

# CernVM Workshop 2024



## Report of Contributions

Contribution ID: 2

Type: **not specified**

## Welcome and Practicalities

*Monday 16 September 2024 14:00 (15 minutes)*

**Presenter:** VOLKL, Valentin (CERN)

**Session Classification:** Monday afternoon: CERN team presentations

Contribution ID: 3

Type: **not specified**

## Workshop Introduction

*Monday 16 September 2024 14:15 (15 minutes)*

**Presenters:** GANIS, Gerardo (CERN); BUNCIC, Predrag (CERN)

**Session Classification:** Monday afternoon: CERN team presentations

Contribution ID: 4

Type: **not specified**

## **Zstd: A new compression algorithm for CVMFS**

*Monday 16 September 2024 16:00 (20 minutes)*

**Presenter:** PROMBERGER, Laura (CERN)

**Session Classification:** Monday afternoon: CERN team presentations

Contribution ID: 5

Type: **not specified**

## **CernVM-FS at 15: technological retrospective**

*Monday 16 September 2024 14:30 (30 minutes)*

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

**Session Classification:** Monday afternoon: CERN team presentations

Contribution ID: 6

Type: **not specified**

## New Developments in LibFUSE

*Tuesday 17 September 2024 09:00 (40 minutes)*

**Presenter:** SCHUBERT, Bernd

**Session Classification:** Tuesday morning: FUSE and Varnish guest speakers

Contribution ID: 7

Type: **not specified**

## Jump Trading and Varnish

*Tuesday 17 September 2024 10:55 (40 minutes)*

**Presenters:** HARVEY, Matt (Jump Trading); CANCE, Stephane; Mr BOUDEBOUDA, Walid

**Session Classification:** Tuesday morning: FUSE and Varnish guest speakers

Contribution ID: 9

Type: **not specified**

## **ATLAS Installations on CVMFS**

*Tuesday 17 September 2024 16:25 (20 minutes)*

**Presenter:** DE STEFANO JR, John Steven (Brookhaven National Laboratory (US))

**Session Classification:** Tuesday afternoon: experiments and sites



Contribution ID: **10**

Type: **not specified**

## **CMS deployments on CernVM-FS**

**Presenter:** VALENZUELA RAMIREZ, Andrea (CERN)

Contribution ID: 11

Type: **not specified**

## **CVMFS Container Snapshotter Benchmarking**

**Presenter:** FATOUROS, Max (Paul Scherrer Institute (CH))

**Session Classification:** Monday afternoon: CERN team presentations

Contribution ID: 12

Type: **not specified**

## **New overlayfs features in cvmfs\_server; Fuse-t on MacOS**

*Monday 16 September 2024 16:20 (20 minutes)*

**Presenter:** BIELIKOV, Yuri (Taras Shevchenko National University of Kyiv (UA))

**Session Classification:** Monday afternoon: CERN team presentations

Contribution ID: 13

Type: **not specified**

## CVMFS + Kubernetes

**Presenter:** THUNDIYIL, Amal Santosh (Sardar Patel Institute Technology of Technology (SPIT) (IN))

Contribution ID: **14**

Type: **not specified**

## **Cvmfs at GSI**

**Presenter:** FLEISCHER, Soren Lars Gerald (GSI - Helmholtzzentrum für Schwerionenforschung GmbH (DE))

Contribution ID: 15

Type: **not specified**

## The European Environment for Scientific Software Installations (EESSI)

The European Environment for Scientific Software Installations (EESSI, pronounced as “easy”) is a collaboration between different HPC sites and industry partners, with the common goal to set up a shared repository of scientific software installations that can be used on a variety of systems, regardless of which flavor/version of Linux distribution or processor architecture is used, or whether it is a full-size HPC cluster, a cloud environment or a personal workstation, and without compromising on the performance of the software.

