The use of CernVM and CernVM-FS for Software Environment Preservation

J Blomer, G Ganis

CERN EP-SFT

3 February 2016
DPHEP workshop, ISCTE-IUL, Lisbon, Portugal

An ecosystem for virtualization in HEP

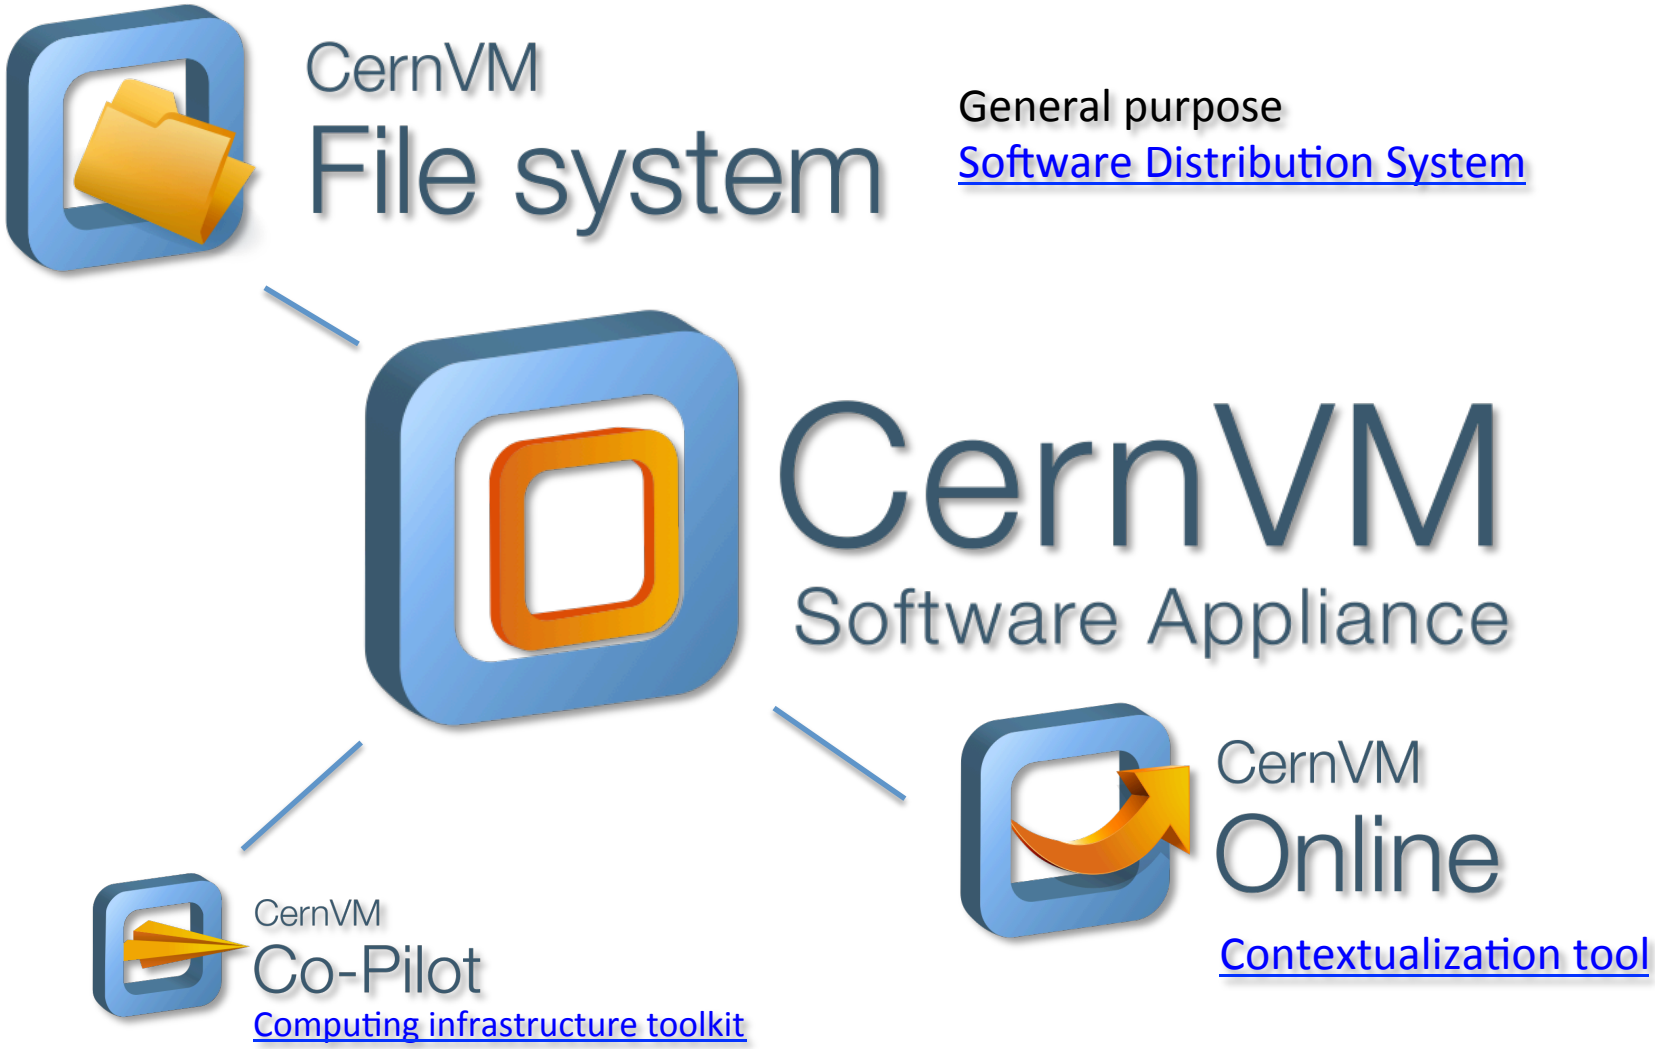

CernVM and Long-Term Analysis Preservation

- Solution to process data w/ any given software version and regenerated VM developed
  - CHEP 2012, DPHEP 2013, CHEP 2013 (NA61 test)
- Main ingredients
  - CernVM-FS embedded versioning (à la time machine)
  - Well defined and versioned *recipes* to regenerate VMs

