



Status and Plans for CernVM

G Ganis / CERN EP-SFT

For the CernVM team

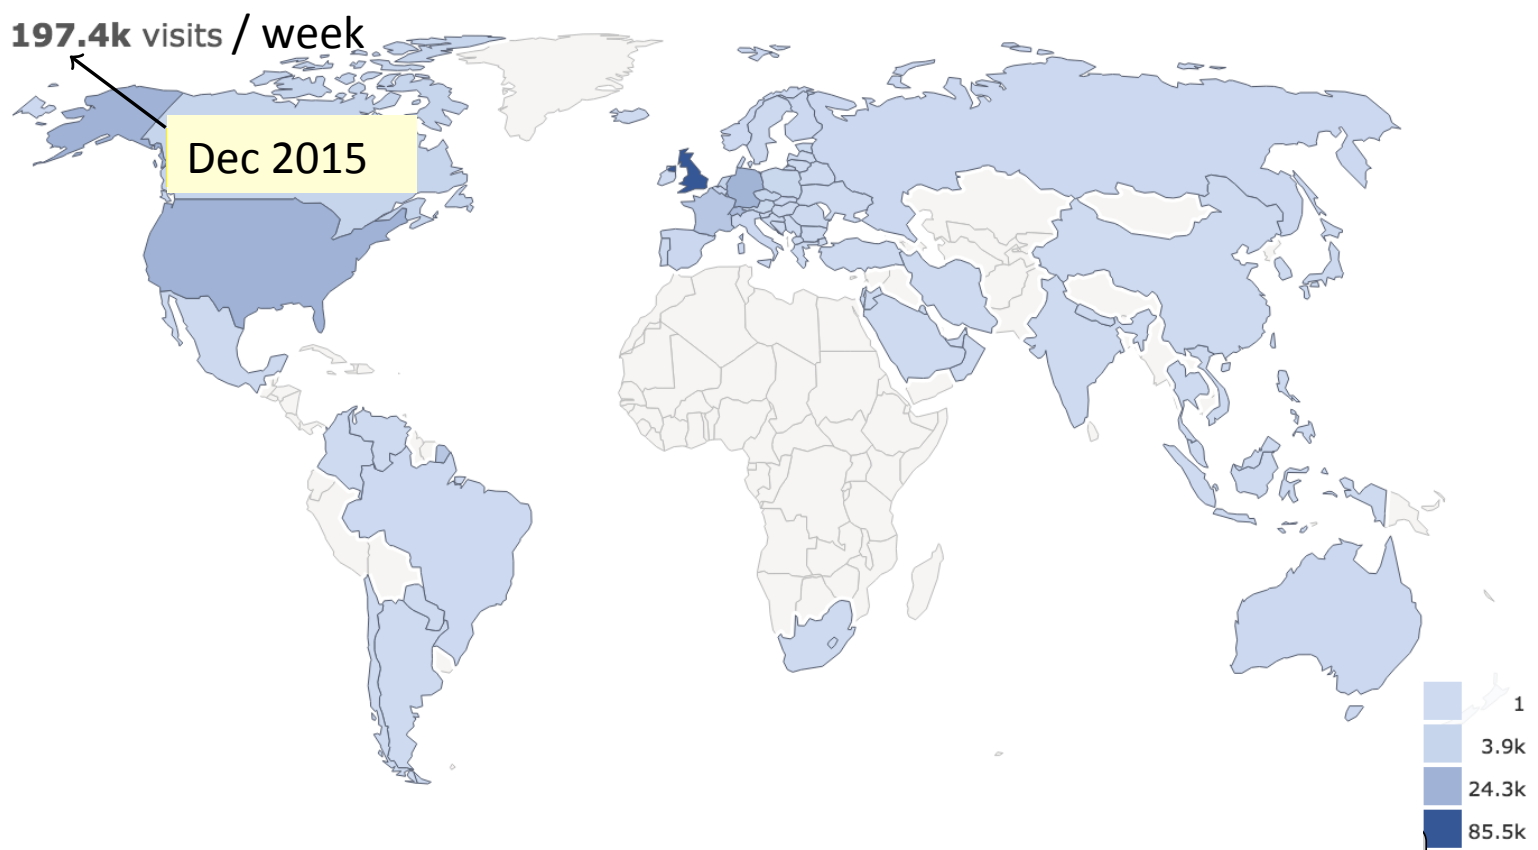
6 June 2016

CernVM Users Workshop, RAL, UK

Agenda

- Deployment overview
- Status
- Developments since the last workshop
 - Containerization
- CernVM 4

