

# **CernVM plans for 2017**

G. Ganis for the CernVM Team  
16 January 2017

# Outline

- Human resources
- Summary of activities in 2016
- Plans for 2017

# People

		2016	2017	2018
G. Ganis	Sf	50%	<b>50%</b>	50%
J. Blomer	Sf	100%	<b>100%</b>	100%
R. Popescu	Fw	33%	<b>100%</b>	66%
P. Jirout	TS	50%	<b>20%</b>	
N. Hardi	TS (KT)	25%	<b>75%</b>	
N. Kousi	Fw	15%	<b>15%</b>	
R. Meusel	Fw	40%		
J. Molina	TS	25%		
P. Buncic	St	5%	<b>5%</b>	5%
Total / FTE		3.43	<b>3.65</b>	2.15+

# Summer Students in 2016

- Tim Schaffer (OpenLab)
  - *HEP Application delivery on HPC resources*
- Felix Scheffler
  - *CernVM porting to Arm64*

# The CernVM ecosystem

- Main-stream

- CernVM-FS
- CernVM-Appliance

All usages  
D, C, V , DP

- Side-stream

- CernVM-Online
- (WebAPI)

D, C, DP  
D, V

- Discontinued

- Copilot
  - Used in LHC@Home (T4T), now replaced by DataBridge

# CernVM(-FS) reminder

- CernVM-FS
  - Optimized for large scale software distribution
  - Mission **critical** for WLCG
    - Also used by ~45 non-LHC experiments, including EUCLID
- CernVM
  - Designed to seamlessly setup computing environments for HEP experiments
  - Using bootloader technology
    - OS loaded from CernVM-FS, package management by RPM
  - **SL7, SL6, SL5, SL4** compatible
  - Used by all 4 LHC experiments
    - Laptops / desktops, Computing clouds, HLT farms, Data preservation, Open data, Master Classes



# Current Scale of Deployment

- > 350 million files under management
- > 50 repositories

Running jobs: 334639  
Active CPU cores: 460014  
Transfer rate: 13.29 GiB/sec

- 10 k new VMs / month




Image Landsat  
Data SIO, NOAA, U.S. Navy, NGA, GEBCO  
Image IBCAO  
Image U.S. Geological Survey

Google earth

Terms of Use  
eye alt 15825.12 km

# Summary of 2016 activities

- CernVM-FS
- CernVM
- AoB



# CernVM-FS in 2016 (1)

- 2.2.0 -> 2.2.3, Feb-May 2016
  - Move to semantic versioning
  - Support for OverlayFS (replacement for AUFS)
    - Requires kernel > 4.2: Fedora 23, Ubuntu 15.10, EL7.3
  - Support for extended attributes
  - Support for SHAKE-128 (SHA-3 derived) hash algorithm
  - Support for new platforms
    - Mac OS X El Capitan, AArch64, Power-8
  - Consolidation and bug-fixing

# CernVM-FS in 2016 (2)

- 2.3.2, Oct 2016
  - Support for data distribution
    - Collaboration with Univ. of Nebraska (B. Bockelman)
  - Plug-in interface for client authorization helpers
  - Improved garbage collection
  - New platforms
    - Fedora 24, SLES12 on x86\_64, CentOS 7 on AArch64
  - Reduced memory footprint
  - Further consolidation and bug fixing

# CernVM-FS in 2016 (3)

## Progress on new publishing interface

- Completed
  - Requirement analysis
    - Includes initial discussion with stakeholders (ALICE, SFT, IT)
  - Initial design of the new components
    - Including protocol, choice of technology
- In progress
  - Implementation of the new back-end components
  - Integration of new components with existing server-side tools

# CernVM-FS in 2016 (4)

- Successful submission to the KT fund of a project to integrate CernVM-FS with Docker
  - Project approved in April
  - Request submitted at TSC 2016/1
  - Project started on 1 October
- Status
  - Understood Docker internals
    - Familiarity w/ Go language, OCI standard specs
  - Identified required modifications
  - Submitted a RFC on Docker GitHub to get in touch with the community

