

# **SYSTEMS plans for 2016**

G. Ganis for the Systems Team  
25 January 2016

# SYSTEMS activities

- CernVM ecosystem
- PROOF

## Outline

- People
- For each subproject
  - Summary of 2015
  - Plans for 2016

# People

		2015	<b>2016</b>	2017
G. Ganis	St / R, C, X	100%	<b>100%</b>	100%
J. Blomer	St / C	100%	<b>100%</b>	100%
R. Meusel	FI / C	100%	<b>50%</b>	
J. Molina	TS/ C	90%	<b>25%</b>	
P. Buncic	St / C	5%	<b>5%</b>	5%
P. Jirout	TS/ C		<b>50%</b>	20% ?
<b>FTE</b>		3.95	<b>3.30</b>	2.05+

C = CernVM, R = ROOT/PROOF, X = XRootD legacy responsibility (0%-10%)

# Guests & Visitors in 2015

- I. Charalampidis (11M)
  - Fellow under the “Citizen Cyberlab” EU Project Virtual Atom Smasher
    - Volunteer thinking educational game platform
    - [CERN computing challenge](#)
  - Theodoridou Lioumpa (4M)
    - Helping Ioannis with visualization tasks
- Summer Students
  - Petra Mazdin (NMS + trainee, CernVM boot time optimization)
  - Enrico Guiraud (OpenLab + trainee, multiproc)
- GSoC
  - Rafayel Mkrtchyan (Evaluation of HTTP2)

# CernVM ecosystem

- Main

- CernVM-FS
- CernVM-Appliance
- CernVM-Online
- (WebAPI)

All usages

D, C, V, DP

D, DP

D, V

- Legacy or auto-pilot

- Copilot

V

D = Desktop

C = Cloud

V = Volunteer

DP = Data Preservation

# Dissemination activity 2015

- CernVM Users Workshop, CERN, March 2015
- CernVM-FS hackathon at ISGC2015
  - Within EGI Community Forum, Taipei, Taiwan
- CHEP contributions involving team members
  - *Experiences on File Systems: Which is the best FS for you?* (JB, **plenary**)
  - *Large-Scale Merging of Histograms using Distributed In-Memory Comp.* (JB)
  - *Status and Roadmap of CernVM* (GG)
  - *CernVM WebAPI - Controlling Virtual Machines from the Web* (IC)
  - *Lightweight scheduling of elastic analysis containers ...* (pres. by D Berzano/ALICE)
  - *PROOF-based analysis on ATLAS Italian Tiers ...* (GG, poster)
  - *Using S3 cloud storage with ROOT and CvmFS* (pres. by M Rios/IT)
  - *Engineering CernVM-FS ... for Auxiliary Physics Data* (JB)
- Talks
  - GG, *Virtualization in HEP*, CCT, Toulouse, France, Jan 2015 (**EUCLIDE** workshop)
  - GG, *The Ins & Outs of XRootD AuthE & AuthZ*, XRootD Workshop, UCSD, Jan 2015
  - GG, *Update on Software Environment Preservation with CernVM*, DPHEP, June 2015
  - JB, *CernVM(-FS) as possible choice for efficient code and virtual image distribution*, **ICARUS** software meeting, July 2015
  - RM, JM, *Experience on Q&A for CernVM-FS*, 1st Developer Workshop, CERN, Sep 2015
  - JB, *The CernVM Ecosystem*, AEI, Hanover, Oct 2015 (**Advanced LIGO**)

# Journal publication

In collaboration with GROW-FS authors (D Thain, I Sfiligoi)

[The evolution of Global Scale Filesystem for Scientific Software Distribution, IEEE Comp. in Science and Engineering 17\(6\), Nov/Dec 2015](#)

## The Evolution of Global Scale Filesystems for Scientific Software Distribution

Jakob Blomer, Predrag Buncic, René Meusel, and Gerardo Ganis | CERN

Igor Sfiligoi | University of California, San Diego

Douglas Thain | University of Notre Dame

Delivering complex software across a worldwide distributed system is a major challenge in high-throughput scientific computing. To address this problem in high-energy physics, a global scale filesystem delivers software to hundreds of thousands of machines around the world.

Delivering software across a worldwide distributed system is a major challenge in high-throughput scientific computing. Within one computing center, deploying sci-  
stable releases and small footprints that can be easily packaged.

# CernVM Users Workshop in 2015

- March 5-6, 2015
  - <http://indico.cern.ch/event/348657/>
- Essential Input/Feedback from users, experiments and sites
- Technology session with invited speakers
  - Google, Amazon, Citrix, Basho, Univ. Notre-Dame
- Improved understanding of containers
- Strengthened relationship to the Parrot dev team
  - Central for HPC setups

# CernVM-FS in 2015

- Commissioning of v 2.2.0
  - Consolidation of new features in the server
    - Garbage collection
  - Improved support for HPC
    - Fixes in libcvmfs and Parrot connector, cvmfs\_preload
  - Support for OverlayFS (replacement for AUFS)
    - Requires kernel > 4.2: Fedora 23, Ubuntu 15.10
      - But not EL7: publishing nodes must stay on SLC6
  - Several fixes in the MacOSX client
    - Including support for El Capitan
  - Experimental support for data distribution
    - Collaboration with Univ. of Nebraska (B. Bockelman)
- Improved monitoring and code house keeping
  - See contribution to 1st Developer Workshop and JM talk PH-SFT, June 29th, 2015

