



# Potential of Virtualization Technology for Long-term Data Preservation

J Blomer on behalf of the CernVM Team

[jblomer@cern.ch](mailto:jblomer@cern.ch)

CERN PH-SFT

## Potential of Virtualization Technology

Preserve the historic **data analysis environment**

*(This capability is only a part of long-term data preservation)*

## Potential of Virtualization Technology

Preserve the historic **data analysis environment**

*(This capability is only a part of long-term data preservation)*

### Motivation:

#### ① Process legacy data


- Data formats are typically self-describing or convertible
- Software implicitly encodes knowledge about the correct interpretation of the data
- After substantial upgrades and modifications of the detector, the new software might lose this legacy knowledge


#### ② Validation of new software versions

- Otherwise, if the new software can process legacy data, comparison with historic version provides input for validation



**CernVM Goal:** *a uniform and portable environment for developing and running LHC data processing applications*





```
> cmsRun DiPhoton_Analysis_cfg.py
```



```
> cmsRun DiPhoton_Analysis_cfg.py
```



Individual  
Analysis Code

*0.1 MLOC*

Experiment  
Software Framework

*4 MLOC*

High Energy Physics  
Libraries

*5 MLOC*

Compiler  
System Libraries  
OS Kernel

*20 MLOC*



```
> cmsRun DiPhoton_Analysis_cfg.py
```



Individual  
Analysis Code

0.1 MLOC

Experiment  
Software Framework

4 MLOC

High Energy Physics  
Libraries

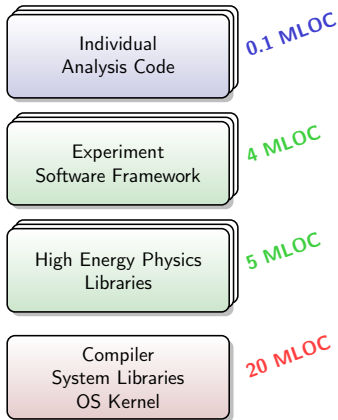
5 MLOC

Compiler  
System Libraries  
OS Kernel

20 MLOC

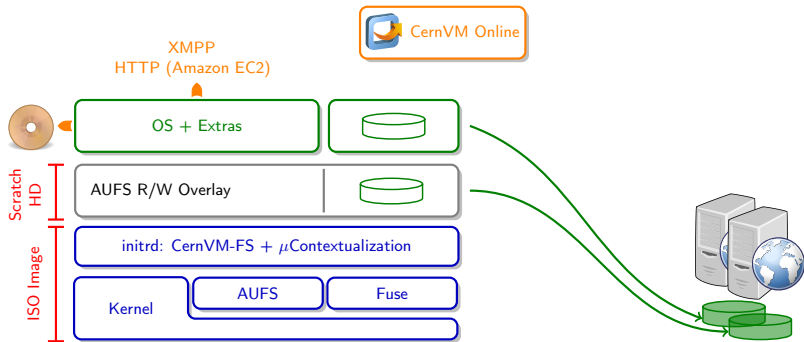


```
> cmsRun DiPhoton_Analysis_cfg.py
```



- Amplifying
- changing
- Frequent Updates
  - Not a single binary – a development environment
  - Hundreds of libraries with partially untracked dependencies
  - Not easily chunkable
  - *Not easily packagable*
- stable

