CernVM: Towards a Containerized Compatibility Layer

Jakob Blomer

CernVM Virtual Workshop 2021

1 February 2021

The CernVM Virtual Appliance at a Glance

CernVM (Container or VM)

- Curated Linux platform with all the dependencies to run LHC applications
- RHEL compatible
- "Batteries included": ready for most IaaS clouds
- Strongly versioned: `/cvmfs/cernvm-prod.cern.ch/.cvmfs/snapshots/cernvm-system-4.1.0.0`
- Minimal image ($\sim$ 20 MB)
- Graphical (development/outreach environment) and batch flavors
- Customizable through contextualization
- Launcher for easy image instantiation with VirtualBox

Singularity-ready root file system:

`/cvmfs/cernvm-prod.cern.ch/cvm3` (EL6, frozen)
`/cvmfs/cernvm-prod.cern.ch/cvm4` (EL7, production)

Reminder: Building Blocks of CernVM

Twofold system: μCernVM boot loader + OS delivered by CernVM-FS

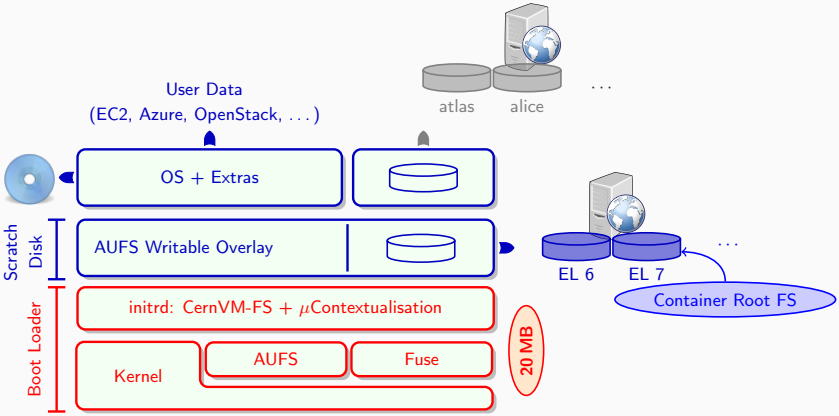

CernVM Full Virtualization Usage Q4/2020

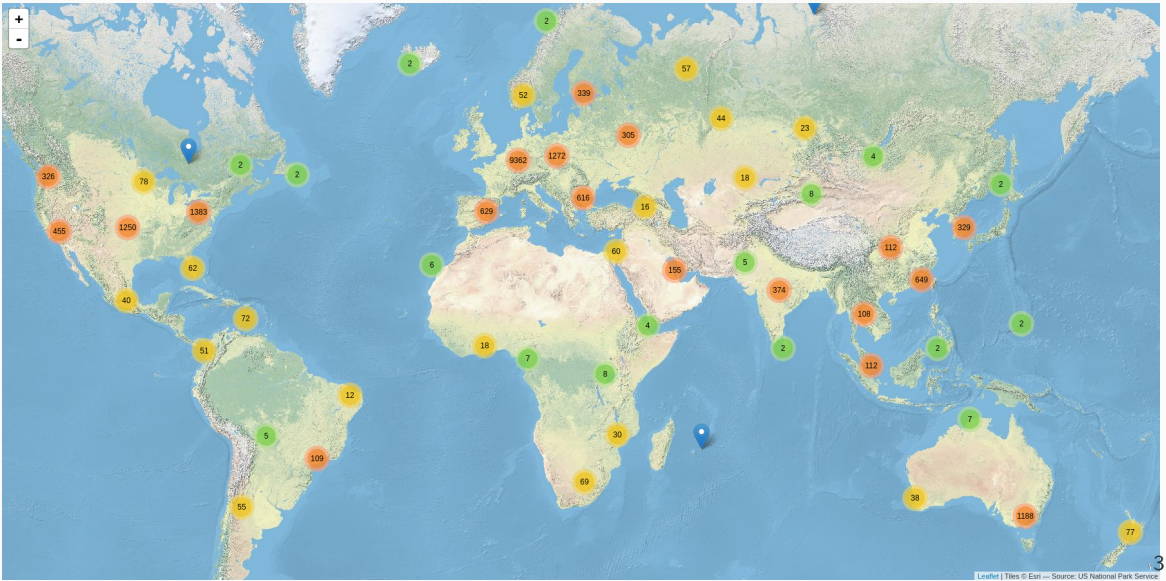

CernVM Full Virtualization Usage Q4/2020

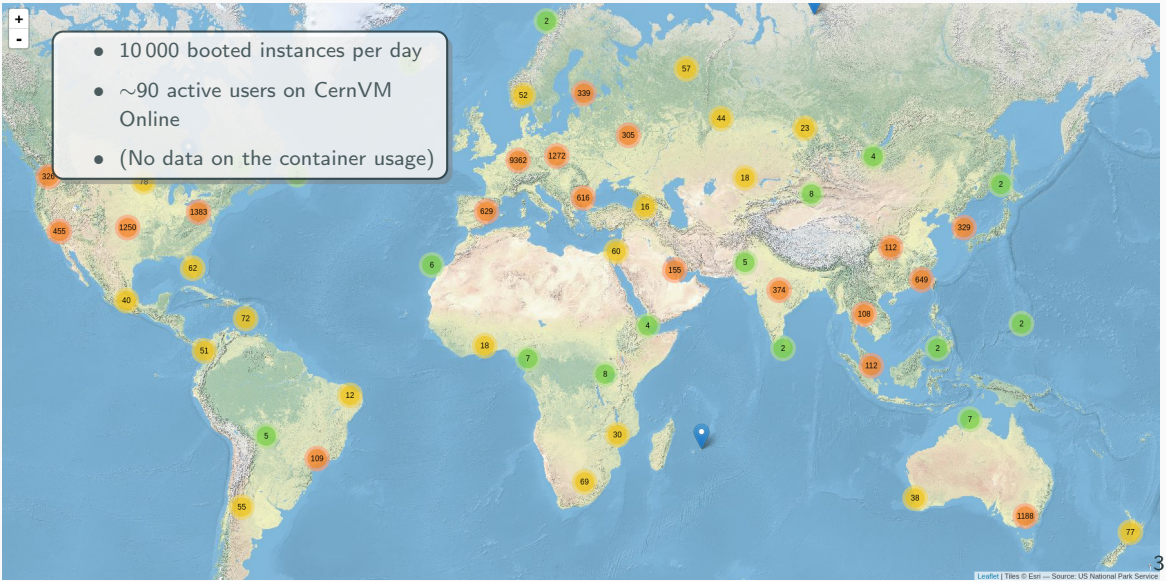

- 10 000 booted instances per day
- ∼90 active users on CernVM Online
- (No data on the container usage)

Current Use Cases I: IaaS Clouds

- ATLAS Sim@P1 on-demand cloud on trigger farm
- Through CloudScheduler and VAC at various academic and commercial clouds