CernVM Virtual Appliance



Boot of new CernVM

- ~70000 / day, ~75 countries (~40% UK)

CernVM Virtual Appliance

Complete and portable environment
for **developing** and **running** HEP applications

Use cases:

- Development environment
 - On desktops, laptops
- IaaS Clouds
 - Public, commercial, hybrid
- Volunteer Computing
 - Common LHC framework based on CernVM (L. Field et al.)
- Long-Term Analysis Preservation
- Outreach & Education

CernVM Desktop usage

Interactive usage through [CernVM-Online](http://cernvm.org)

In 2015:
~300 users,
~65 external users

The screenshot shows the CernVM Online dashboard. At the top, there is a navigation bar with links for About, Dashboard, Marketplace, Documentation, Downloads, and Publications. Below the navigation bar is a menu with options: Dashboard, Create Context, Pair an instance, Marketplace, Create Cluster, and Logout. To the right of the menu, the main content area is titled "Dashboard" and "Your context definitions". It displays a table of context definitions with columns for Name, Operations, and WebAPI. The table lists several contexts, each with a "Clone" button, a "Publish" button, a "Withdraw" button, and a "Launch now" button. At the bottom right of the table, there is a "Create new context" button.

Name	Operations	WebAPI
CernVM-SLC4	Clone Withdraw	Launch now
Cvm3-LDT-1	Clone Publish	Launch now
Cvm3-OS	Clone Publish	Launch now
cvm3-test	Clone Publish	Launch now
ecsft	Clone Publish	Launch now
ecsft2	Clone Publish	Launch now
ecsft3	Clone Publish	Launch now
Test-SLC4	Clone Publish	Launch now
TutorialLP	Clone Publish	Launch now
TutorialVM	Clone Publish	Launch now

CernVM IaaS usage

Infrastructure-as-a-Service

Various cloud usage

- ATLAS HLT farm
- Cloud resources seamlessly integrated with experiment task queues (e.g. ATLAS CloudScheduler, LHCb VAC)
- ALICE Software Release testing on CERN OpenStack
- Commercial providers (“Helix Nebula”)
- ...

Used by all LHC experiments in a way or the other, though not exclusively

CernVM in Volunteer Computing

Pioneered use of virtualization in VC with Test4Theory in LHC@Home

Basic ingredient for the LHC common approach to VC under development -> See I Reid talk

Building block of the CERN Computing Challenge

The Virtual Machine is booting Log-in and keep track of your progress

Activity	Progress	Contribution	Ranking
How actively your virtual machine utilizes your CPU.	What fraction of the current job is completed.	How many simulation jobs your virtual machine has processed.	Your ranking among the volunteers in this challenge, by number of jobs completed.
10 % 0 %	1 % 0 %	0	0

While waiting... why don't you check if antimatter falls up?

Your virtual accelerator is initializing

The first time you start your machine it will take **up to 10 minutes** until it's ready. Please be patient...

We just started a [Virtual Computer](#) in your machine. It is a [CernVM Linux](#) machine that runs the experimental software CERN and its collaborators have developed. This *virtual computer* is going to simulate particle-particle collisions, exactly like the ones that occur in the [Large Hadron Collider](#) at CERN.

