

# **CernVM Users Workshop**

## **Report of Contributions**

Contribution ID: 1

Type: **not specified**

## Opening

*Tuesday 30 January 2018 09:00 (15 minutes)*

**Presenters:** GANIS, Gerardo (CERN); BLOMER, Jakob (CERN)

Contribution ID: 2

Type: **not specified**

## Open Session and Closing

*Wednesday 31 January 2018 18:00 (30 minutes)*

Contribution ID: 3

Type: **not specified**

## CernVM: Ten Years After

*Tuesday 30 January 2018 09:15 (45 minutes)*

Keynote by the founder of the project

**Presenter:** BUNCIC, Predrag (CERN)

**Session Classification:** Technology Outlook

Contribution ID: 4

Type: **not specified**

## CVMFS Build and Release Pipeline Using Docker Microservices

*Wednesday 31 January 2018 16:30 (15 minutes)*

IceCube is a cubic kilometer neutrino detector located at the south pole. CVMFS is a key component to IceCube's Distributed High Throughput Computing analytics workflow for sharing 500GB of software across datacenters worldwide. Building the IceCube software suite across multiple platforms and deploying it into CVMFS has until recently been a manual, time consuming task that doesn't fit well within an agile continuous delivery framework.

Within the last 2 years a plethora of tooling around microservices has created an opportunity to upgrade the IceCube software build and deploy pipeline. We present a framework using Kubernetes to deploy Buildbot. The Buildbot pipeline is a set of pods (docker containers) in the Kubernetes cluster that builds the IceCube software across multiple platforms, tests the new software for critical errors, syncs the software to a containerized CVMFS server, and finally executes a publish. The time from code commit to CVMFS publish has been greatly reduced and has enabled the capability of publishing nightly builds to CVMFS.

**Primary author:** SKARLUPKA, HEATH (University of Wisconsin Madison)

**Co-author:** SCHULTZ, David (University of Wisconsin-Madison)

**Presenter:** SKARLUPKA, HEATH (University of Wisconsin Madison)

**Session Classification:** Feedback from Users

Contribution ID: 5

Type: **not specified**

## **New Features in CernVM-FS**

*Tuesday 30 January 2018 10:30 (30 minutes)*

**Presenter:** POPESCU, Radu (CERN)

**Session Classification:** News from the Development Team

Contribution ID: 6

Type: **not specified**

## **CernVM-FS Graph Driver Plugin for Docker**

*Tuesday 30 January 2018 11:00 (30 minutes)*

**Presenter:** HARDI, Nikola (University of Novi Sad (RS))

**Session Classification:** News from the Development Team

Contribution ID: 7

Type: **not specified**

## News from the CernVM Appliance

*Tuesday 30 January 2018 11:30 (30 minutes)*

**Presenter:** BLOMER, Jakob (CERN)

**Session Classification:** News from the Development Team



Contribution ID: 8

Type: **not specified**

## ALICE Feedback

*Tuesday 30 January 2018 14:00 (10 minutes)*

**Presenter:** BERZANO, Dario (CERN)

**Session Classification:** Feedback from Users

Contribution ID: 9

Type: **not specified**

## Infrastructure Perspective

*Tuesday 30 January 2018 15:20 (20 minutes)*

Combined feedback from Stratum 1 administrators

**Presenter:** CONDURACHE, Catalin (STFC - Rutherford Appleton Lab. (GB))

**Session Classification:** Feedback from Users

Contribution ID: **10**

Type: **not specified**

## **Message based Publishing**

**Session Classification:** Feedback from Users

Contribution ID: 11

Type: **not specified**

## Deploying to CVMFS via Gitlab-CI

*Tuesday 30 January 2018 15:40 (15 minutes)*

**Presenters:** PETRIC, Marko (CERN); SAILER, Andre (CERN)

**Session Classification:** Feedback from Users

Contribution ID: 12

Type: **not specified**

## Running native CVMFS on a Cray supercomputer

*Tuesday 30 January 2018 16:30 (20 minutes)*

**Presenters:** GILA, Miguel (CSCS (CH)); CONCIATORE, Dino (Eidgenoessische Technische Hochschule Zuerich (ETHZ) (CH))

**Session Classification:** Focused Topics

Contribution ID: 13

Type: **not specified**

## HPC in ATLAS

*Tuesday 30 January 2018 16:50 (25 minutes)*

**Presenter:** BENJAMIN, Doug (Duke University (US))

**Session Classification:** Focused Topics

Contribution ID: 14

Type: **not specified**

## Evolution of FUSE and OverlayFS

*Wednesday 31 January 2018 08:30 (40 minutes)*

OverlayFS is the “union filesystem” solution that is now available as part of the Linux kernel. OverlayFS is currently in active development. POSIX compliance, NFS export and improved performance are currently being worked on. There are plans to add user namespace and unprivileged mounting support.