→ Chris' presentation

Current Use Cases II: LHC@Home

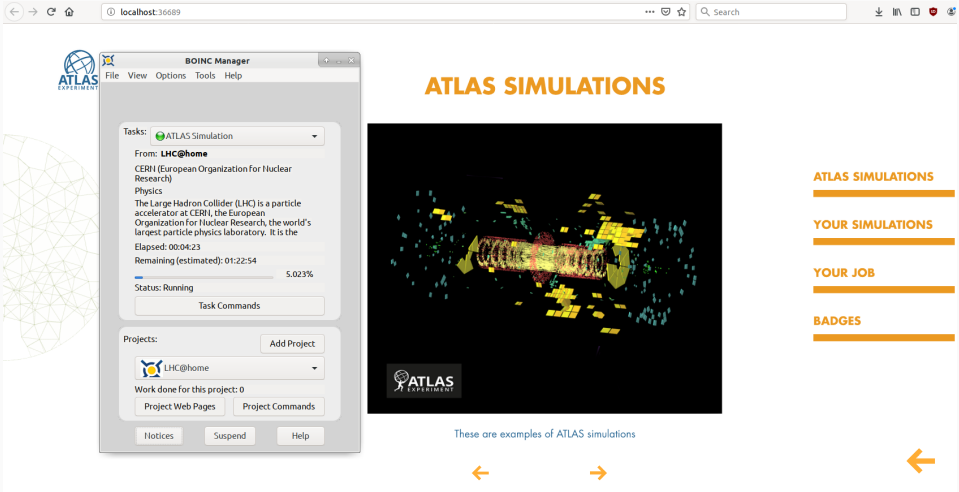

Current Use Cases III: OpenData

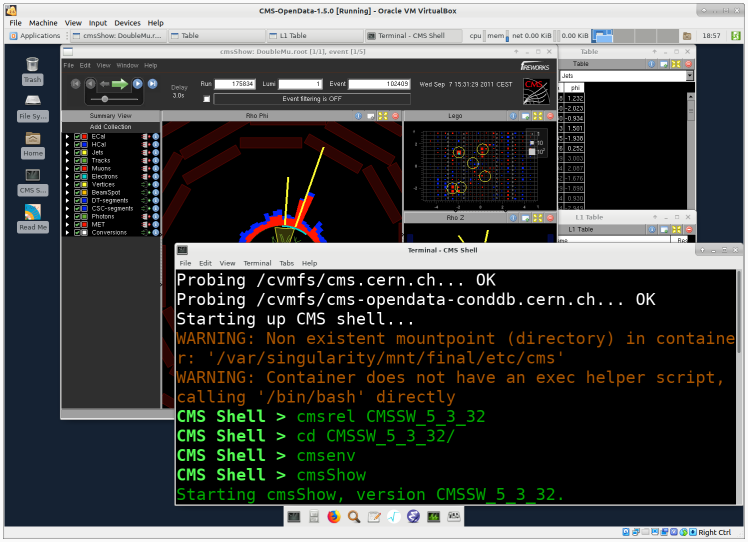

Current Use Cases IV: Data Preservation

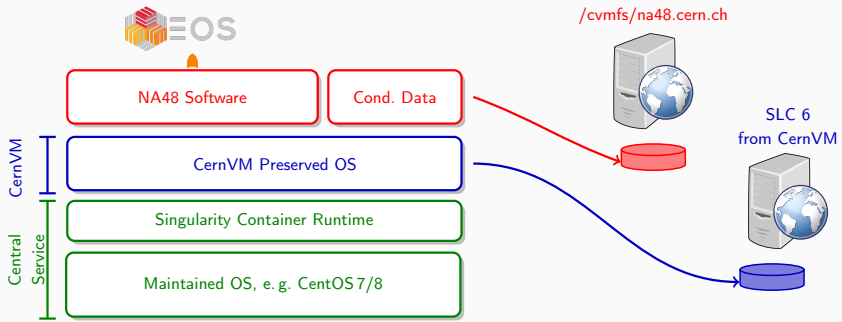


```
[jblomer@lxplus7] singularity exec -B /cvmfs -B /eos \
>     /cvmfs/cernvm-prod.cern.ch/slc6 /cvmfs/na48.cern.ch/my_application
```

Status and Plans

Updates since the 2019 Workshop

- Freeze of CernVM 3 (EL 6 based)
- Minor updates to CernVM 4 and the bootloader to follow upstream updates
- Documentation now available on ReadTheDocs

Plans for the next months

- Maintenance of CernVM 4, bootloader, launcher
- Documentation update to match hypervisor and cloud platform updates
- Updating CernVM Online to the new Single Sign-On system