The EESSI codebase (<https://github.com/EESSI>) is open source and heavily relies on various other open-source software projects, including Ansible, archspec, CernVM-FS, Cluster-in-the-Cloud, EasyBuild, Gentoo Prefix, Lmod, ReFrame, Singularity, and Terraform.

The concept of the EESSI project was inspired by the Compute Canada software stack, and consists of three main layers:

- a filesystem layer leveraging CernVM-FS, to globally distribute the EESSI software stack;
- a compatibility layer using Gentoo Prefix, to ensure compatibility with different client operating systems (different Linux distributions, macOS, Windows Subsystem for Linux);
- a software layer, hosting optimized installations of scientific software along with required dependencies, which were built for different processor architectures, and where archspec, EasyBuild and Lmod are leveraged.

In this talk, we will introduce you to EESSI, outline the use cases it enables, present recent developments, and give an outlook to the promising future of EESSI.

### About the speaker

Kenneth Hoste is a computer scientist and FOSS enthusiast from Belgium. He holds a Masters (2005) and PhD (2010) in Computer Science from Ghent University. His dissertation topic was “Analysis, Estimation and Optimization of Computer System Performance Using Machine Learning”.

Since October 2010, he is a member of the HPC team at Ghent University where he is mainly responsible for user support & training. As a part of his job, he is also the lead developer and release manager of EasyBuild, a software build and installation framework for (scientific) software on High Performance Computing (HPC) systems.

Since 2020, he is actively involved with the European Environment for Scientific Software Installations (EESSI) project, which aims to provide a central stack of scientific software installations that can be used across a wide range of systems, without compromising on performance.

In his free time, he is a family guy and a fan of loud music, frequently attending gigs and festivals. He enjoys helping people & sharing his expertise, and likes joking around. He has a weak spot for stickers.

**Presenter:** RÖBLITZ, Thomas

Contribution ID: **16**

Type: **not specified**

## **LHCb installations on CVMFS**

*Tuesday 17 September 2024 14:20 (20 minutes)*

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

**Session Classification:** Tuesday afternoon: experiments and sites

Contribution ID: **18**

Type: **not specified**

## **CVMFS-CSI: CVMFS in Kubernetes Deployments at CERN-IT**

*Tuesday 17 September 2024 17:40 (20 minutes)*

**Presenter:** MUNDAY, Jack Charlie

**Session Classification:** Tuesday afternoon: experiments and sites



Contribution ID: 20

Type: **not specified**

## Use of FUSE in EOS

*Tuesday 17 September 2024 10:30 (25 minutes)*

**Presenter:** SINDRILARU, Elvin Alin (CERN)

**Session Classification:** Tuesday morning: FUSE and Varnish guest speakers

Contribution ID: 23

Type: **not specified**

## **Pelican and the Open Science Data Federation**

*Tuesday 17 September 2024 15:15 (20 minutes)*

**Presenter:** BOCKELMAN, Brian Paul (University of Wisconsin Madison (US))

**Session Classification:** Tuesday afternoon: experiments and sites

Contribution ID: 27

Type: **Presentation**

## **CVMFS in EUCLID**

*Tuesday 17 September 2024 15:35 (20 minutes)*

**Presenter:** CASENOVE, Pierre (CNES)

**Session Classification:** Tuesday afternoon: experiments and sites

Contribution ID: 29

Type: **Presentation**

## **CVMFS Stratum 1 Performance Benchmarks**

*Tuesday 17 September 2024 14:55 (20 minutes)*

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

**Session Classification:** Tuesday afternoon: experiments and sites

Contribution ID: **30**

Type: **not specified**

## CernVM-FS Hacking Session

*Monday 16 September 2024 10:00 (2 hours)*

On-site session for discussions and coding. The core team will be present and participants are welcome to join us at the venue.