Components of the CernVM blueprint

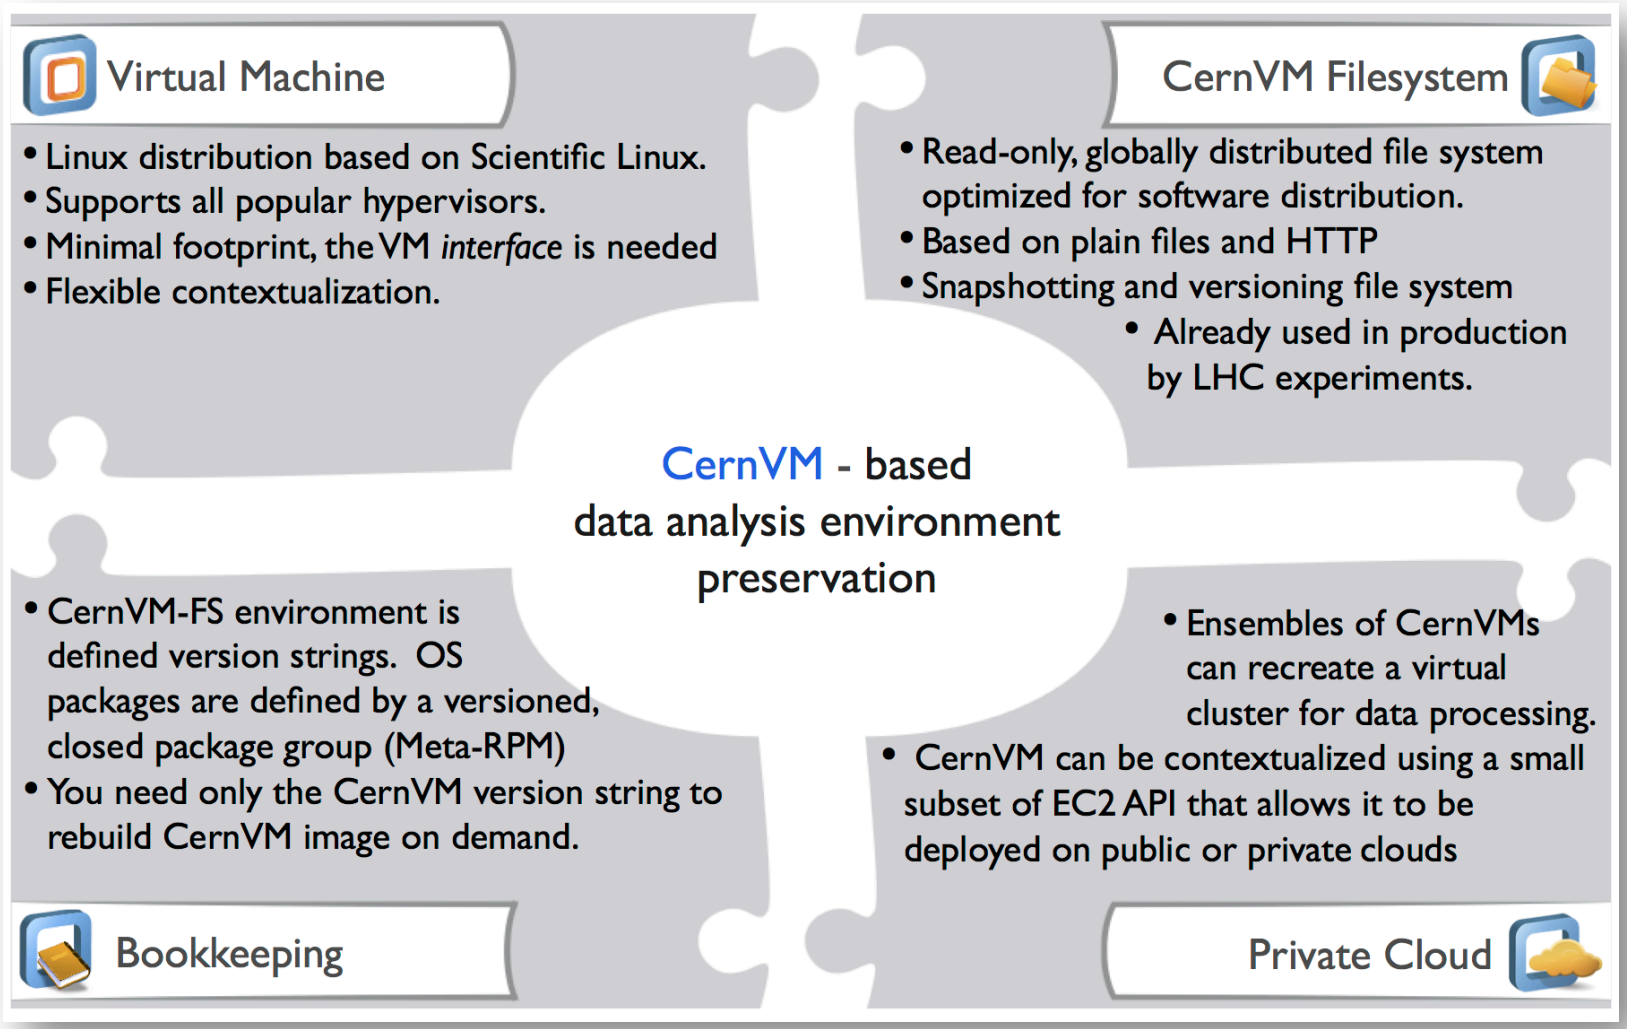

CernVM and Long-Term Analysis Preservation

- Solution to process data w/ any given software version and regenerated VM developed
  - CHEP 2012, DPHEP 2013, CHEP 2013 (NA61 test)
- Main ingredients
  - CernVM-FS embedded versioning (à la time machine)
  - Well defined and versioned *recipes* to regenerate VMs
- Novel *bootloader* technology (CernVM v3 and higher) opened new ways
  - Operating System template on CernVM-FS
    - Versioned and preserved as experiment software
  - Potential to recreate any older Linux environment
  - Support / Integration with container / Docker technology

DPHEP 2015

"The Need for a Versioned Data Analysis Software Environment", 2nd (WSSSPE2), 2014

The *bootloader* technology

µCernVM bootloader (≈20 MB) + OS delivered by CernVM-FS (≈100 MB)

