CernVM
SoFTware Appliance

Software Packaging: Can Virtualization help?

Predrag Buncic (CERN/PH-SFT)

CernVM
SoFTware Appliance

Virtualization

- Virtualization was first implemented more than 30 years ago by IBM as a way to logically partition mainframe computers into separate virtual machines (original CERNVM)
- Modern (multi and many core) Intel and AMD CPUs offer support for virtualization (already 2nd CPU generation)
- Virtual machine
  - tightly isolated software container that can run its own operating systems and applications as if it were a physical compute
- Benefits
  - Compatibility
  - Isolation
  - Encapsulation
  - Hardware independence
- Enabling Cloud computing

CernVM
SoFTware Appliance

Virtualization R&D

- Provide a complete, portable and easy to configure user environment for developing and running LHC data analysis locally and on the Grid independent of physical software and hardware platform (Linux, Windows, MacOS)
  - Code check-out, edition, compilation, local small test, debugging, …
  - Grid submission, data access…
  - Event displays, interactive data analysis, …
  - Suspend, resume…

- Decouple application lifecycle from evolution of system infrastructure
- Reduce effort to install, maintain and keep up to date the experiment software
- Web site: http://cernvm.cern.ch

From Application to Appliance

Starting from experiment software…

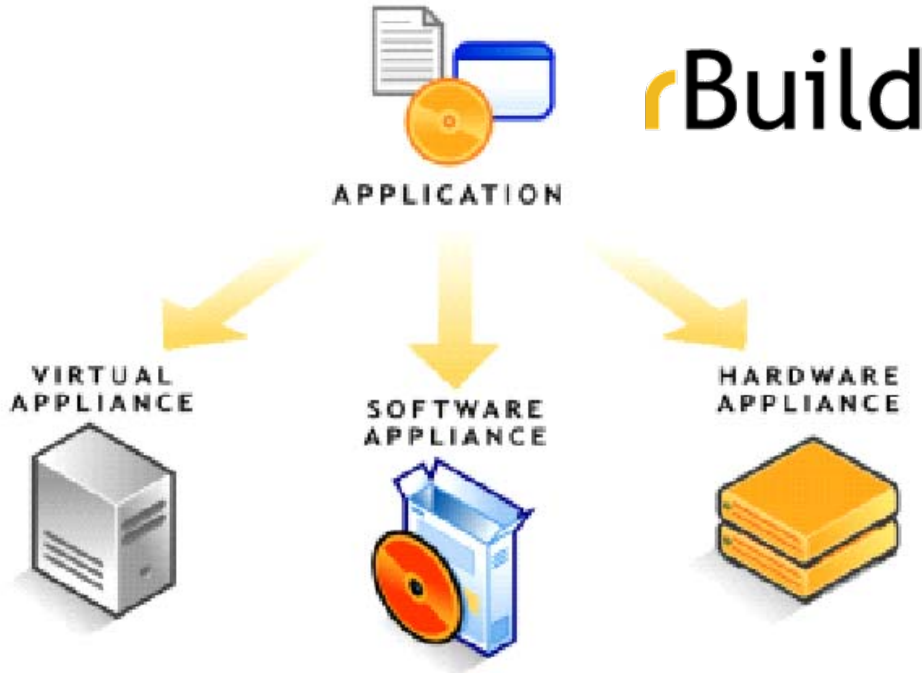

…ending with custom Linux specialised for a given task

Build types

- Installable CD/DVD
- Stub Image
- Raw Filesystem Image
- Netboot Image
- Compressed Tar File
- Demo CD/DVD (Live CD/DVD)
- Raw Hard Disk Image
- Vmware ® Virtual Appliance
- Vmware ® ESX Server Virtual Appliance
- Microsoft ® VHD Virtual Apliance
- Xen Enterprise Virtual Appliance
- Virtual Iron Virtual Appliance
- Parallels Virtual Appliance
- Amazon Machine Image
- Update CD/DVD
- Appliance Installable ISO