# CernVM in 2016 (1)

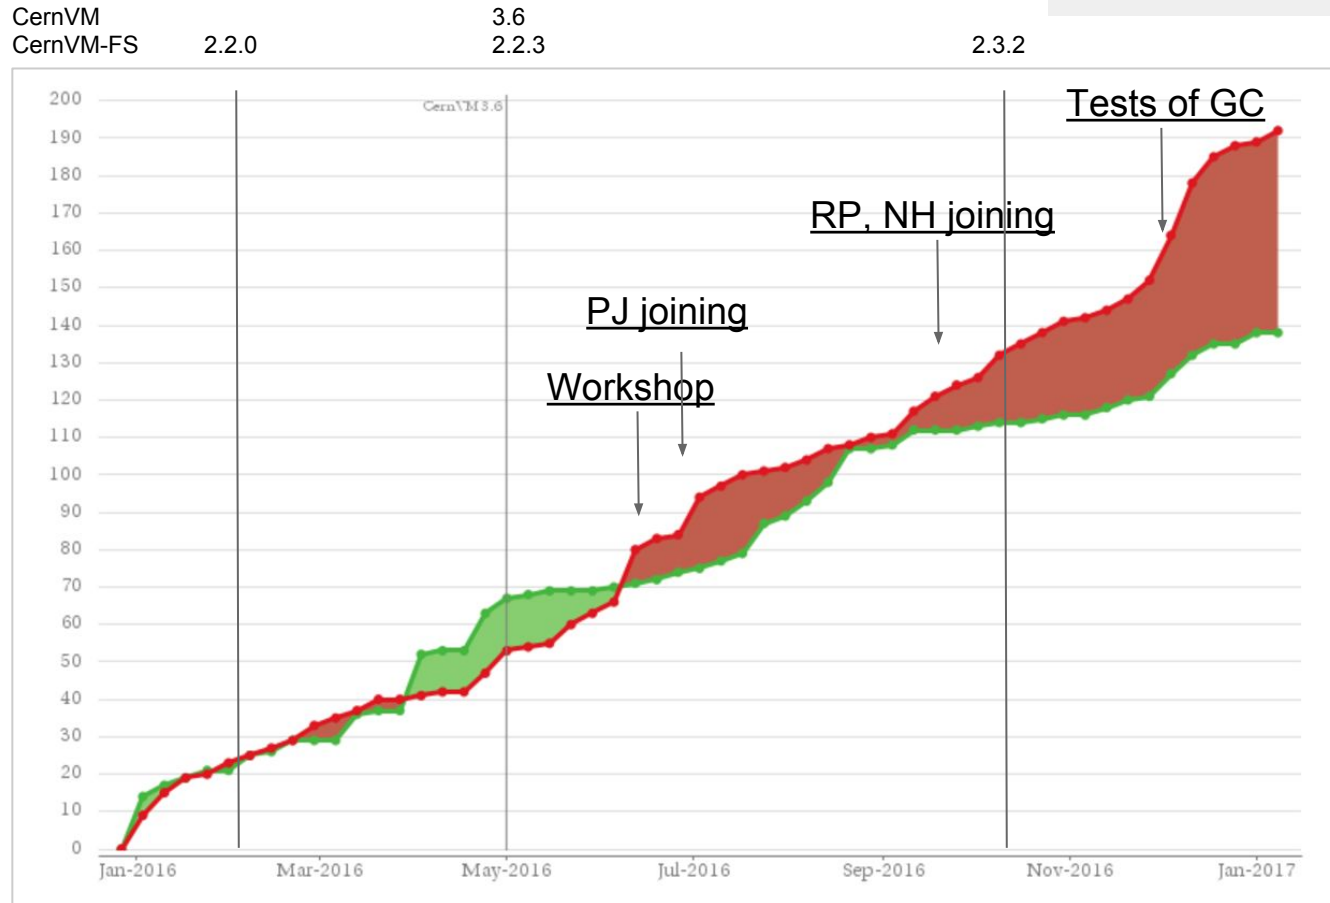
- **Bootloader (u-CernVM) upgrade**
  - Kernel 4.1.23
    - Includes fix against 'dirty cow'
- **Version 3, RHEL 6 based**
  - Commissioning of v3.6 (SL 6.6), May 2016
    - Kernel 4.1.23
    - Support for Vagrant, glideinWMS
    - Improved boot time, use cvmfs proxy for contextualization
  - 14 security hotfixes (current version 3.6.5.14, Nov 25th)
- **Version 4, RHEL 7 based**
  - Consolidation of current preview state
    - Commissioning synchronized with lxplus/lxbatch
  - Boot time optimization (factor ~2 faster than v3)