Stop [Settings] [Trash]

BOINC | About



CernVM in Data Preservation

ALEPH Software
in CernVM / SL 4

Demonstrates that
VMs can bridge 10+
years; see also

[DPHEP Status Report](#)

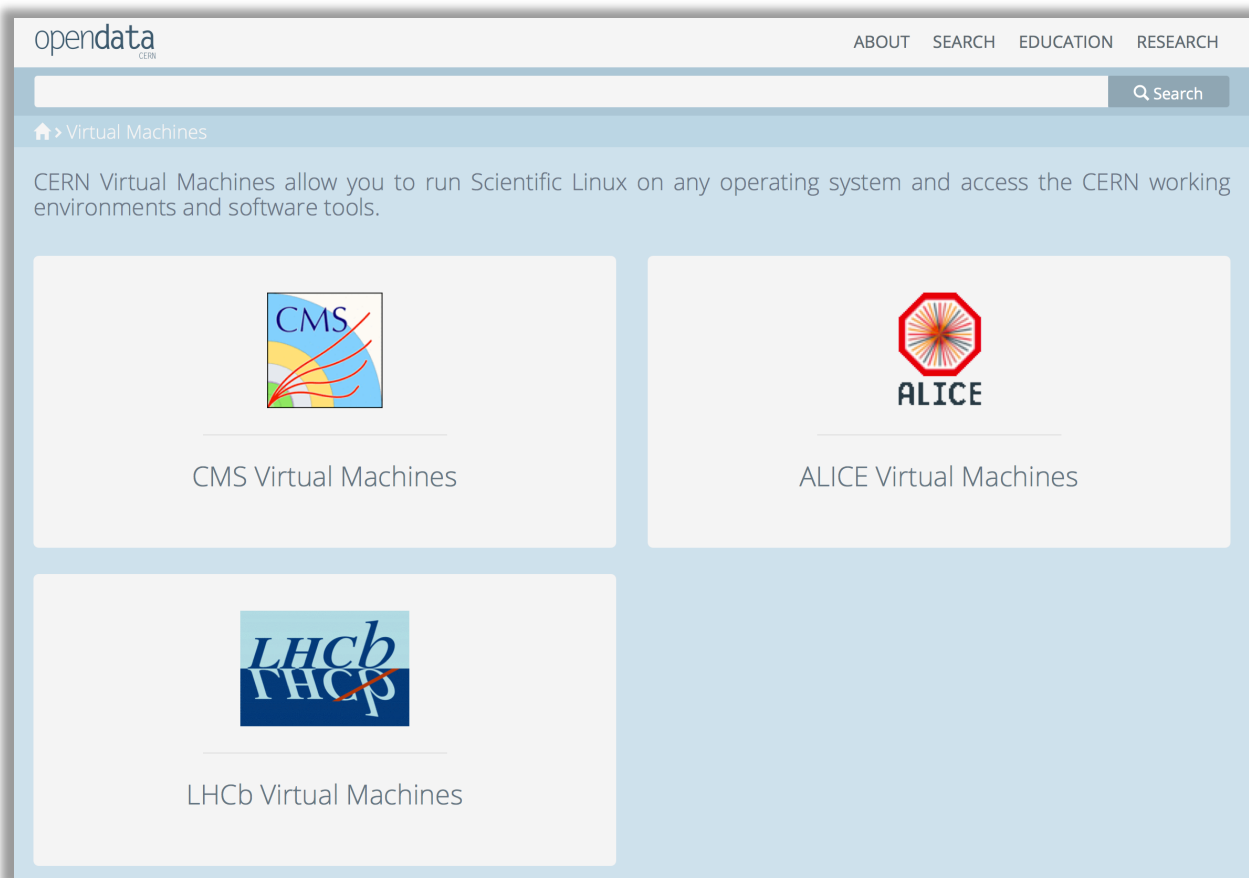
```
jakob — aleph@cernvm-aleph01:~/test/ALPHA — ssh — 66x18
pb-d-128-141-134-74:~ jakob$ ssh -X aleph@cernvm-aleph01
aleph@cernvm-aleph01's password:
[aleph@cernvm-aleph01 ~]$ source setaleph.sh
[aleph@cernvm-aleph01 ~]$ cd test/ALPHA/
[aleph@cernvm-aleph01 ALPHA]$ sh alpha.sh
*****
*****                ALPHA RUN                ***** 11.6 *****
*****
*****

Wed Mar 19 16:10:27 CET 2014

*****
***   Compilation and creation of the makefile 6lep.mk
*****
gmake -f /home/aleph/test/ALPHA/6lep.mk
gmake: `6lep' is up to date.
```