CernVM
SoFTware Appliance

“Thin” Virtual Machine

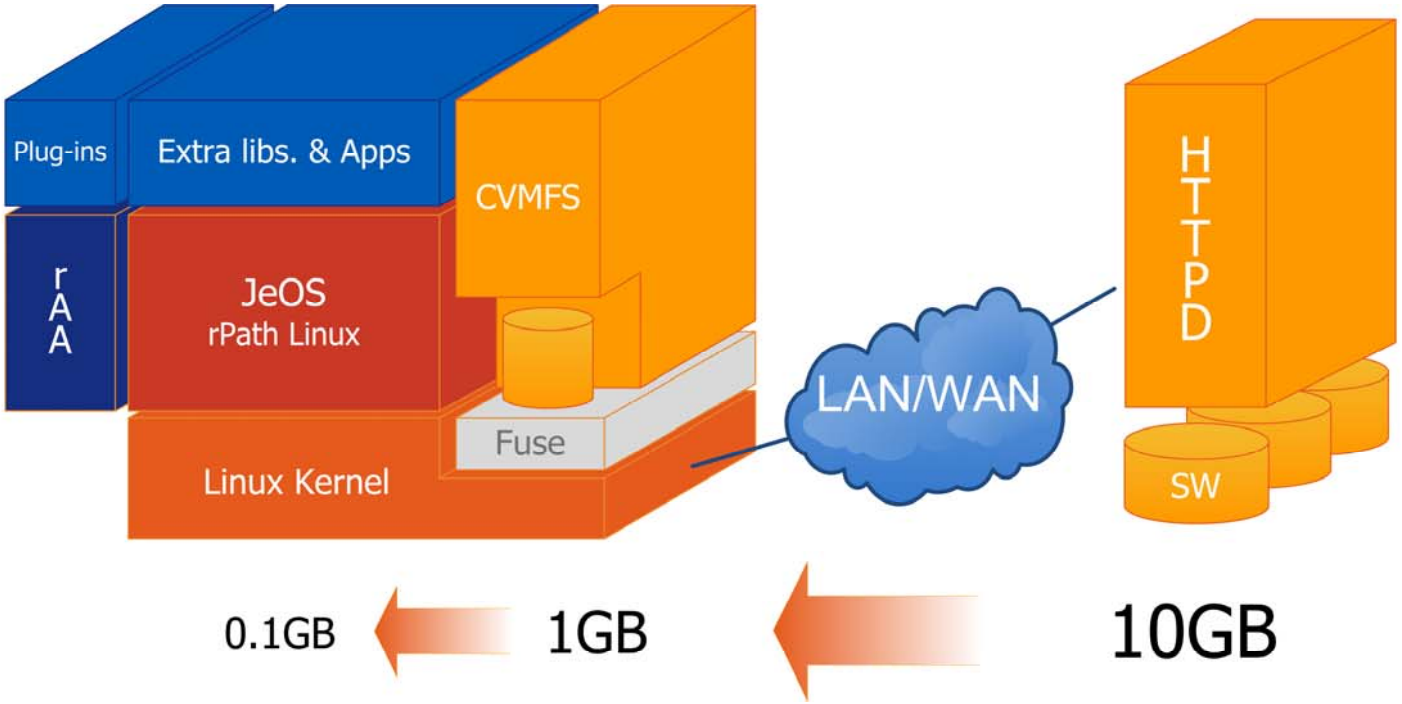

- The experiment are packaging a lot of code
  - but really use only fraction of it at runtime
- CernVM downloads what is needed and puts it in the cache
  - Does not require persistent network connection (offline mode)