# CernVM in 2016 (2)

- New ports
  - ARM64 (F. Scheffler)
  - 32-bit for Volunteer Computing activities
- Support for Amazon EC2 HVM
  - Almost native performance
    - Direct access to NIC and GPU
- First release of launcher tool
  - À la Vagrant/Docker interface to control CernVMs
  - Preview at <http://cernvm.cern.ch/portal/launch>

# JIRA activity in 2016

Open/Close issues



Since Jan 2016

CernVM-FS: 143 new issues (43 bugs), 59 resolved (33 bugs)  
CernVM: 23 new issues (8 bugs), 8 resolved (4 bugs)

# Other activities 2016

- Continued contacts with EUCLID
  - Technical meeting at CERN (April)
  - Official adoption of CernVM-FS for software distribution
- Established contacts with LIGO
  - CERN recognized experiment
  - Using CernVM-FS for data federation
- New expression of interest
  - MESOSPHERE, PIVOTAL
    - Demo of container distribution with CernVM-FS



# Other activities 2016

- New website
  - Migrated to Drupal 7, same look and feel
- DPHEP
  - Update presented to February DPHEP workshop
  - Contributed to long paper to iPRES 2016, Bern, Oct 2016
    - “CERN Services for Long Term Data Preservation”
    - Presented by J Shiers
- CernVM Users Workshop, RAL/UK, June

# CernVM Users Workshop in 2016

- [June 6-8, 2016, RAL/UK](#)
- Essential Input/Feedback from users, experiments and sites
- Technology session with invited speakers
  - Mesosphere, VMware, Linaro, IBM, Pivotal
- Main outcome
  - Development plans fit the needs of stakeholders
    - Improve publication process
    - Solutions for HPC
    - Flexibility in cache setups
  - CernVM appliance used by all LHC experiments
    - In a way or another

# Other dissemination activity in 2016

- CHEP 2016 contributions involving team members
  - *New directions in the CernVM File System* (JB, PB, GG, RM, RP; poster)
  - *Accessing Data Federations with CernVM-FS* (JB, RM et al.)
  - *Web Proxy auto-discovery for WLCG* (JB et al.)
- Talks
  - GG, JB *Update on Software Environment Preservation with CernVM*, DPHEP, February
  - GG, JB, *CernVM ecosystem status*, LHCb workshop, June
  - GG, JB, *CernVM-FS as Distribution Engine for containers*, EP KT Innovation Day, October
  - JB, RP, *CernVM-FS devs*, ALICE (October) and IT (December)
  - JB et al., *A Novel Approach for Distributing and Managing Container Images*, MesosCon, Denver, US, June
  - JB, *Global Software Distribution with CernVM-FS*, CCL Workshop On Scalable Computing, Notre-Dame, US, October
  - JB, *Cloud Technologies in HEP data processing*, ICFDT, Frascati, Italy, March
  - JB et al., *x32-ABI Re-evaluation*, Technology meeting, January
- Papers
  - JB, GG et al., *CERN Services for Long Term Data Preservation*, iPRES 2016, Bern, Switzerland