CernVM in Outreach

CERN OpenData
Project enabling
technology

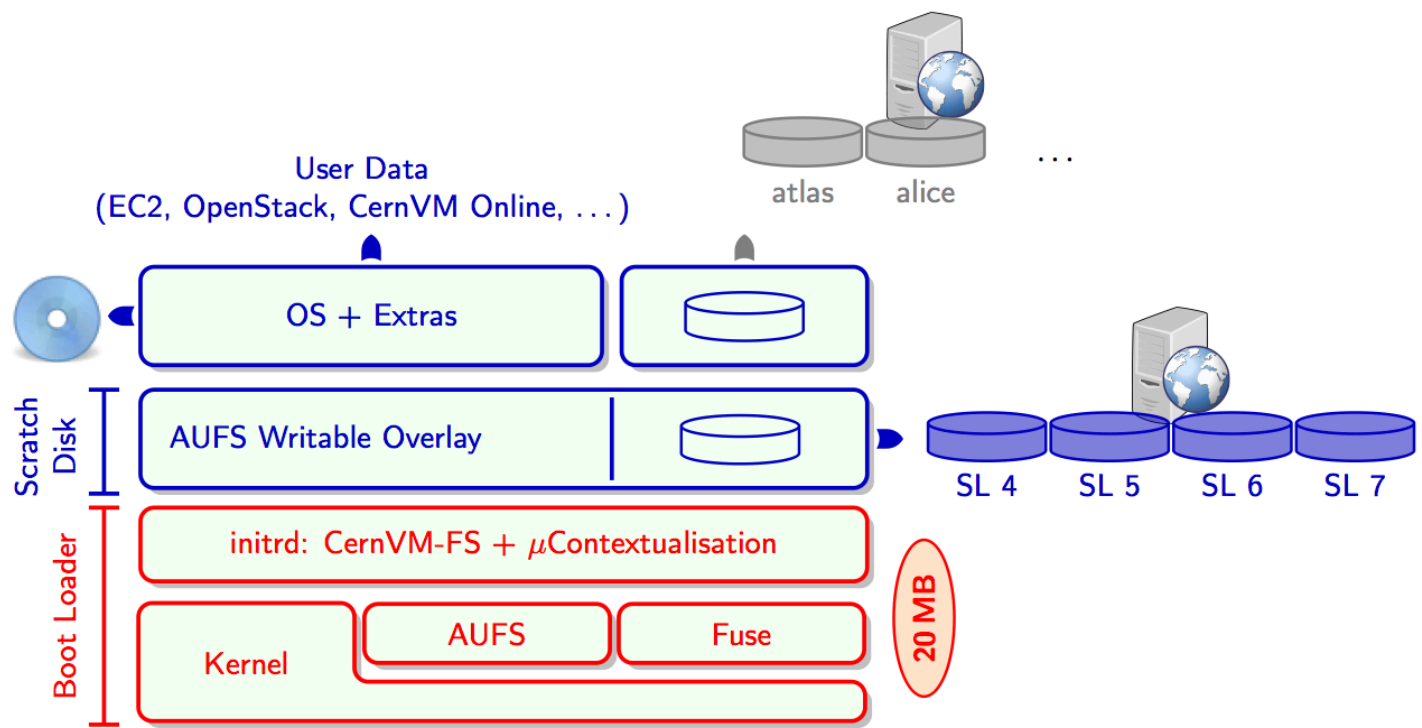


The screenshot shows the 'opendata.cern' website. The header includes navigation links for 'ABOUT', 'SEARCH', 'EDUCATION', and 'RESEARCH', along with a search bar. The main content area is titled 'Virtual Machines' and contains three cards: 'CMS Virtual Machines' with the CMS logo, 'ALICE Virtual Machines' with the ALICE logo, and 'LHCb Virtual Machines' with the LHCb logo.