CernVM 5: Proof-of-Concept

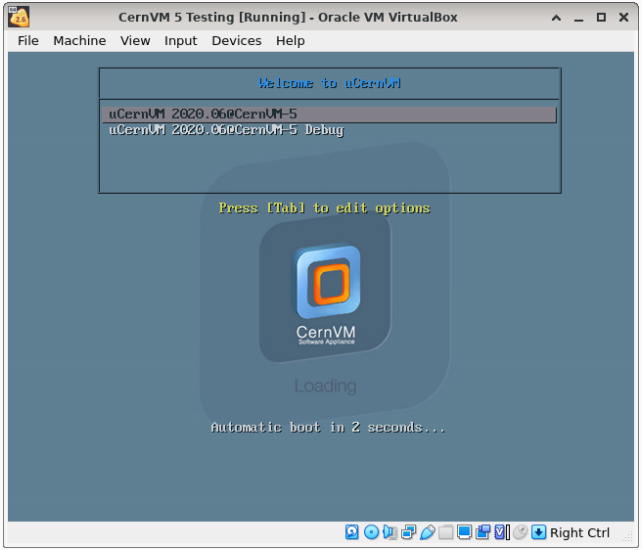

CernVM 5: Proof-of-Concept

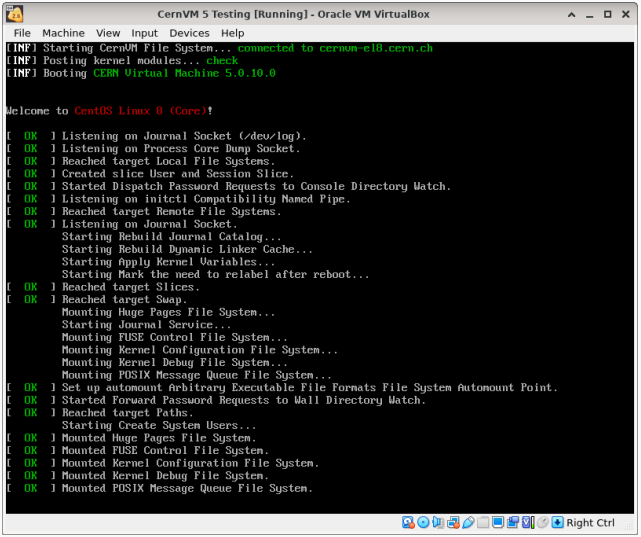

CernVM 5: Proof-of-Concept

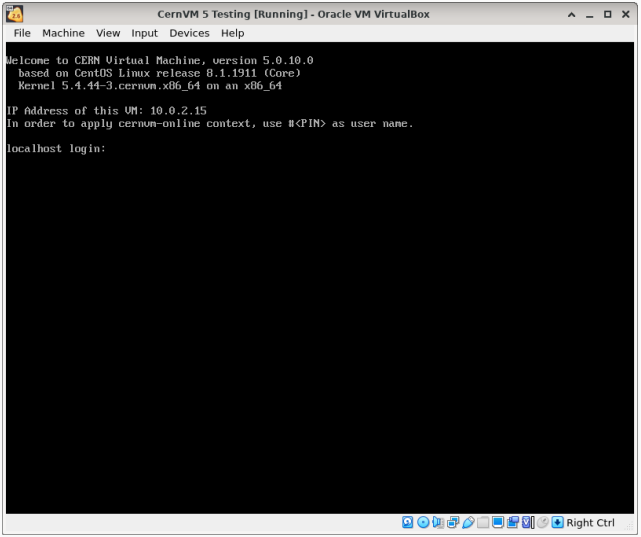

Outlook: CernVM 5

- Back to the "Just enough Operating System" (JeOS) approach would allow us to deploy CernVM as a **virtual machine** for fully virtualized clouds and as a **regular container** for commercial k8s clusters
- Target: < 500 MB CernVM base container, bootable in < 30s

CernVM 5 Layering

Experiment application from /cvmfs

Full app universe provided by /cvmfs

Minimal Compatibility Container
glibc + ε

Bootloader (only for full virtualization)

Minimal container contents (ε)

- CernVM-FS client
- cloud-init
- HepOSLibs (?)

App universe from /cvmfs/sft.cern.ch

- Desktop environment (*Enlightenment*?)
- Python, compilers, ROOT, XrootD client
- ...

Backup Slides

Software and Platform Lifecycle

- Our standard Linux platform, Red Hat Enterprise Linux (aka Scientific Linux, CentOS), has a life time of ∼10 years per release
- No security updates once out of maintenance, hence availability on central services stops (lxplus, lxbatch, . . .)