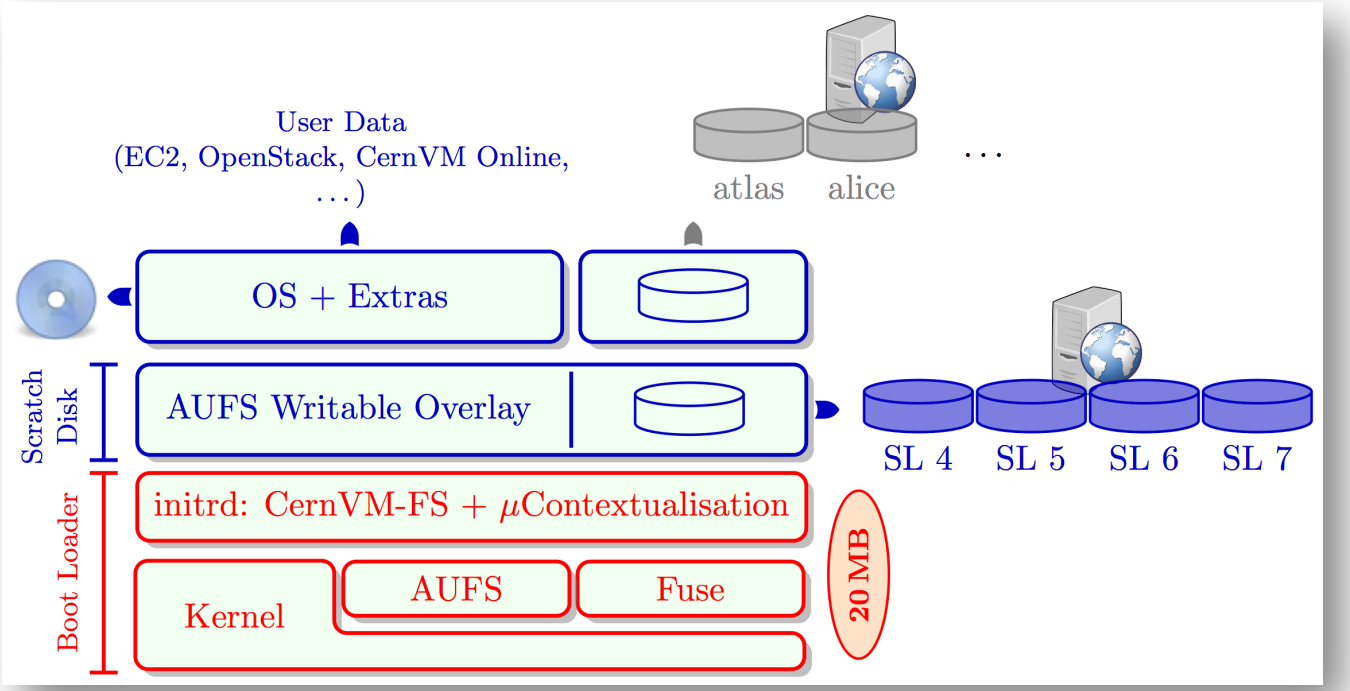

J Blomer et al., "Micro-CernVM: slashing the cost of building and deploying Virtual Machines", 2014 J Phys Conf Ser 513 032009 (CHEP 2013)

G Ganis et al., "Status and Roadmap of CernVM", 2015 J Phys Conf Ser 664 022018 (CHEP 2015)

The *bootloader* technology (2)

- Stable, Long-Term Supported Linux kernel
  - 3.10, 3.18, 4.2
- Operating System templates on CernVM-FS
  - Installed with OS package manager
- After boot, basically a 'chroot' to the chosen template
  - Writable area provided by a overlay FS (AUFS)
  - Like in *live CD* trial installations
- Possibility to choose a flavor and its version
  - Exploiting CernVM-FS versioning features
- Default template SL 6
  - Available: SL 4, SL 5, SL 7 (preview)
  - Not limited to RedHat-based systems, in principle

Demonstrators for Software Environment Preservation

- ALEPH
  - Software was last validated on Scientific Linux 4
  - Dedicated VM and containers
- CMS Open Data Pilot
  - Released data were taken in 2011
  - Requires Scientific Linux 5 from same period
- CernVM 3 is an enabling technology for the CERN Open Data Project
  - http://opendata.cern.ch

ALEPH: regenerating SLC4

```
/cvmfs/cernvm-slc4.cern.ch                    (OS template)
/cvmfs/aleph.cern.ch                          (ALEPH software)
/cvmfs/sft.cern.ch/lcg/external/cernlib       (CERNlib)
```

