusted process
    - » Some sites may accept "random" user generated images, but most won't
  - images are "contextualised" to connect to local site workload management system
    - » But at least one site (other than CERN) is interested in seeing images connect directly to experiment workload management system.

No root access by end user during image generation.

Recipient site controls how "payload" ends up in the image

Working group areas & Status

- Generation
- Transmission
- ~~Expiry & Revocation~~
- Contextualisation
- Support for multiple Hypervisors

Image endorser required to revoke images in case of security issues and the like.

Working group areas & Status

- Generation
  - Led by Dave Kelsey & Keith Chadwick
  - Likely to produce
    - » Policy proposal for image generation process. If sites can demonstrate they meet the requirements of the policy then their images should be trusted for execution at remote sites
    - » Recommendations for hypervisor configuration to ensure maximum security.

Sites anyway expected to follow best practices.

- Transmission
- Expiry & Revocation
- Contextualisation
- Support for multiple Hypervisors

Current discussion is around roles and endorsers for the different components ("base" operating system and VO software) and about who can be trusted.

Working group areas & Status

- Generation
- Transmission
  - Led by Owen Synge
  - Likely to produce
    - » Recommendation for basic transport protocol(s) to be supported
      - ➢ Prescriptive for sites wishing to generate images
    - » Proposal for optional protocols to improve transmission efficiency
      - ➢ E.g. transmission of only differences w.r.t. a reference image
      - ➢ Status of "interesting" protocols such as bitTorrent likely to be an issue.
  - ~~Unlikely to~~ Will not comment on intra-site image transmission
- Expiry & Revocation
- Contextualisation
- Support for multiple Hypervisors

Current model is tagged images distributed in manner akin to mechanism used for VO software today.

Working group areas & Status

- Generation
- Transmission
- Expiry & Revocation
  - Status a little unclear
    - » a mix of standalone area and generation policy?
  - "Image Revocation List" a la CRL?
    - » Technical proposal required
- Contextualisation
- Support for multiple Hypervisors

Image endorser required to revoke images in case of security issues and the like.

Working group areas & Status

- Generation
- Transmission
- Expiry & Revocation
- Contextualisation
  - Led by Sebastien Goasguen
  - Likely to produce
    - » Proposal for mechanism allowing site to configure image
      - File system mounted at image instantiation and automated invocation of scripts on the file system during the initialisation.
      - Final job/payload will not execute as root
    - » Restrictions on aspects sites are allowed to configure
      - No changes to C compiler, perl, python, … to be allowed
- Support for multiple Hypervisors

Only basic discussions so far. Contentious issue is kernel patching. Group conclusion is that this is not allowed; sites who have security concerns with an image must refuse to run this and must notify the endorser to allow wider revocation. This ensures that all sites are protected.

Working group areas & Status

- Generation
- Transmission
- Expiry & Revocation
- Contextualisation
- Support for multiple Hypervisors
  - Led by Andrea Chierici
  - Produce, if possible,
    - Recommendations/recipe(s) to enable sites to generate images that can be used with a range of hypervisors
      - Perhaps a limited set of all possible, however,…
      - Poll underway to identify most popular hypervisors

Little discussion in the group so far. We have identified the hypervisors of interest (kvm and both Xen modes). Andrea is testing extensively at present.