FUSE is a userspace interface for developing filesystems. FUSE started out on Linux, but is now available on other platforms as well. FUSE is mostly in the maintenance mode at the moment, but there are plans for adding user namespace support, improving operation for distributed filesystems and performance improvements to keep in pace with the developments of fast, memory based storage.

This talk aims to give an overview of FUSE and OverlayFS features past, present and future. The target audience is userspace developers familiar with the UNIX filesystem interface.

### **About the speaker:**

Miklos Szeredi is a Linux kernel hacker working for Red Hat. He has been interested in virtual filesystems for a long time, starting several open source projects including Filesystem in Userspace (FUSE) and the Overlay Filesystem. Prior to joining Red Hat, he has worked at SUSE Labs and at Ericsson. Miklos is currently living in a small town near Budapest in Hungary with his family of six, twins being the latest addition.

**Presenter:** SZEREDI, Miklos (Red Hat)

**Session Classification:** Technology Outlook

Contribution ID: 15

Type: **not specified**

## Designing the Git Virtual File System (GVFS)

*Wednesday 31 January 2018 09:10 (40 minutes)*

We've built a virtual file system that enables the Windows team to work in a Git repository that is a few orders of magnitude larger than what Git was previously able to support. In this talk we'll cover a high level overview of the scale challenges we faced with Git, how we designed our virtual file system on top of NTFS, and some of the difficulties we ran into while building a file system that is correct, lazy, and performant.

### **About the speaker:**

Saeed Noursalehi is on the Visual Studio Team Services team at Microsoft, focused on helping some very large teams within Microsoft migrate to Git. Among other things, this means solving some hard scale problems in Git, which is a great source of fun. He also enjoys rock climbing, road biking, and music.

**Primary author:** NOURSALEHI, Saeed (Microsoft)

**Presenter:** NOURSALEHI, Saeed (Microsoft)

**Session Classification:** Technology Outlook



Contribution ID: 16

Type: **not specified**

## Building Reproducible Science with Singularity Containers

*Wednesday 31 January 2018 10:10 (40 minutes)*

One of the biggest problems in scientific HPC is ensuring that results are reproducible. That is, the code a scientist runs locally must be able to run identically on any computational resource. Until recently, the job of ensuring that fell to system administrators who needed to manage a complex web of tools and dependencies on those resources. However, with the introduction of HPC containers via Singularity, the ability to mobilize the compute environment has never been easier. Singularity allows anybody to run their own containers on HPC, ushering in a new era of computational mobility, validity, and reproducibility.

### **About the speaker:**

Michael Bauer first began working with containers at GSI national lab in Darmstadt, Germany, in 2017 while taking a semester off of school at the University of Michigan. Michael met Greg Kurtzer, project lead of Singularity, during his time at GSI and he began contributing heavily to the Singularity project. At the start of summer 2017, Greg hired Michael to work at the Silicon Valley startup RStor, where he continued to work on the Singularity container technology. After 6 months at RStor, the Singularity team left RStor to create their own company, SyLabs, Inc., where Michael, Greg and several other developers now work full time on developing Singularity.

**Presenter:** BAUER, Michael

**Session Classification:** Technology Outlook

Contribution ID: 17

Type: **not specified**

## Tooling for Using Linux

*Wednesday 31 January 2018 10:50 (40 minutes)*

LinuxKit is a framework for building small, modular, immutable Linux systems that was open sourced last year by Docker. It came via a different design process than CernVM but shares much of the same philosophy. This talk looks at similarities and differences, and shows how to construct systems with LinuxKit, and future developments. It will also cover containerd, the new container runtime that LinuxKit and Docker use, and container image distribution.