See also:
CERN@School,
T Whyntie talk

The bootloader technology

μ CernVM 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](https://arxiv.org/abs/1307.2009) (CHEP 2013)

Status of the building blocks

Boot-loader (image)

- 3.10 -> 3.18 -> 4.1 Linux Kernel (LTS) + “guest additions”
- Root file system (/) from CernVM-FS + writable overlay
- Contextualization to choose OS, proxy, ...

OS on CernVM-FS

- Based on Scientific Linux
 - Configuration and tuning, Contextualization, Extra packages
- Strongly versioned: meta package determines all package versions
- Default repository: based on **SL 6.6** (CernVM 3.6)
 - Production grade, 3 major releases, ~30 security hot fixes
- Also available:
 - **SL4** : prototype for LEP experiments
 - **SL5** : Stable, GUI and contextualization
 - **SL7** : preview, under consolidation

Hypervisor Support Status

Almost all cloud environments: key factor to success

Hypervisor	
VirtualBox	✓
VMware	✓
KVM	✓
Xen	✓
Microsoft Hyper-V	✓
Parallels	✗ ¹
Docker	✓

(Cloud) Controller	
Openstack	✓
OpenNebula	✓
CloudStack	✓
Amazon EC2	✓ ²
Google Compute Engine	✓ ³
Microsoft Azure	✓
Vagrant	✓
VMware vCloud Air	4

Recent additions (includes support for glideinWMS-VM used by CMS)

- ¹ Unclear license of the guest additions
- ² Only tested with ephemeral storage, not with EBS backed instances
- ³ Only amiconfig contextualization
- ⁴ Planned

Vagrant

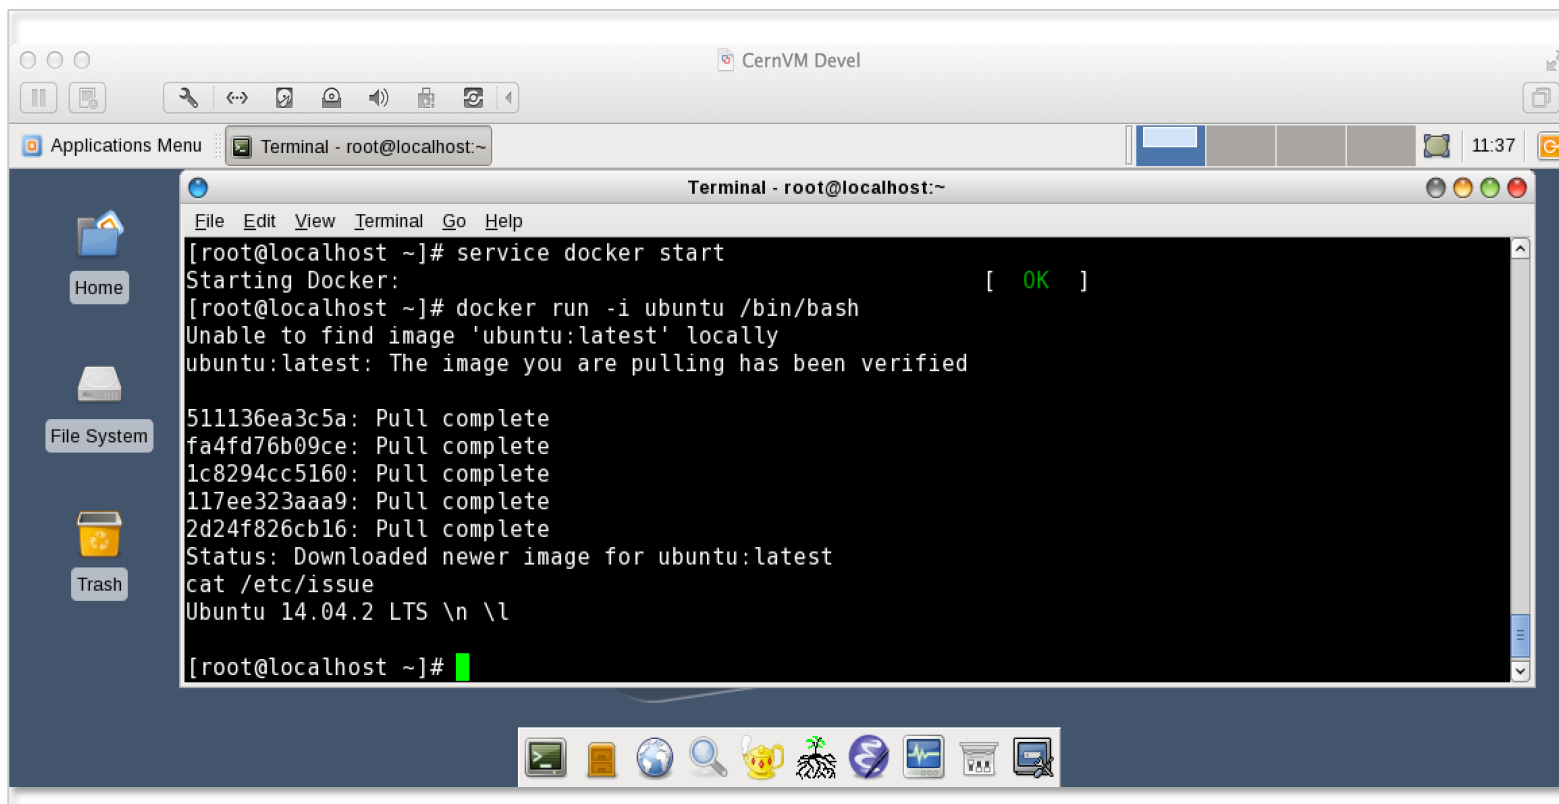
- Facilitates the instantiation of reproducible and portable development environments using virtual machines.
- Wrapper around hypervisors and cloud controllers, e. g. around VirtualBox.
- For Linux, Windows, OS X