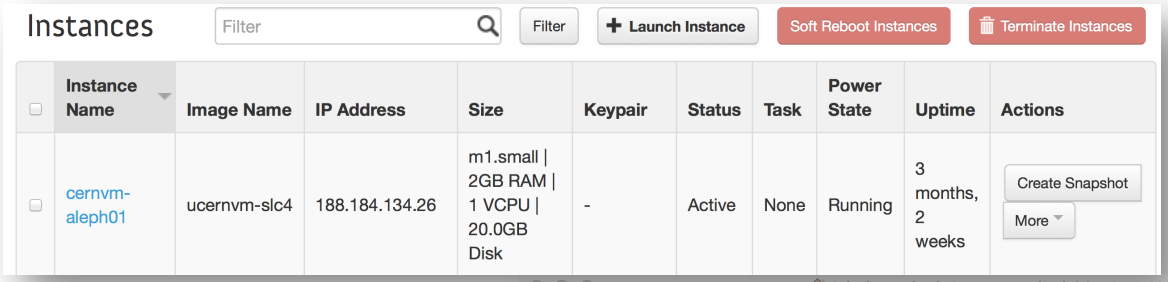

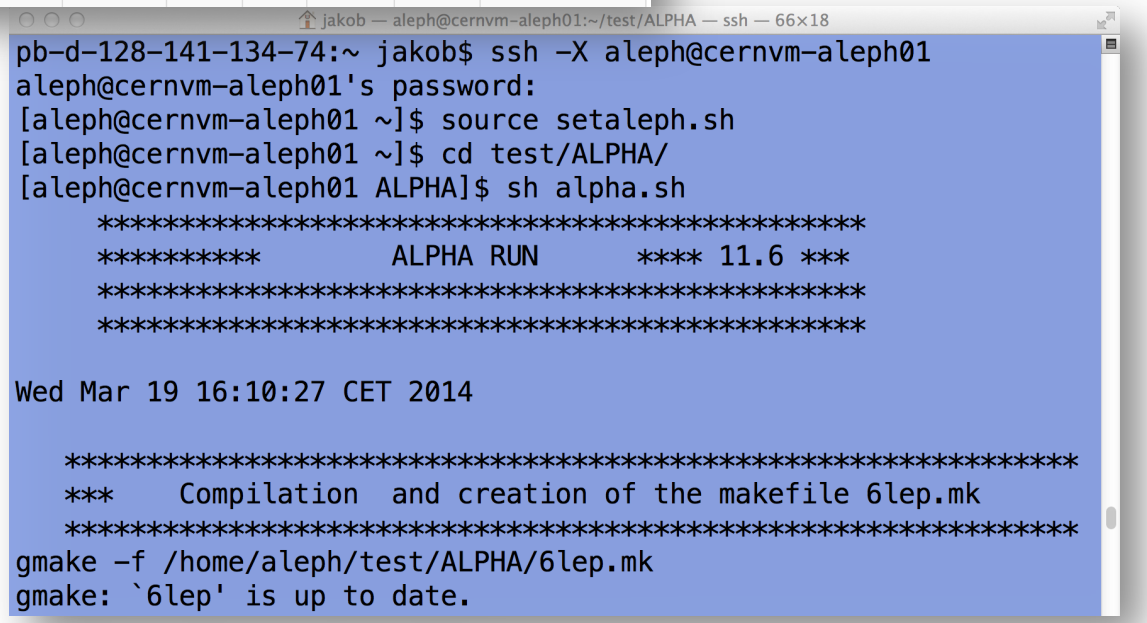

cernvm-aleph01
VM machine on CERN openstack

CMS Open Data Pilot context

Provide an easy-to-use virtual machine with CMS computing environment for CMS Open Data

- Data:
  - Frozen set, access via XRootD, DPHEP portal
- Software:
  - Frozen CMSSW version (4.2.8.patch7)
- SL 5 compatible VM
  - Full Analysis environment: compile + run
  - Graphical environment, easy-to-install, easy-to-use

CMS Open Data Pilot: deployment

```
/cvmfs/cernvm-slc5.cern.ch          (OS template)
/cvmfs/cms.cern.ch                  (CMS software)
```

- Deployed as OVF/OVA bundle[1]
  - Open specification for bundling VMs, stable since 2009
  - OVA: tarball with a hard disk image and an XML specification