## CernVM Blueprint for Preserving the Data Processing Environment

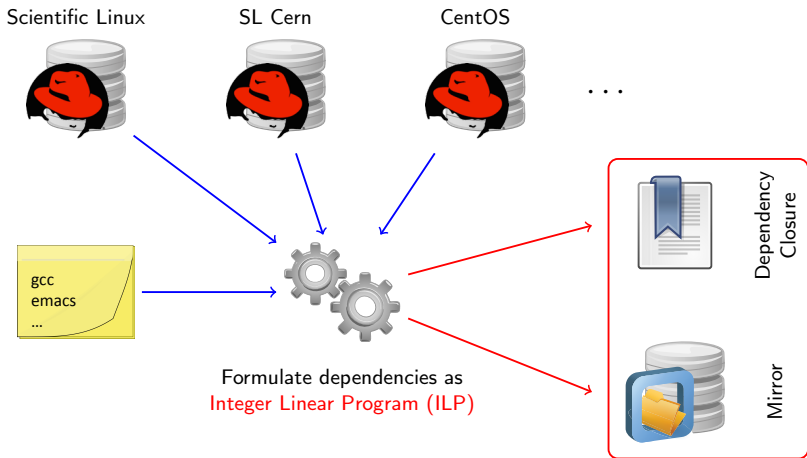


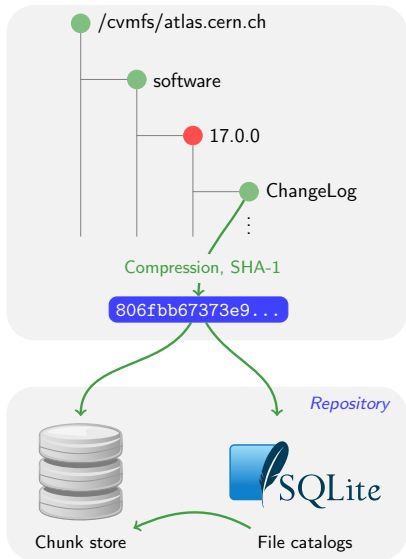
Provide the analysis environment **including** the operating system on CernVM-FS.

- No packaging, the environment is provided on demand
- CernVM-FS is a *snapshotting* and *versioning* file system
- Only 2 CernVM-FS “version strings” describe the data analysis environment

Maintenance of the repository **must not** become a Linux distributor's job

**Idea:** Automatically generate a **fully versioned, closed** package list  
 from an unversioned "shopping list" of packages  
*(Standard package managers are not designed for preservation!)*





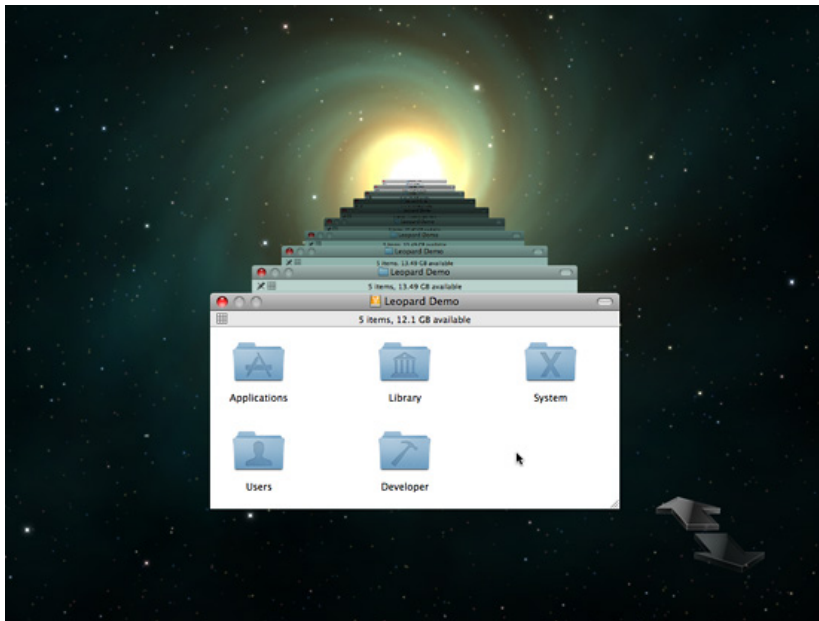
## Data Store

- Compressed chunks (files)
- Eliminates duplicates
- Never deletes

## File Catalog

- Directory structure, symlinks
- Content hashes of regular files
- Digitally signed  
⇒ integrity, authenticity
- Plain files, stored as chunks in the data store

The *root hash* (40 characters) defines a file system snapshot (similar to git)  
**Intrinsic and explicit versioning required**

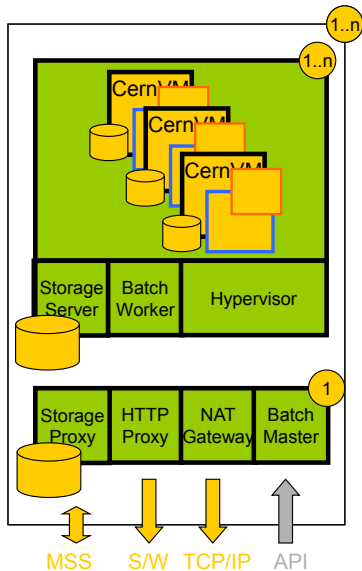


## Simple API

- Instantiate + Contextualize
- Terminate
- List instances, list images

## Amiconfig plug-ins

- Credentials (ssh, X.509)
- Condor head & batch services
- Squid server
- XrootD storage proxy
- Monitoring & directory service agents
- Network configuration & tuning