Example:

- `vagrant init cernvm/3-prod`
- `vagrant up`
- `vagrant ssh`

Full support for Lxc and Docker

- Allows to run other linux flavours in CernVM, e.g. Ubuntu
- Available starting from CernVM 3.4




```
Terminal - root@localhost:~  
File Edit View Terminal Go Help  
[root@localhost ~]# service docker start  
Starting Docker: [ OK ]  
[root@localhost ~]# docker run -i ubuntu /bin/bash  
Unable to find image 'ubuntu:latest' locally  
ubuntu:latest: The image you are pulling has been verified  
511136ea3c5a: Pull complete  
fa4fd76b09ce: Pull complete  
1c8294cc5160: Pull complete  
117ee323aaa9: Pull complete  
2d24f826cb16: Pull complete  
Status: Downloaded newer image for ubuntu:latest  
cat /etc/issue  
Ubuntu 14.04.2 LTS \n \l  
[root@localhost ~]#
```

Sharing of multi-core VMs w/ containers

The [cernvm-fork](#): utility allows for simple creation of resource containers with isolated environments

```
$ cernvm-fork fork03 -new -fast -cvmfs=sft.cern.ch \  
  -run=/cvmfs/sft.cern.ch/my_app
```

- Facilitates sharing of multi-core virtual machines
- CernVM-FS cache shared across containers



CernVM as container

Issues

- Make CernVM-FS repositories available inside the container
- Work around missing union file system for writable overlay

Options for CernVM-FS (1)


1: Fuse, mapped from host

- Shared cache
- Collaboration from host

2: Using Parrot CernVM-FS

- Pure user-space (ptrace)
- Less performant, some limitations (e.g. suid binaries)