Easy auto-installation in VirtualBox

Same solution for LHCb @ OpenData http://opendata.cern.ch

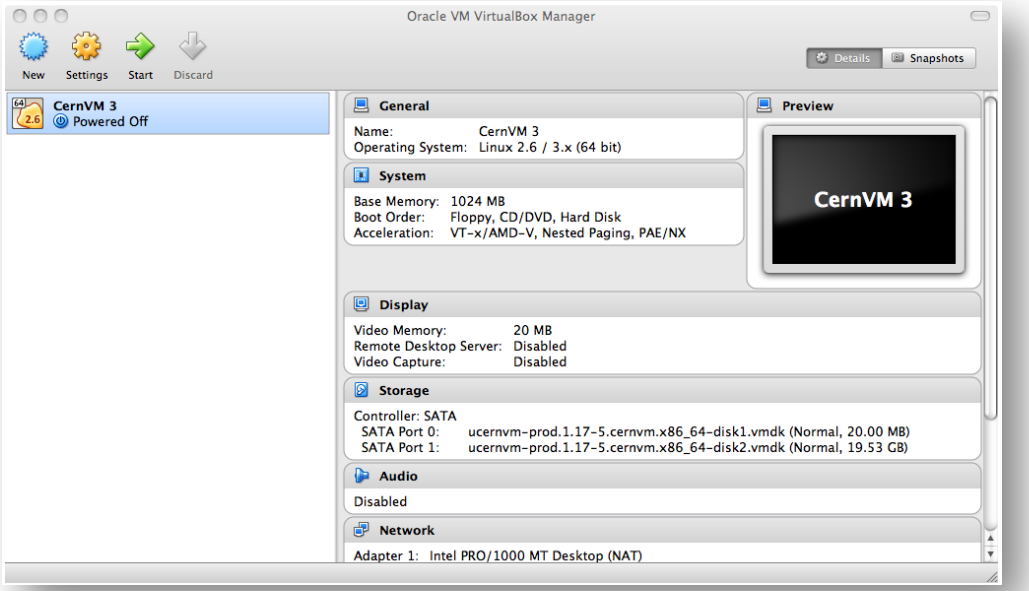

[1] Open Virtualization Format / Open Virtual Appliance: http://www.dmtf.org/standards/ovf

Containers

- Lots of interest, seen as way to optimize use of resources
- Kubernetes, Mesos, Openstack, …
- All LHC experiments gave a try
  - Basic approach: pack everything in the container
    - Huge containers
  - ALICE investigating the concept of pilot container
  - CernVM-FS efficient integration/support essential

