Exploiting Virtualization for Software Preservation

DPHEP Topical Workshop on "Full Costs of Curation”
Pere Mato for the PH-SFT group at CERN

14/01/2014

LTDP Challenges

- "Bit preservation" is probably not very difficult
  - You need to ensure that data is migrated from a media/format to a more recent media/format at regular basis
  - Easily findable, fully usable by designated communities with clear access policies
- Data interpretation and knowledge preservation is much more difficult
  - HEP data is, in general, not easy to understand by people outside the immediate community, specially true for unprocessed data (e.g. RAW, DST, ...)
    - Details of the detector, details of the experimental conditions, details of the calibrations, details of ...
  - The **experimental software** is essential for the interpretation of data

Knowledge and Software needs "also" to be preserved

Software Preservation Approaches

- Porting Software
  - Keep software alive by porting from system to system, compiler to compiler, etc.
  - Difficult and very expensive in manpower (e.g. migrating CERNLIB to 64 bit, ROOT 3 to ROOT 6, etc.)
  - Validation, validation, validation
- Freezing the Software
  - No maintenance nor software porting effort required
  - Virtualization is the enabling technology
    - Although virtualization still too young to know how it will "age"
  - The experiment software need to be frozen together with the OS
  - A number of services supporting the application will need to be adapted and frozen together with the application
  - Nothing should be left uncontrolled

Main Ingredient: Virtualization

- Virtualization technology (in same cases with emulation) enables the possibility of running the **old experiment software** in an **old OS version**
  - This is not sufficient since current applications are more complex and require a bunch of 'services' for performing their functions
- Performance degradation is not an issue
  - Time and Moore's law will help us, specially for the 'emulation' case
- Virtualization will stay with us for a long while
  - Industry have also need to be also to run 'legacy' systems in virtualized environments
  - Also the main ingredient for Cloud computing
  - Instead of preserving the actual VM images we need to preserve the recipes or the meta-images that will allow us to build the VM images adequate to the current hypervisor flavor/version
- See next presentation on CernVM[FS] Technology

‘Canned’ Production System

- One single VM running the old experiment's software will not be sufficient (or you will have to wait years to get the results)
  - We need to put many VMs working in parallel
  - so, we need to have a sort of workload management system or batch system
  - … and some user interface with some monitoring, etc.
- Old experiment’s software require several services
  - **Data services** - input and output data files
  - **Metadata catalogs** - bookkeeping
  - **Detector conditions service** - full set of calibrations/alignments
  - **Data monitoring and validation** - comparisons to reference histograms
  - etc.
- We need a complete ‘production system’ build from a set connected nodes **isolated** from the world
  - Remember that we will be running old and insecure software+OS