CernVM
SoFTware Appliance

Publishing Releases

1. Each experiment is given a VM to install and test their software using own installation tools

2. Publishing is atomic operation

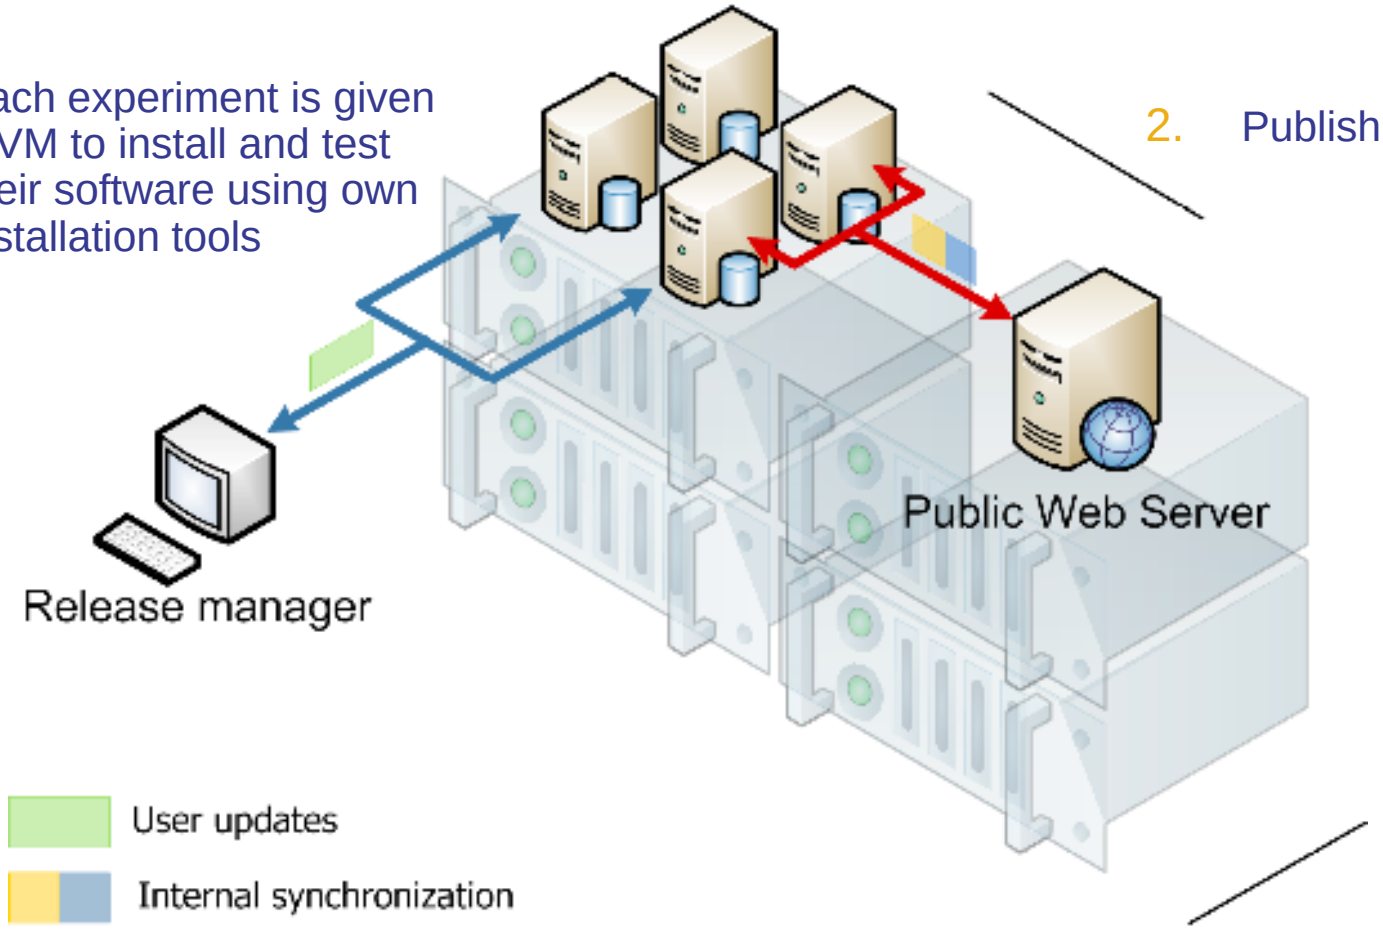

CernVM
SoFTware Appliance

As easy as 1,2,3

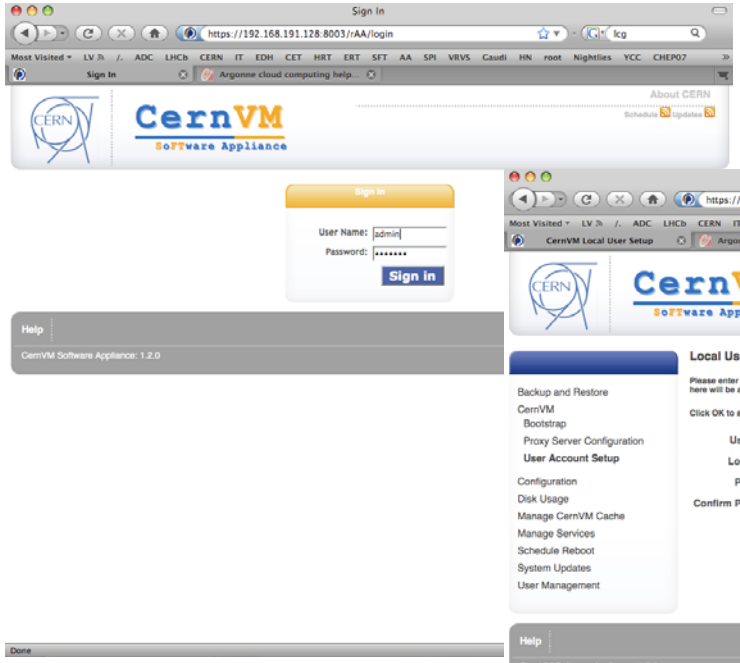

1. Login to Web interface

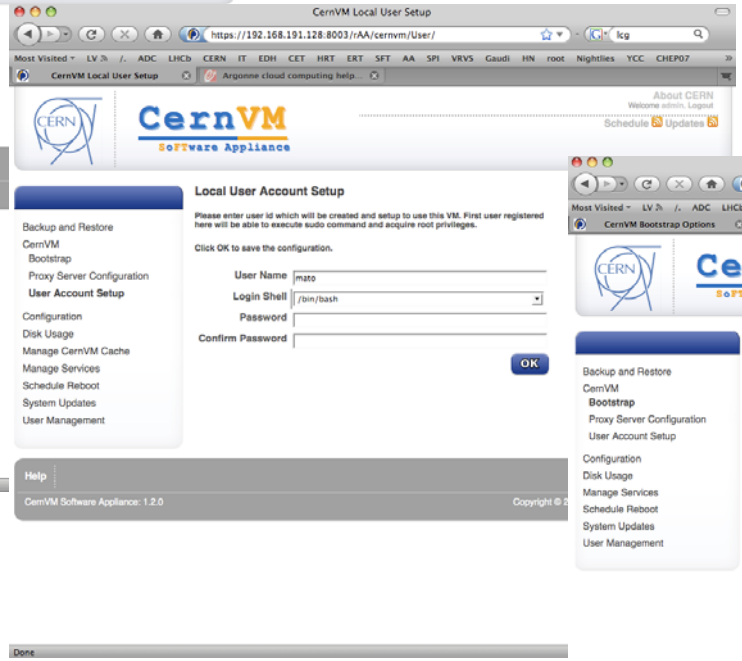

2. Create user account

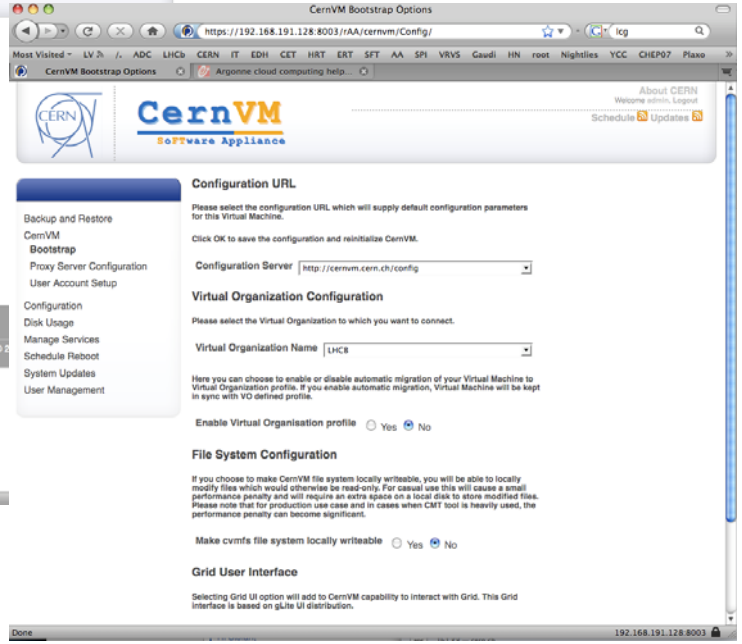

3. Select experiment, appliance flavor and preferences

1.3.0 Development release

- First attempt to include LC software in CernVM
  - Thanks to André Sailer and Christian Grefe
- Available now for download from
  - http://rbuilder.cern.ch/project/cernvm-devel/releases
- Can be run on
  - Linux (VMware Player, VirtualBox)
  - Windows(VMware Player, VirtualBox)
  - Mac (Fusion, VirtualBox)
- Appliance can be configured and used with ALICE, LHCb, ATLAS, CMS and LCD frameworks
- This release comes in two editions
  - Basic (text development environment, suitable for ssh login, ~250MB)
  - Desktop (full desktop environment, works on VMware & VirtualBox, ~500MB)