Containers and Docker

CONTAINER

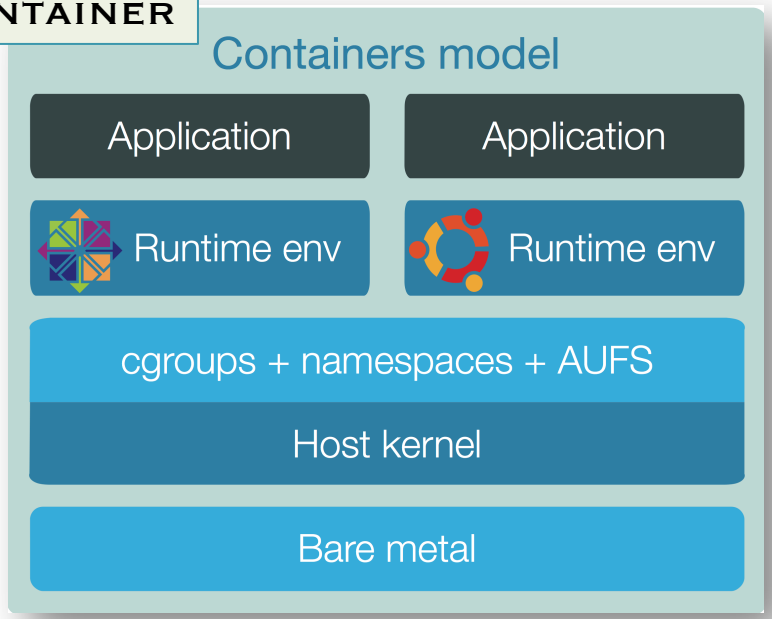

VIRTUAL MACHINE

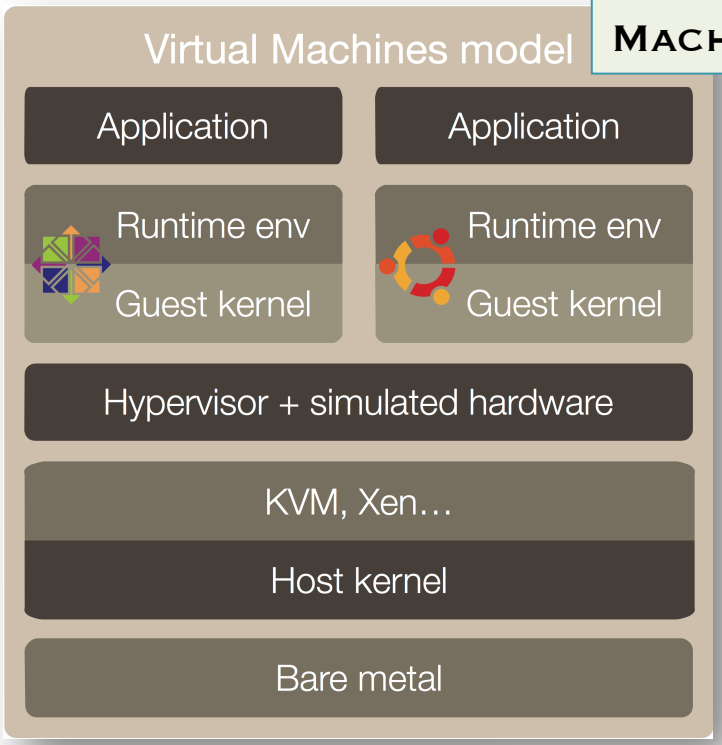

https://www.docker.com

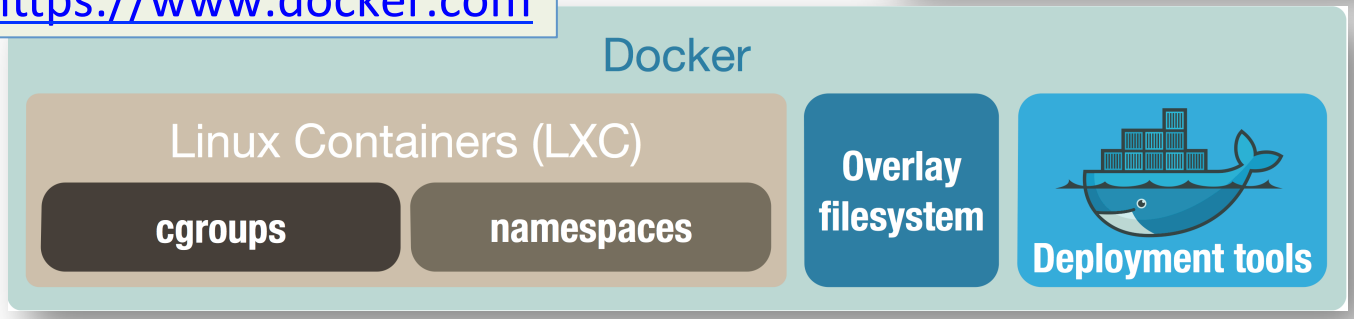

DOCKER

Containers within VM

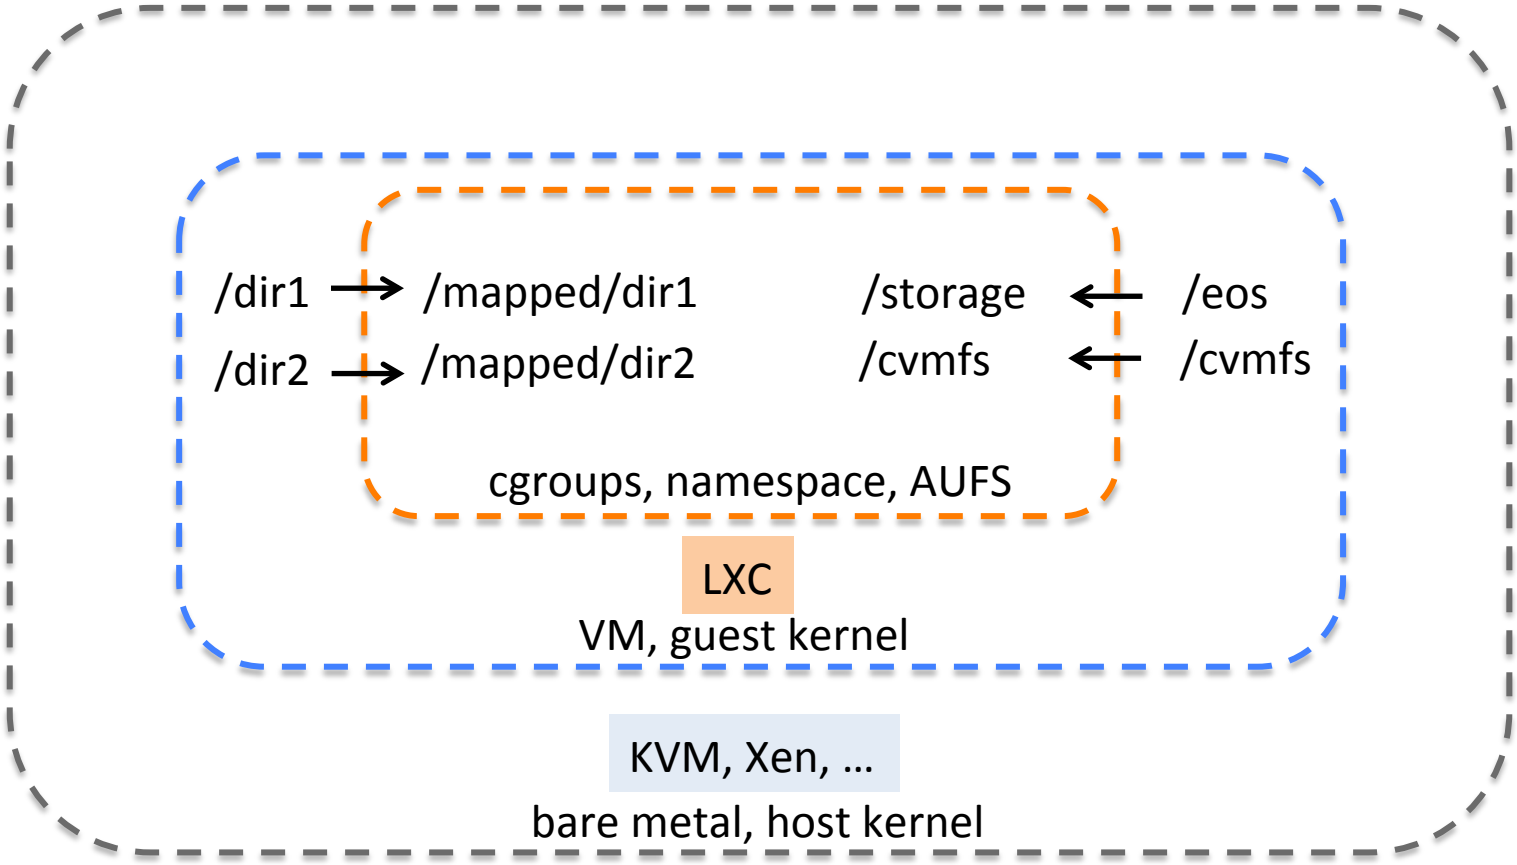

CernVM-docker SLC4 for ALEPH

Details at http://cernvm.cern.ch/portal/cvm3-slc4-docker

```
docker run -i -t \
     -v /cvmfs/cernvm-slc4.cern.ch/cvm3:/cernvm \
     -v /cvmfs/aleph.cern.ch:/cvmfs/aleph.cern.ch  \
     -v /cvmfs/sft.cern.ch:/cvmfs/sft.cern.ch \
     -v $HOME/local/aleph/MIT:/MIT \
      ucvm-slc4-docker /init
```

Bind mount relevant volumes from /cvmfs, /eos, …

Also working area, for convenience

/init contains instructions how to setup the system