‘Canned’ Production System

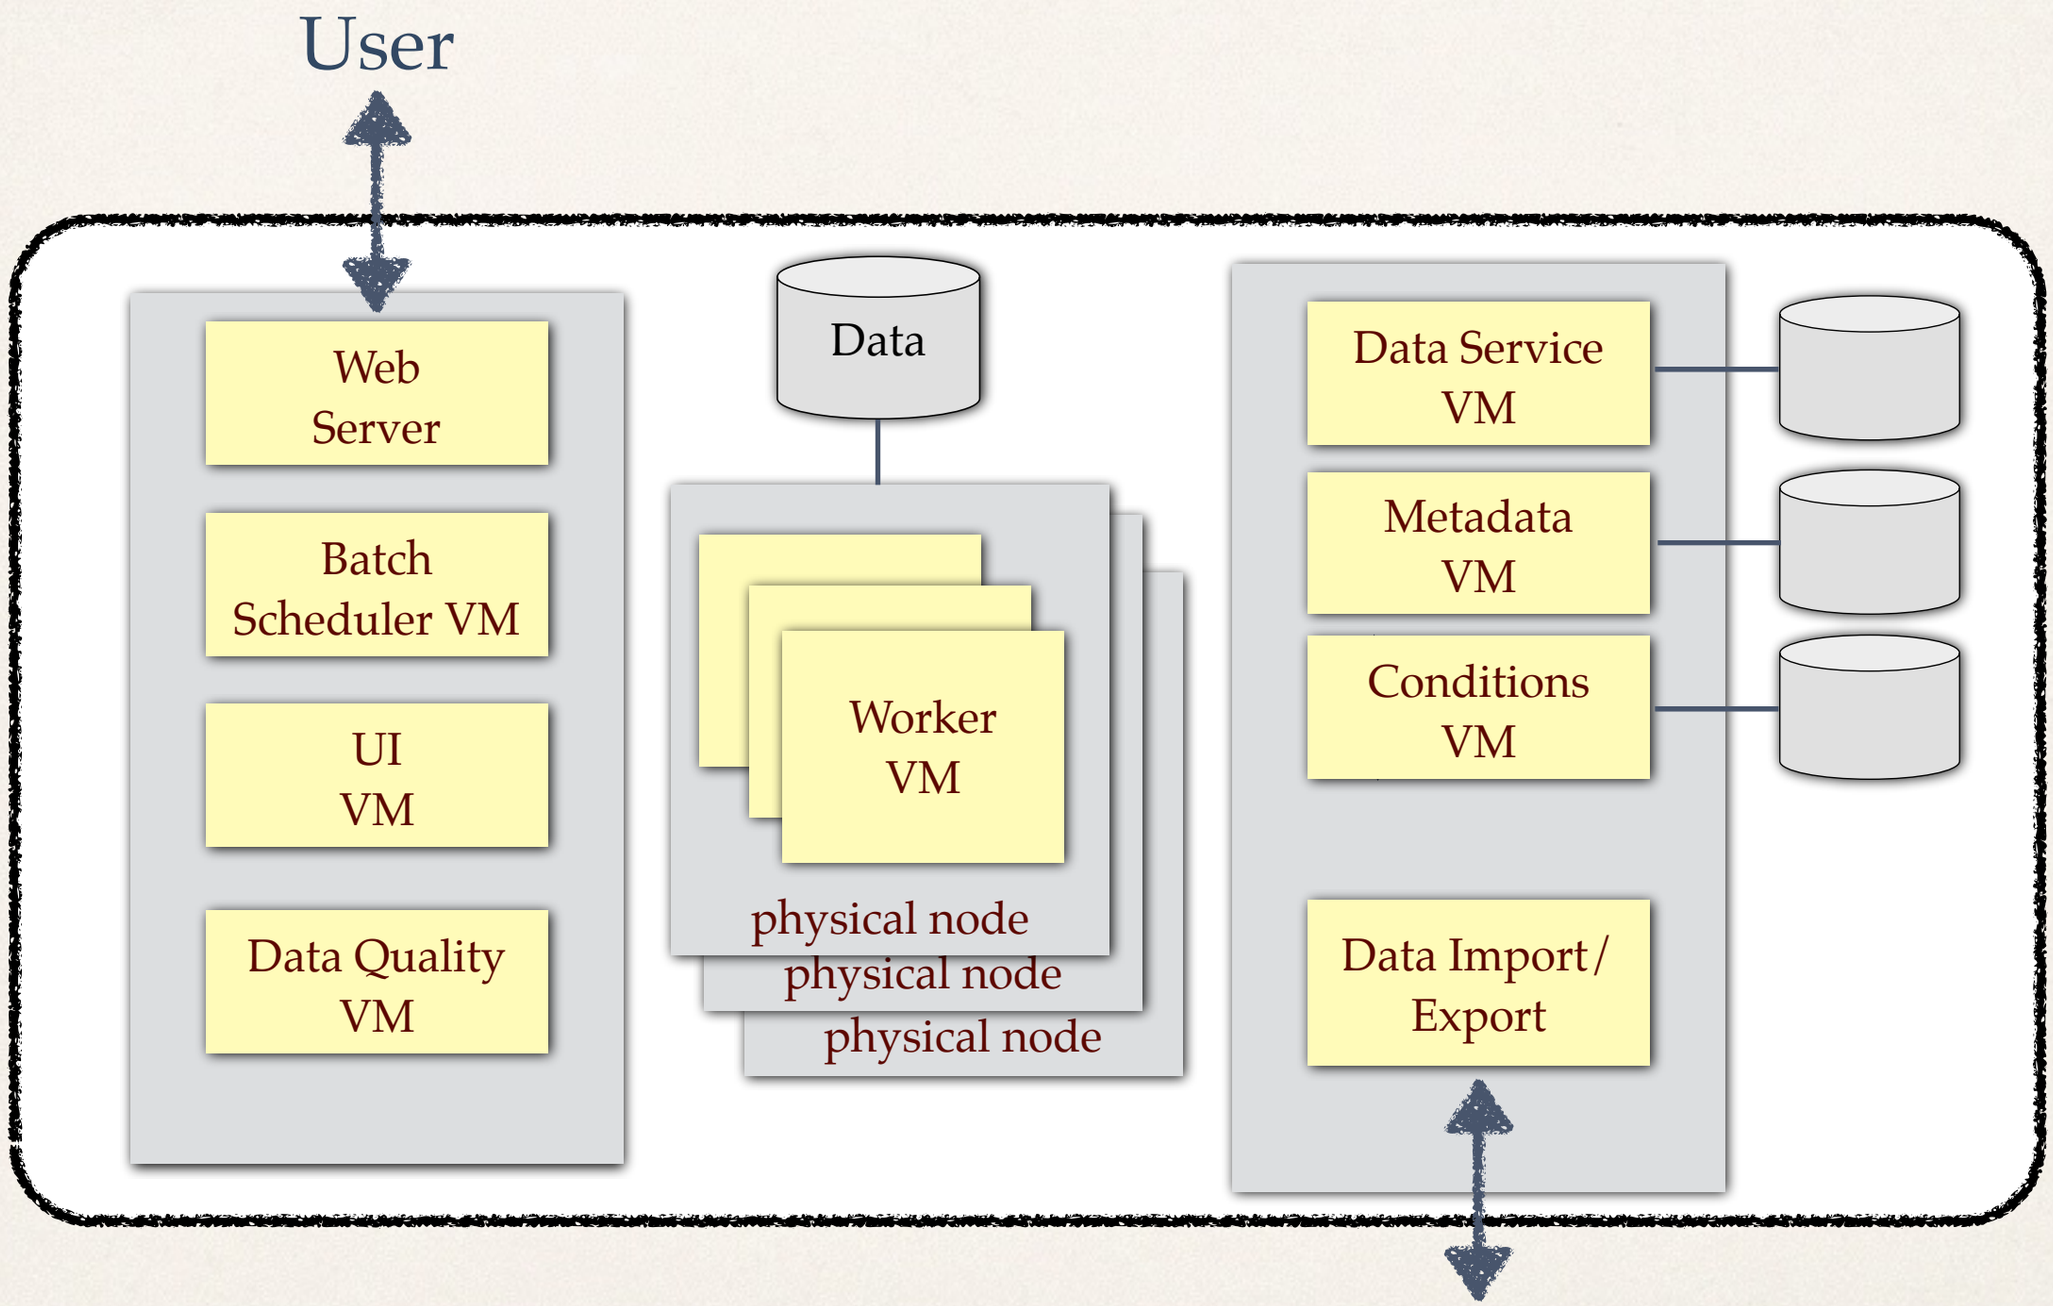

Preservation ‘Exercises’ are needed

- In order to validate the software freezing option we need to perform a number of **controlled experiments to validate main assumptions**
  - Be able to run very old OS versions in new hypervisors
  - Be able to reproduce knows key data analysis and validate the results
  - Run productions (simulation+reconstruction) without any need of an external service
  - Be able to modify old experiment code adding new algorithms and code
  - etc.
- Take very old (~20 years) experiment software (plus compiler, OS, etc.) and resurrect it to perform a number of ‘use cases’
- Learn from the encountered problems
- Provide guidelines to new and running experiments

Conclusions

- Virtualization technology has the potential for providing long-term software persistency by freezing it together with the OS
- Virtualization needs to be accompanied by
  - 'Time Machine' capable file system
  - VM meta-images
  - Canned data production systems
  - etc.
- To validate the technology and the assumptions is essential to design a number of exercises
  - This would allow to validate assumptions and learn what to do for current and new experiments

Cost/Effort Estimates

- Virtualization+FS
  - Service: 1 FTE
  - Software Maintenance: 0.5 FTE
- Canned Data Production
  - Development: 2 FTE * 2 years
  - Maintenance: 0.5 FTE
  - 0.5 FTE * N experiments
- Exercises/Demonstrators
  - 1 FTE * 0.5 year * N exercises