**About the speaker:**

Justin Cormack is a software engineer working for Docker in Cambridge, UK. He is a maintainer for Docker and LinuxKit, and works across the container ecosystem.

**Presenter:** CORMACK, Justin (Docker)

**Session Classification:** Technology Outlook

Contribution ID: **18**

Type: **not specified**

## Open HTC Content Delivery Network

*Wednesday 31 January 2018 14:30 (20 minutes)*

**Presenter:** DYKSTRA, Dave (Fermi National Accelerator Lab. (US))

**Session Classification:** Focused Topics

Contribution ID: **19**

Type: **not specified**

## **XCache Overview**

*Wednesday 31 January 2018 14:50 (20 minutes)*

**Presenter:** HANUSHEVSKY, Andrew Bohdan (SLAC National Accelerator Laboratory (US))

**Session Classification:** Focused Topics

Contribution ID: 20

Type: **not specified**

## XCache in CernVM-FS

*Wednesday 31 January 2018 15:10 (10 minutes)*

**Primary author:** YANG, Wei (SLAC National Accelerator Laboratory (US))

**Co-author:** POPESCU, Radu (CERN)

**Presenter:** POPESCU, Radu (CERN)

**Session Classification:** Focused Topics

Contribution ID: **21**

Type: **not specified**

## **CernVM-FS for Data**

*Wednesday 31 January 2018 17:00 (20 minutes)*

**Presenters:** BOCKELMAN, Brian Paul (University of Nebraska Lincoln (US)); WEITZEL, Derek John (University of Nebraska Lincoln (US))

**Session Classification:** Feedback from Users

Contribution ID: 22

Type: **not specified**

## Magnum and HNSciCloud

*Wednesday 31 January 2018 17:40 (20 minutes)*

**Presenter:** BRITO DA ROCHA, Ricardo (CERN)

**Session Classification:** Focused Topics

Contribution ID: 23

Type: **not specified**

## Containers in LHCb

**Session Classification:** Focused Topics



Contribution ID: 24

Type: **not specified**

# Provisioning Complex Software Environments

*Tuesday 30 January 2018 17:30 (30 minutes)*

**Presenter:** THAIN, Douglas

**Session Classification:** Focused Topics

Contribution ID: 25

Type: **not specified**

## Automated conversion of Docker images to CVMFS for LIGO and the Open Science Grid

*Tuesday 30 January 2018 15:00 (10 minutes)*

In this lightning talk, I will discuss the development of a webhook-based tool for automatically converting Docker images from DockerHub and private registries to CVMFS filesystems. The tool is highly reliant on previous work by the Open Science Grid for scripted nightly conversion of images from DockerHub.

**Primary author:** Dr DOWNES, Tom (University of Wisconsin-Milwaukee)

**Presenter:** Dr DOWNES, Tom (University of Wisconsin-Milwaukee)

**Session Classification:** Feedback from Users

Contribution ID: 26

Type: **not specified**

## cernatschool.org's use of CVMFS and the CernVM

*Wednesday 31 January 2018 16:45 (15 minutes)*

cernatschool.org is a very small Virtual Organisation made up of secondary school and university students, and participating organisations in the Institute for Research in Schools.

We use CVMFS to deploy dependencies and Python 3 itself for custom software used for analysing radiation data from Medipix detectors. This software is designed for running on GridPP worker nodes, part of the UK based distributed computing grid.

The cernatschool.org VO also uses the CernVM, for job submission and interacting with the grid. The current use for both CVMFS and the CernVM is for facilitating analysis of 3 years worth of data from the LUCID payload on TechDemoSat-1.

The CernVM looks like it could be particularly useful in the future for a standard system for students to use to program and analyse data themselves with, allowing easy access to any software they might need (not necessarily using GridPP compute resources at all).

**Primary author:** FURNELL, Will

**Presenter:** FURNELL, Will

**Session Classification:** Feedback from Users

Contribution ID: 27

Type: **not specified**

## West-Life, Tools for Integrative Structural Biology

*Wednesday 31 January 2018 15:20 (15 minutes)*

Structural biology is part of molecular biology focusing on determining structure of macromolecules inside living cells and cell membranes. As macromolecules determines most of the functions of cells the structural knowledge is very useful for further research in metabolism, physiology to application in pharmacology etc.