Root file system (/) layout



CernVM as a Docker container

- Realizes uCernVM idea within Docker
- Access CernVM-FS via Parrot or from host

Basic usage

```
$ cat <CernVM Docker tarball> | docker import - my_cernvm
$ docker run -it my_cernvm /init

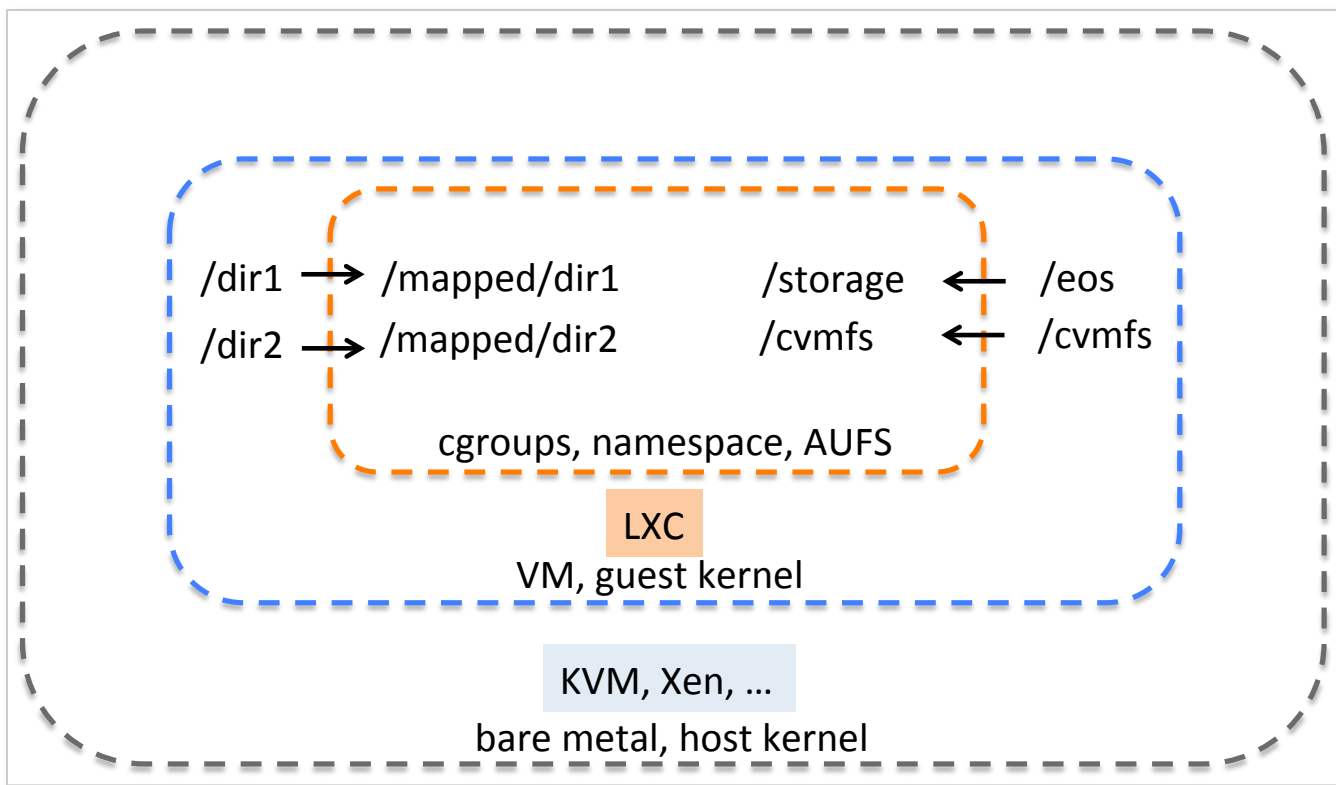
# CernVM-FS from host
$ docker run -v /cvmfs/cernvm-prod.cern.ch:/cvmfs/cernvm-prod.cern.ch
...
```

Limitations

- Runtime environment only
- Cannot be used as a base image to create derived containers

See: <http://cernvm.cern.ch/portal/docker>

Analysis Preservation: containers



- May help disentangle system and application software
 - Host and guest kernels can, in principle, be updated as needed, for example, for security issues