Dashboard | CernVM Online

https://cernvm-online.cern.ch/dashboard

Logged in as icharala | Log out

About Dashboard

## Commands
















- Dashboard
- Pair an instance
- Create Context


## Recent Definitions

- PrivateCloud-MakeflowPool
- PrivateCloud-Worker
- Private-Desktop
- PrivateCloud4ALL
- PrivateCloud-CatalogServer



## Dashboard


### Your context definitions

Name	ID	Operations
 PrivateCloud-CatalogServer	22b13425ef7244c4b7de60dbbca64728	 Remove  Use as template
 PrivateCloud-Worker	a632e7ad3ca64774ac9318fc9e086640	 Remove  Use as template
 PrivateCloud-MakeflowPool	f7e9ba92a55146119ac3cd6141f6d957	 Remove  Use as template
 Private-Desktop	79383bc71d7a4760a8397cc2e6d2a2ed	 Remove  Use as template
 PrivateCloud4ALL	37bbd937803b407e8fbd469d65fac74c	 Remove  Use as template

 Create new context

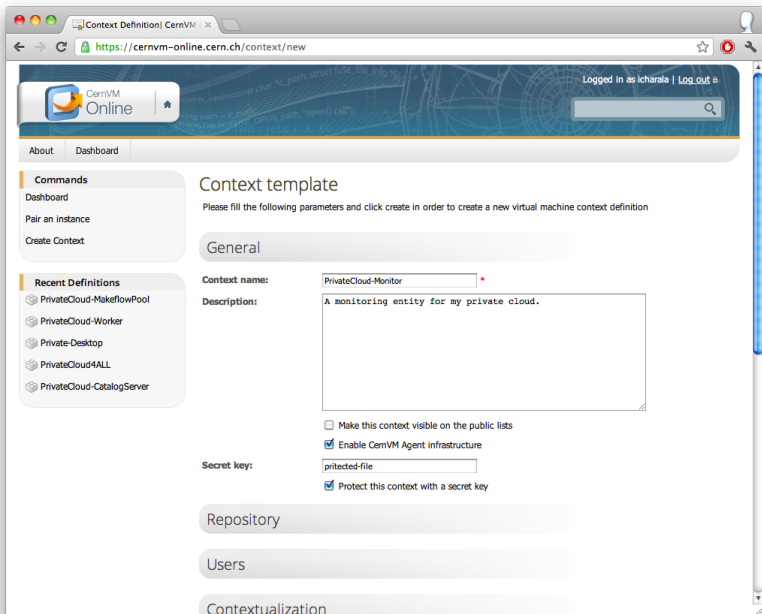
### Your virtual machines

Machine	CernVM	Context	Operations
 128.141.235.19 (b5609c63-a2bb-4ba5-901a-25066733df9c)	2.1.0	PrivateCloud-CatalogServer	 Unmanage

 Pair an instance of CernVM

© Copyright CERN 2012 - PH Department - SFT - CernVM Software appliance





The screenshot shows a web browser window with the URL `https://cernvm-online.cern.ch/context/new`. The page title is "Context Definition | CernVM". The user is logged in as "icharala" and can click "Log out".

The main content area is titled "Context template" and includes the instruction: "Please fill the following parameters and click create in order to create a new virtual machine context definition".

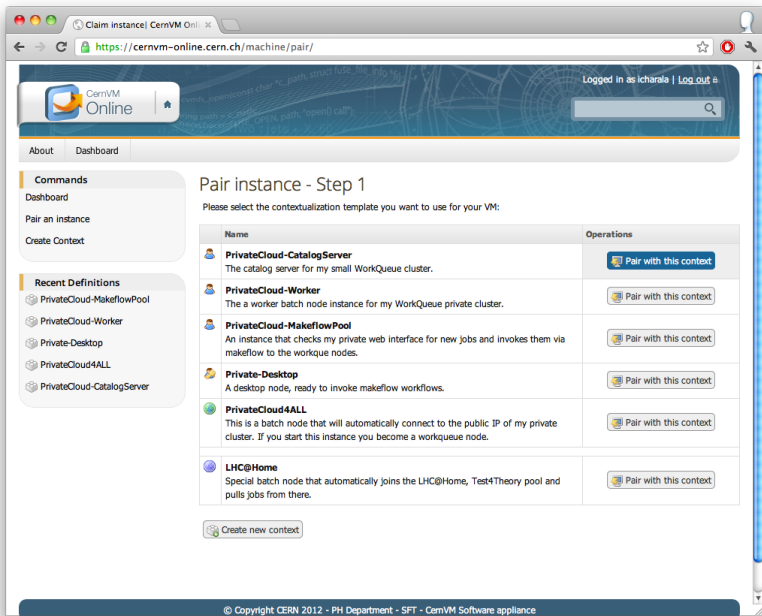
The "General" section contains the following fields and options:

- Context name:**
- Description:**
- Make this context visible on the public lists
- Enable CernVM Agent infrastructure
- Secret key:**
- Protect this context with a secret key

Below the "General" section are three tabs: "Repository", "Users", and "Contextualization".

On the left sidebar, there are two sections:

- Commands:** Dashboard, Pair an instance, Create Context
- Recent Definitions:**
  - PrivateCloud-MakeflowPool
  - PrivateCloud-Worker
  - Private-Desktop
  - PrivateCloud4ALL
  - PrivateCloud-CatalogServer



Claim instance | CernVM Online | <https://cernvm-online.cern.ch/machine/pair/>

Logged in as icharala | [Log out](#)

About | Dashboard

**Commands**







- Dashboard
- Pair an instance
- Create Context

**Recent Definitions**

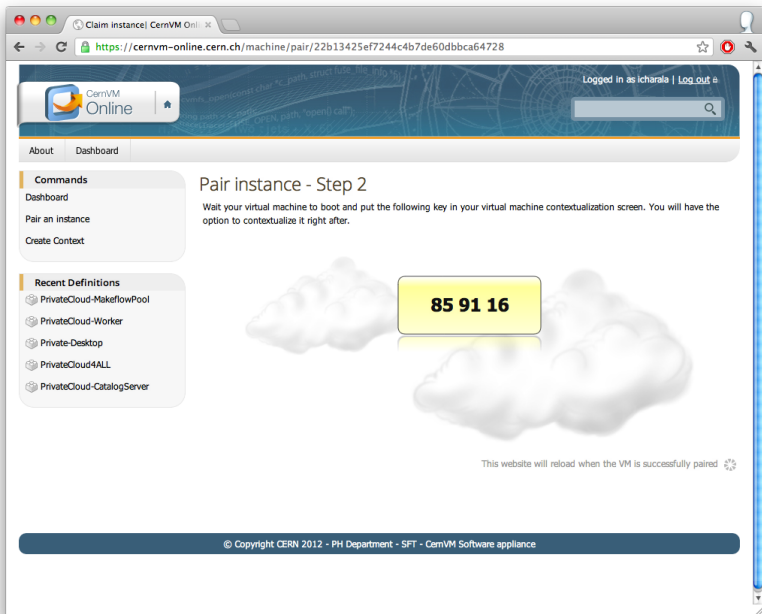
- PrivateCloud-MakeflowPool
- PrivateCloud-Worker
- Private-Desktop
- PrivateCloud4ALL
- PrivateCloud-CatalogServer

## Pair instance - Step 1

Please select the contextualization template you want to use for your VM:

Name	Operations
 <b>PrivateCloud-CatalogServer</b> The catalog server for my small WorkQueue cluster.	<input type="button" value="Pair with this context"/>
 <b>PrivateCloud-Worker</b> The a worker batch node instance for my WorkQueue private cluster.	<input type="button" value="Pair with this context"/>
 <b>PrivateCloud-MakeflowPool</b> An instance that checks my private web interface for new jobs and invokes them via makeflow to the workque nodes.	<input type="button" value="Pair with this context"/>
 <b>Private-Desktop</b> A desktop node, ready to invoke makeflow workflows.	<input type="button" value="Pair with this context"/>
 <b>PrivateCloud4ALL</b> This is a batch node that will automatically connect to the public IP of my private cluster. If you start this instance you become a workqueue node.	<input type="button" value="Pair with this context"/>
 <b>LHC@Home</b> Special batch node that automatically joins the LHC@Home, Test4Theory pool and pulls jobs from there.	<input type="button" value="Pair with this context"/>

© Copyright CERN 2012 - PH Department - SFT - CernVM Software appliance



Claim instance | CernVM Online

https://cernvm-online.cern.ch/machine/pair/22b13425ef7244c4b7de60dbbca64728

Logged in as ichtarla | Log out

CernVM Online

About Dashboard

**Commands**

- Dashboard
- Pair an instance
- Create Context

**Recent Definitions**

- PrivateCloud-MakeflowPool
- PrivateCloud-Worker
- Private-Desktop
- PrivateCloud4ALL
- PrivateCloud-CatalogServer

## Pair instance - Step 2

Wait your virtual machine to boot and put the following key in your virtual machine contextualization screen. You will have the option to contextualize it right after.