As macromolecules are too small to be observed directly by light microscope, there are other methods used to determine the structure including nuclear magnetic resonance (NMR), X-Ray crystallography, cryo electron microscopy and others. Each method has it's advantages and disadvantages in the terms of availability, sample preparation, resolution.

West-Life project has ambition to facilitate integrative approach using multiple techniques mentioned above. As there are already lot of software tools to process data produced by the techniques above, the challenge is to integrate them together in a way they can be used by experts in one technique but not experts in other techniques.

One product of the West-Life project is a data management service - virtual folder. It delivers a uniform way to integrate scattered data from different storage providers.

Another product is a virtual machine, which may allow to launch specific software tools to process user's data in virtual folder.

CernVM with option to be launched with graphical user interface is used as a basic template to contextualize virtual machine with additional structural biology software suites such as CCP4, Scipion and others. CernVM-FS is used to distribute updates of structural biology software suites as well as West-Life specific services - virtual folder and newly repository.

The virtual machine templates are available in EGI's APP DB as well as within STFC cloud computing infrastructure.

**Primary author:** KULHANEK, Tomas (STFC Daresbury Laboratory)

**Co-author:** CONDURACHE, Catalin (STFC-Rutherford Appleton Laboratory (GB))

**Presenter:** KULHANEK, Tomas (STFC Daresbury Laboratory)

**Session Classification:** Focused Topics

Contribution ID: **28**

Type: **not specified**

## **ATLAS Feedback**

*Tuesday 30 January 2018 14:10 (10 minutes)*

**Presenter:** DE SALVO, Alessandro (Sapienza Universita e INFN, Roma I (IT))

**Session Classification:** Feedback from Users

Contribution ID: 29

Type: **not specified**

## CMS Feedback

*Tuesday 30 January 2018 14:20 (10 minutes)*

**Presenter:** MUZAFFAR, Shahzad Malik (Fermi National Accelerator Lab. (US))

**Session Classification:** Feedback from Users

Contribution ID: **30**

Type: **not specified**

## **LHCb Feedback**

*Tuesday 30 January 2018 14:30 (10 minutes)*

**Presenter:** COUTURIER, Ben (CERN)

**Session Classification:** Feedback from Users

Contribution ID: **31**

Type: **not specified**

## **Belle II Feedback**

*Tuesday 30 January 2018 14:40 (10 minutes)*

**Presenter:** SOBIE, Randy (University of Victoria (CA))

**Session Classification:** Feedback from Users



Contribution ID: **32**

Type: **not specified**

## **EUCLID Feedback**

*Tuesday 30 January 2018 14:50 (10 minutes)*

**Presenters:** Mr PONCET, Maurice (CNES); LE BOULC'H, Quentin (CNRS)

**Session Classification:** Feedback from Users

Contribution ID: **33**

Type: **not specified**

## **LIGO Feedback**

**Presenter:** DOWNES, Tom (University of Wisconsin-Milwaukee)

**Session Classification:** Feedback from Users

Contribution ID: 34

Type: **not specified**

## SuperMUC

*Tuesday 30 January 2018 17:15 (15 minutes)*

**Presenter:** WALKER, Rodney (Ludwig Maximilians Universitat (DE))

**Session Classification:** Focused Topics

Contribution ID: 35

Type: **not specified**

## **CernVM facilitates offline processing on the ATLAS HLT farm**

*Wednesday 31 January 2018 15:35 (10 minutes)*

**Presenter:** BERGHAUS, Frank (University of Victoria (CA))

**Session Classification:** Focused Topics

Contribution ID: 36

Type: **not specified**

## Compute Canada Software Installation and Distribution

*Wednesday 31 January 2018 17:20 (20 minutes)*

**Presenter:** TAYLOR, Ryan (University of Victoria (CA))

**Session Classification:** Focused Topics

Contribution ID: 37

Type: **not specified**

## SFT Feedback

*Tuesday 30 January 2018 15:10 (10 minutes)*

**Presenter:** MENDEZ LORENZO, Patricia (CERN)

**Session Classification:** Feedback from Users

Contribution ID: **38**

Type: **not specified**

## **Containerized CernVM-FS Server**

*Wednesday 31 January 2018 15:45 (10 minutes)*

**Presenter:** VAN DER STER, Dan (CERN)

**Session Classification:** Focused Topics