A Vision for Virtualisation in WLCG

Tony Cass

WLCG GDB, 9/9/9

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Goals

- Enable experiments/users to choose environment for job execution.
- Ensure sites have control/traceability over resource usage.

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Approach

- Step-by-step: Build on
  - established successes
  - established trust
- But end goal in view. Prepare for this now with
  - technical agreements/developments
  - user behaviour (especially explicit statement of resource requirements)

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Approach

- Five steps
- Steps 1-3
  - realistic
  - relatively uncontroversial(?)
  - achievable by end-2010?
- Steps 4 & 5
  - kite-flying
  - probably controversial
  - interesting

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Step 1

- Users can choose between virtual images created at sites.
- Not really any different from now; could be rephrased “sites provide virtual machines for job execution, not real hardware”.
- Key issue is (full) understanding of resource requirements
  - OS type, memory, (range of) #cores, ...

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Step 2

- Distribution of virtual machine images between sites (or from CERN...).
  - Image limited to minimalist operating system (SL4/5/6...)
- Requires
  - transparent process for image generation guaranteeing content
  - mechanism for sites to hook into local monitoring and batch scheduling.
  - trusted and verifiable method of image distribution

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Step 3

- Distributed virtual image includes experiment software environment
  - So users can choose ATLAS version X on OS Y.
- Requires “transparent process for image generation” to be extended to include experiment software.
  - Snapshot of experiment build servers at CERN?
- Removes need for pilot jobs to verify (or create) correct environment.

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

What about CernVM?

- Instantiation of CernVM machines being discussed between IT and PH teams; could be an option at CERN.
- But scalability and verifiability of CernVM distribution for widespread use as remote batch image is far from evident.
  - Not excluded, but more likely after successful experience with static images.

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Step 4

- Distributed virtual image includes client to connect directly to experiment pilot job framework (Dirac, PanDA).
- Initially with virtual machine images instantiated according to jobs arriving at sites.
- Later, sites instantiate virtual machines according to observed load and local policy
  - Lots of busy ATLAS machines? Start more...
- Requires some way for pilot job frameworks to know (remaining) lifetime of virtual machine.
  - VM unlikely to be updated (security patches...), so lifetime will be limited.

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Step 5

- Experiment pilot job frameworks replaced by commercial/public domain schedulers.
  - Virtual LSF cluster for ATLAS
  - Virtual SGE cluster for CMS
  - …
  - …

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Not vision, but...

- ALICE jobs, at least, require outgoing network connections.
- CERN will provide public (routed) IP addresses for virtual machines. This option may not be available to all...
- IP masquerading also tested at CERN
  - Internal traffic routed, external traffic NATted by the hypervisor.
  - Works, but no way to initiate remote connection to a virtual machine.
    - Should not be a problem based on previous statements.
- Would be good if need for outgoing connections could be removed...
  - Or at least requirements documented per experiment.

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it