Two options for experiment application software to manage the operating system change

1. Porting & validation
   can be challenging wrt. legacy dependencies such as CERNLIB
2. Freezing & sandboxing
   Captures legacy software plus OS and compilers using virtualization technology

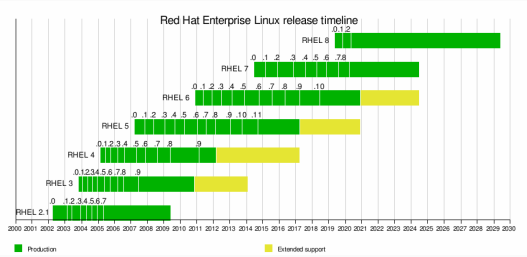

OVA Images

- Image bundled with user data
- Single-click import in VirtualBox
- Can be created on CernVM:

```
cvm2ova \
  -n CMS-OpenData-1.5.1 \
  -i cernvm4-2018.10-1.hdd \
  -u user-data.txt
```

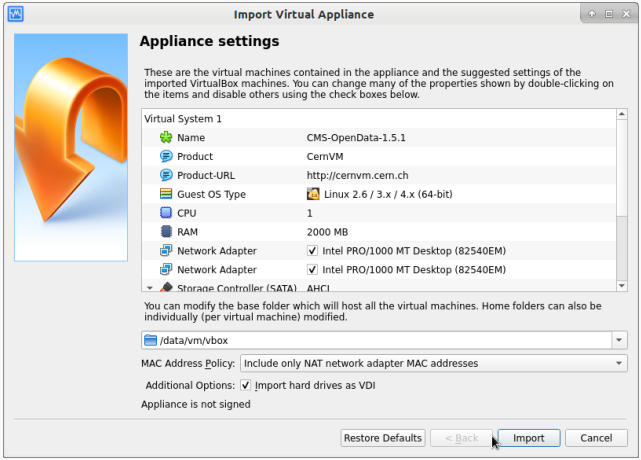

cernvm-launch

`cernvm-launch` command-line utility

- Steers VirtualBox to instantiate CernVMs
- Copyable binary
- Interface similar to Docker and Vagrant
- Requires only the user-data text file ("context")

Implementation

- Uses `libcernvm`, the core of CernVM Online's WebAPI
- Creates the ISO image from context
- Downloads and caches CernVM image from the web
- Calls VirtualBox command line interface

↦ https://cernvm.readthedocs.io/en/stable/cpt-launch.html

↦ https://github.com/cernvm/public-contexts

Example Invocation of cernvm-launch

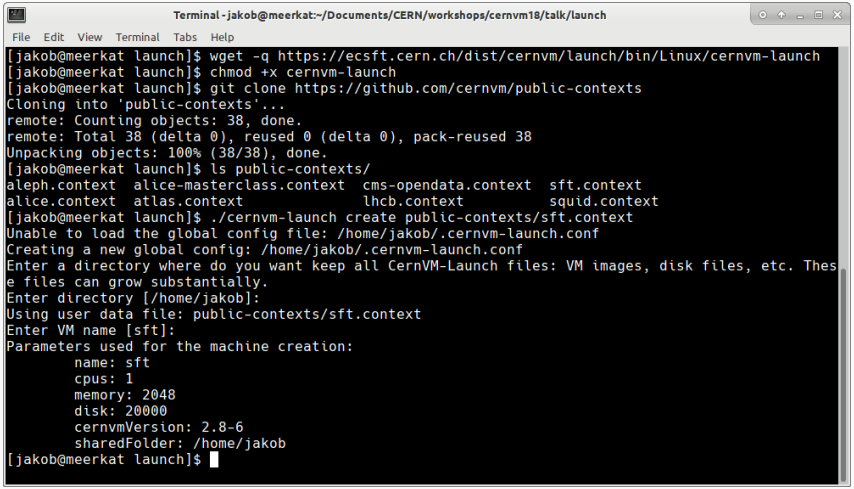

Example Invocation of cernvm-launch