CernVM WebAPI: simplifying volunteers life

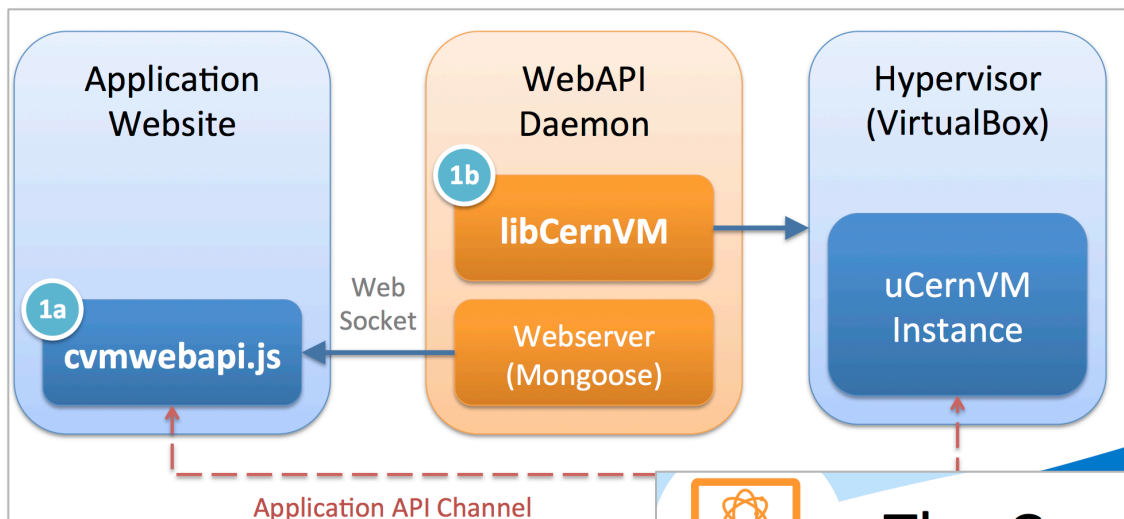
The problem

- In Volunteer Computing virtualization helped in solving the application porting problem
- But for end-users meant extra-work, manual intervention ... a burden

Goals

- Offload all possible work from users
- Assist them with manual interventions
- Standard for interfacing applications inside the VM

CernVM WebAPI



cvmwebapi.js:
High-level abstraction API

libCernVM:
Standalone C++ library to interact with a local hypervisor

- Already used in
- CernVM-Online
 - CERN Computing Challenges
 - Citizen Grid

Source, work: I Charalampidis



The CernVM Online Interface

Dashboard

Your context definitions

Name	Operations	WebAPI
CopilotVM	Clone Withdraw	Launch now
CopilotVM-Agent	Clone Withdraw	1 CPU / 1 GB RAM / 10 GB disk 1 CPU / 2 GB RAM / 10 GB disk 2 CPU / 2 GB RAM / 20 GB disk
LXC v5	Clone Withdraw	
T4T-Client-23	Clone Withdraw	Launch now
T4T-Client-44	Clone Withdraw	Launch now

CernVM 4 : RHEL 7 compatible

Roadmap

- ✓ CernVM-FS support for **capabilities** (extended attribute)
- ✓ CernVM-FS **systemd** integration as a “low level storage daemon”
- ✓ System configuration and tuning with *systemd*
- ~25 minor issues open

Boot time optimization

Scenario	CernVM 3	CernVM 4
ADSL	120s	55s
CERN	60s	40s
Warm Cache	37s	17s

Future steps on CernVM 4

- Consolidation
 - Contextualization streamlining
- Testing, performance tuning
 - Better understanding of usage (analytics)
- Platform for the LCG builds
- Images available on CERN Openstack

Recent releases at a glance

CernVM 3.4, June 2015,

[SL 6.6, support for Docker and lxc](#)

CernVM 3.5, August 2015

[SL 6.7, support for Microsoft Azure, Cloudstack, glideinwms-vm](#)

CernVM 3.6, April 2016

[Kernel 4.1, SL 6.7, consolidated bootloader](#)

[Available in Vagrant, Docker](#)

CernVM 4 preview, November 2015

[RHEL 7 compatible](#)



Questions