Screenshots

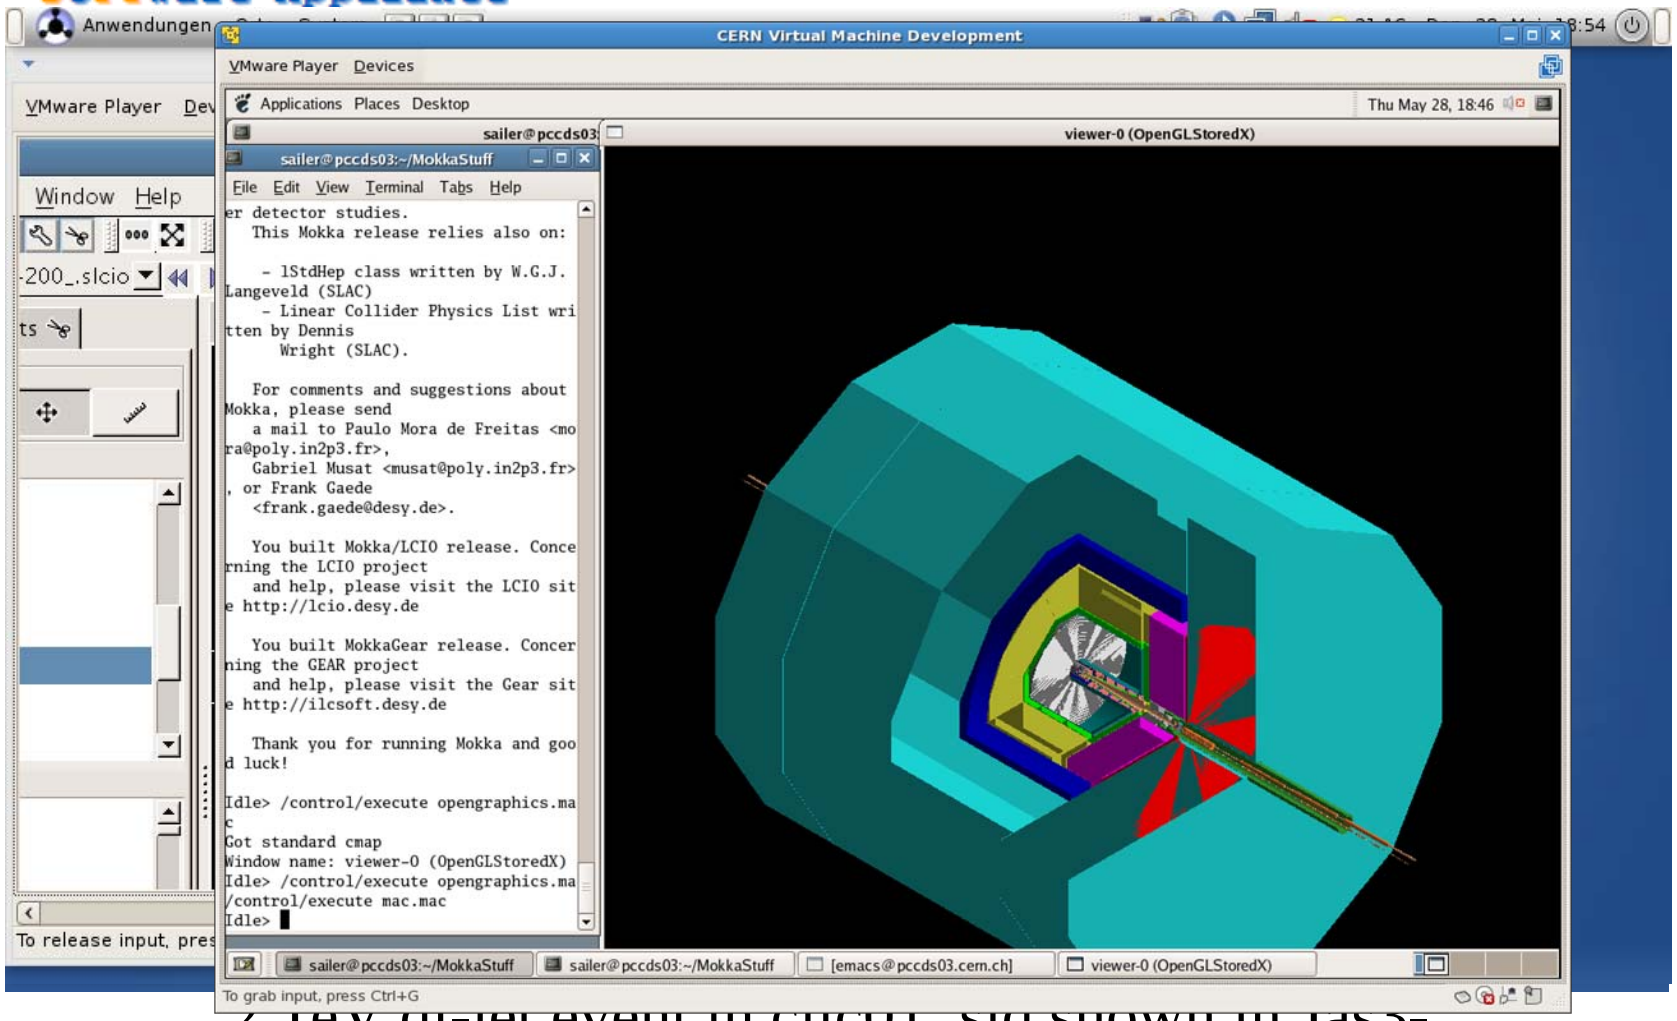

Drawing of clic01_sid.in using Geant4 visualization

2 TeV qq-jet event in clic01_sid shown in Jas3-wired event display

ILD type Detector Simulated in Mokka Software

CernVM
SoFTware Appliance

Conclusions

- Lots of interest from LHC experiments and huge momentum in industry
  - Hypervisors are nowadays available for free (Linux, Mac and Windows)
- CernVM approach solves problem of efficient software distribution using a special file system
  - One image fits all
- What is this good for?
  - Performance penalty ~5% (~1% with the latest CPU generation)
  - To develop and test your code on your local desktop/laptop without having to worry about installation and updates of software framework
  - Grid User Interface
  - Compatible with Cloud
- Beware
  - There will always be performance penalty
  - Hypervisors and CernVM are still in development
  - Not yet clear how to deploy virtual machines as batch/grid job containers