**85 91 16**

This website will reload when the VM is successfully paired

© Copyright CERN 2012 - PH Department - SFT - CernVM Software appliance

```

CernVM-15 [Running]
Welcome to CernVM Virtual Machine, version 2.6.0
Machine UUID 2f2f1157-4f45-405a-bf1c-3a5b309c87d5
To contextualize your VM log-in to http://cernvm-online.cern.ch/

Instance pairing pin: 859116

Changing password for user cernvm.
passwd: all authentication tokens updated successfully.
INIT: Switching to runlevel: 5
INIT: Sending processes the TERM signal
Starting CernVM: Shutting down CernVM-FS:           [ OK ]
Starting CernVM-FS:                                  [ OK ]
Starting CernVM-FS:                                  [ OK ]
Starting CernVM-FS:                                  [ OK ]
Starting vmcontext_hepix: Starting vmcontext_hepix ...
-
  
```



## Virtual Machine

- Linux distribution based on Scientific Linux.
- Supports all popular hypervisors.
- Minimal footprint, the VM *interface* is needed
- Flexible contextualization.

## CernVM Filesystem



- Read-only, globally distributed file system optimized for software distribution.
- Based on plain files and HTTP
- Snapshotting and versioning file system
  - Already used in production by LHC experiments.

## CernVM - based data analysis environment preservation

- CernVM-FS environment is defined version strings. OS packages are defined by a versioned, closed package group (Meta-RPM)
- You need only the CernVM version string to rebuild CernVM image on demand.

- Ensembles of CernVMs can recreate a virtual cluster for data processing.

- CernVM can be contextualized using a small subset of EC2 API that allows it to be deployed on public or private clouds



## Bookkeeping

## Private Cloud





# NA61 Production Jobs in Belgrade

## Integration of a CernVM cloud with a data provenance system


The screenshot shows a desktop environment with two main windows. On the left is a Mozilla Firefox browser window titled "Under construction - Mozilla Firefox" displaying the NA61/SHINE production status page. The page has a green header "LINE - Status of production" and a search box with "year-2012-exp-[p@120GeV[C\_2cm]". Below the search box are several status boxes: "23% space left in Calibration directory" (green), "web acronym checked in 1 minutes ago" (green), "lxbatch-CERN - not connected" (pink), "cernvm-CERN checked in 1 minute ago" (green), and "cernvm-Belgrade checked in 1 minute ago" (green). At the bottom of the browser window, there is a small graph showing a red line with peaks. On the right is a Virtual Machine Manager window titled "Virtual Machine Manager (on localhost.localdomain)". It shows a list of virtual machines with columns for Name, CPU usage, and Host CPU usage. The VMs listed include "batch102", "na61headnode", "batch201", "batch202", "batch301", "batch401", "na61batch501", "na61batch601", "batch701", "batch702", and "batch703". Each VM has a small graph showing its CPU usage. A tooltip at the bottom right of the VM manager window reads "NA61/SHINE production testin page. Under construction - Mozilla Firefox". Red arrows point from the browser window to the VM manager window and from the VM manager window to the browser window.

virtual machines manager on Belgrade's cluster.

Belgrade's cluster registered and monitored on  
NA61/SHINE production page (*under construction*)

<http://dmaletic.web.cern.ch/dmaletic/cgi-bin/na61prod>

- Plain virtualization is not sufficient to preserve the data processing environment
- CernVM and CernVM-FS technologies provide handles to re-create a **data processing environment identified by a few version strings**
- As such, it is easy to integrate the definition of the data processing environment in data provenance systems
- The exposed interface is very slim: CernVM clusters run on private and public clouds without grid infrastructures
- Such virtual machines are easy to use and they can be given to “interested citizens” (see also *LHC@Home 2.0 volunteer cloud*)



Please contact us for any further questions!