```
$ ./local/etc/start-slc4.sh
**************************************
*** Welcome to SLC4 in CernVM+Docker ***
**************************************
```

```
root@ucvm-slc4:/MIT $ source aleph-env.sh
root@ucvm-slc4:/MIT $ ./Compile.sh
root@ucvm-slc4:/MIT $ ./MIT
 ------------------------------------------------
                A L P H A   126.22   14.52.47 07/06/15
                          running on LINUX
                       ALEPHLIB Version  316.20
 ------------------------------------------------
0_QUIBOS_ Init BOS with  5000000 words working space
  *** ALFMT*** BANKAL.FMT on unit  77 contains  1331 bank names
0BREADC-----  NAME........NR  +------------------------------+
```

Next steps

- Further demonstrate on a cloud (Openstack) the setup of a fully functional cluster farm for ALEPH
  - Data from EOS
  - Real QCD event shapes analysis
- Implement some provenance
  - Record changes between versions, …
  - Contacts with CRISTAL (J Shamdasani, A Branson)

Summary

- CernVM ecosystem: natural potential for software environment preservation
  - CernVM-FS embedded versioning of SW and OS
- The *bootloader* technology has simplified the way we can regenerate the target VM
  - Demonstrated with ALEPH and CMS Open Data Pilot
- Container technology very promising
  - Potential to add smoothness to the whole process
  - Demonstrating with ALEPH/SLC4
- On the way for a full provenance system