# CernVM 2017 Plan Of Work

- Guidelines
  - Support for running services
  - Developments driven by unique needs of main stakeholders

# CernVM-FS devs guidelines

- Address requirements of Run3 and AFS replacement
  - Reduction of publishing time to address more demanding use-cases
    - E.g. ATLAS nightlies, ALICE conditions data
  - Fully exploit version tracking capabilities of the file system
    - Required, e.g., for ALICE conditions data @ Run 3
- Improve flexibility for client deployments
  - Diskless server farms
  - HPC special storage systems (e.g. NERSC @ LBNL/US)
- Improvement of the HTTP distribution chain
  - Improved monitoring, coherent network to push the data around
- Commissioning Docker integration

# CernVM-FS (1)

## *Improving publishing phase (v2.4)*

- Complete the implementation of a working prototype
- Commissioning of the new prototype w/ ALICE
  - During the shutdown period
- Deployment for ALICE conditions data
- Further work to reduce propagation delay
  - To Stratum 1, clients
- Performance measurements and improvements
  - Scalability, ...

# CernVM-FS (2)

*Improve flexibility for client deployments (v2.4)*

- Commission version & snapshot support (\*)
- Support for client cache plug-ins
- Cache plug-in managers for different backends
  - Ceph, shared RAM cache, ...

*Docker integration*

- Realize full read/write Docker workflow
  - Effective Replace Docker registry with CernVM-FS
- Extend to generic container service

(\*) See backup

# CernVM devs guidelines

- Facilitate the creation of on-demand turn-key clusters of VMs
  - E.g. batch clusters (HTCondor), BigData clusters
- Performance studies, measurements and improvements



# CernVM-Appliance

- Consolidation/Maintenance of v 3.6
  - Follow-up security hotfixes for v3.6
  - Preparing for commissioning v4 when Ixplus moves to CC7
- Consolidate cernvm-launch
  - Support creation of cluster
  - Common repository of context files
  - Deprecate cernvm-online and webapi
- Performance benchmarking
  - Prepare context for benchmarking
    - Following benchmark working group (B Panzer)
  - Define set of representative cases
    - CERN Openstack, various configurations
    - Commercial clouds (Amazon, Google, Azure)
      - Coordinate with IT

# CernVM-Appliance (2)

- Analytics (Piwik)
  - Better understanding of usage
- AoB
  - CernVM images available of CERN OpenStack
  - Inclusion in LCG builds

# Miscellanea

- CernVM-Online and WebAPI
  - Identify and extract components that can be managed by the group
- Infrastructure
  - Migration of test HW in bld 157 to IT
- Follow up DPHEP evolution
  - Workshop in March
- Follow-up contacts with potential external users
- Prepare for next workshop
  - Tentative period: Dec 2017 or Jan 2018

# Conference targets

- Linaro Connect (Budapest, March)
  - CernVM port on ARM64
- ACAT 2017 (end August)
  - CernVM port on ARM64
  - Status of CernVM-FS devs
- [Codemesh](#) (London, November)
  - C++ and Erlang in CernVM-FS

# Summary

- Main effort goes to support running services
- Developments driven by
  - Forthcoming needs
    - Run 3, AFS replacement, ...
  - Improved flexibility for clients deployments
    - Diverse cache backends, ...
  - Exploitation of products potential
    - Docker integration, VM cluster handling, ...

Feedback is welcome and may influence our priorities

# Backup

# CernVM-FS: snapshots & branching

- Branches allow non-linear workflows on file system snapshots
  - Can describe relation between derived and parent container
  - Allow fixes to conditions data of a given, past run
- Instant access to snapshots
  - New virtual directory provides access to all named snapshots within a single mount point, `.cvmfs_snapshots` (à la zfs)