# CernVM in 2015

- Consolidation of v3
  - Current release: v3.5, August 8, 2015
  - Systematic import of SLC6 hotfix releases
  - Complete support for private and public cloud controllers
  - Full support for Linux containers
- Preparation of v4
  - SL7 compatible
  - Boot time optimization
- Promotion as general purpose platform
  - ~28000 VMs started per day
- Demonstrated platform for data preservation
  - ([OpenData](#): ALICE, CMS, LHCb and ALEPH)

# Other activities 2015

- Continued contacts with EUCLID
  - Testing SL7
- New expression of interest
  - ICARUS
    - Presentation at their offline meeting
    - Dedicated VM prototype
  - LIGO
    - Applying as CERN recognized experiment to use CernVM(FS)
  - MESOSPHERE
    - Demo of container distribution with CernVM-FS
- CERN-NTNU Screening Week (via KT)
  - Uncovered possible business models for CernVM-FS
  - Most interesting: integration with Docker to optimise distribution of (software inside) containers

# Other activities 2015

- Migrated to Jenkins
  - Substantial effort from JM and RM
  - Extended functionality (wrt EC)
    - Automatic integration
    - Monitor performance, impact of a single patch
- Contribution to DPHEP
  - ALEPH in containers
    - Presented at June DPHEP Workshop
  - Central contribution to blueprint
    - Published in December 2015 (arXiv, Zenodo)  
[Status Report of the DPHEP Collaboration: A Global Effort for Sustainable Data Preservation in High Energy Physics](#)

# CernVM 2016 Plan Of Work

- Guidelines
  - [Support for running services](#)
  - Developments driven by HEP users unique needs
- CernVM Users Workshop, June 6-8, 2016
  - May re-tune priorities
  - Held at RAL:

<https://indico.cern.ch/event/469775/>

# CernVM-FS (target: v2.3.0, Q4/2016)

- Distributed write support
  - Allow experiments to have their own publishing node
- Support for HTTP2
- Enable proxy auto discovery
  - Simplifies installation in new clients
- Default installation on CERN Desktops
  - À la Ixplus
- Revise IPv4/IPv6 selection strategy
  - To be (even more) ready for IPv6 deployment
- Evaluate S3 backend
  - In conjunction with RAL, IT

# CernVM Appliance reminder

- Image contextualizable for specific needs
- Version 3.x
  - uCernVM bootloader technology
    - OS loaded from CernVM-FS, package management by RPM
  - Default: SL6 compatible
    - But also SL4, SL5, ...
  - Support for containers (starting v3.4)
- Version 4.x
  - Supporting **systemd** initialization technology
  - Default: SL7 / CC7 compatible

# CernVM-Appliance

- Commissioning of v 3.6
  - Consolidated bootloader (from experience with v4)
  - Full support for x32 ABI
- Consolidation of CernVM 4
- Improved integration / interplay with light virtualization (containers)
- CernVM images available of CERN OpenStack
- Inclusion in LCG builds
- Analytics (Piwik)
  - Better understanding of usage

Technical Student from July 1st

# Data and Software Preservation

- Towards the full picture
  - Versioning and recording of full stack
    - OS, conditions data, software
  - Recreation of the exact environment in a VM
  - Data provenance
- Follow up DPHEP evolution
- We have a central role in DPHEP
  - CernVM(FS) are at the core of the current activities

# Miscellanea

- CernVM-Online
  - Distribute existing raw contexts
  - Consolidate and Improve WebAPI interface
- Web site migration to IT
  - Development version based on Drupal 7 exists
    - Keeps most of current look & feel
    - CSS remain to be optimized
      - May need iteration/help from (external) experts (Nefeli?)
    - Reverse proxy
- Infrastructure
  - Migration of HW in bld 157 to IT

# Miscellanea

- Maintenance of Ioannis' code relevant to us
  - Future of T4T
- Follow-up contacts with potential external users
- Apply for the KT fund
  - Integration with Docker to optimise distribution of (software inside) containers

# PROOF

## Multi-process ROOT

- Lite
  - Multicore machines
- Standard
  - Based on XRootD connection layer
  - Deployed via [PoD](#) (Proof-on-Demand)
  
- Appeared in 7 CHEP2015 abstracts

# PROOF in 2015

- Full support in ROOT v6
- Multi-proc
  - Replacement and extension of PROOF-Lite
    - See also ROOT POW presentation
  - Merged parallelization effort
- Standard PROOF
  - Maintenance mode
    - Made sure to run with old XrdClient
      - Import XrdClient few class headers in ROOT
  - Started to evaluate a XrdSsi (XRooTD) based connection layer

# PROOF in 2016

- Maintain PoD based setups
  - Reducing unused / duplicated code
- Consolidate and complete multiproc interface
  - Full TSelector support, applied to TTree::Draw
- Study / Prototype extension of multi-proc to local, well connected clusters
  - À la iPython parallel
- Prototype an XrdSsi-based distributed ROOT session control layer

# Conferences

- ICFDT 2016
  - Frascati. Italy, March 30th - April 1st 2016
    - Invited talk on Cloud Based Infrastructures (JB)
- CHEP 2016
  - 1 CernVM, 1 multiproc contributions
- ...

# Questions? Comments?

Don't forget the workshop!

<https://indico.cern.ch/event/469775/>