**Presenters:** BLOMER, Jakob (CERN); PROMBERGER, Laura (CERN); VOLKL, Valentin (CERN)

Contribution ID: 32

Type: **not specified**

## Varnish Hands-On

*Wednesday 18 September 2024 09:30 (1 hour)*

**Presenter:** Mr BOUDEBOUDA, Walid

**Session Classification:** Discussion session / Hands-On

Contribution ID: **34**

Type: **not specified**

## **Discussion session / Hands-On**

*Wednesday 18 September 2024 11:20 (40 minutes)*

**Session Classification:** Discussion session / Hands-On

Contribution ID: 35

Type: **not specified**

## Workshop Closing

*Wednesday 18 September 2024 12:00 (20 minutes)*

**Presenter:** VOLKL, Valentin (CERN)

**Session Classification:** Discussion session / Hands-On



Contribution ID: 39

Type: **Presentation**

## CernVM-FS makes software EESSI to use

*Tuesday 17 September 2024 14:00 (20 minutes)*

`/cvmfs/software.eessi.io` is the production CernVM-FS repository of the European Environment for Scientific Software Installations (EESSI, pronounced as “easy”, see [eessi.io/docs](https://eessi.io/docs)). After several pilot iterations, EESSI has released its first production software stack in the fall of 2023. The growing software stack already includes hundreds of software installations optimised for a variety of CPU microarchitectures (x86\_64 and aarch64 CPU families) that are ready to be used by anyone on any (Linux) machine anywhere in the world in a matter of minutes.

In this talk, we will present the status of EESSI (progress we made since the last CernVM workshop in 2022) and the features we are currently working on. Particularly, we will cover our efforts to automate the building of software for multiple CPU architectures, ongoing work to add support for more CPU architectures (AMD Genoa, Intel Sapphire Rapids and Fujitsu A64FX) and GPUs (NVIDIA and AMD), introduce our growing suite of portable regression tests, describe how we collaborate with software developers (including the recently created `dev.eessi.io` CernVM-FS repository) as well as experiences we made through training and support activities.

We will also show the progress for RISC-V CPUs (available via the CernVM-FS repository `riscv.eessi.io`), work on enabling and fine-tuning EESSI on (Euro)HPC sites and discuss plans for improved monitoring.

**Primary authors:** DRÖGE, Bob (University of Groningen, The Netherlands); HOSTE, Kenneth (Ghent University, Belgium); PEETERS, Lara (Ghent University, Belgium); RÖBLITZ, Thomas

**Presenter:** RÖBLITZ, Thomas

**Session Classification:** Tuesday afternoon: experiments and sites

Contribution ID: 40

Type: **Presentation**

## CVMFS@GSI

*Tuesday 17 September 2024 17:25 (15 minutes)*

GSI is and has been operating a cvmfs instance for many years.

It is used primarily to make prerequisites for physics simulations and analysis of GSI and future FAIR experiments available on our compute cluster.

These include items such as precompiled executables, Spack packages, as well as Apptainer images.

The current cvmfs setup at GSI will be outlined in this talk.

**Primary author:** FLEISCHER, Soren Lars Gerald (GSI - Helmholtzzentrum für Schwerionenforschung GmbH (DE))

**Presenter:** FLEISCHER, Soren Lars Gerald (GSI - Helmholtzzentrum für Schwerionenforschung GmbH (DE))

**Session Classification:** Tuesday afternoon: experiments and sites

Contribution ID: 42

Type: **Presentation**

## New CernVM-FS use cases at CMS

*Tuesday 17 September 2024 16:45 (20 minutes)*

CernVM-FS remains a central service for the distribution of the CMS Offline Software (CMSSW). Traditional use cases include the distribution of CMSSW releases and container images. It also plays a crucial role in the deployment of Integration Builds (IB) and pull request testing as part of the CMSSW CI/CD workflow.

In this contribution, we present three new use cases of CernVM-FS for CMS. Firstly, the distribution of gridpacks to accelerate Monte Carlo generation. Gridpacks are pre-computed diagrams distributed in the form of tarballs that are used as lookup files. The CernVM-FS ingest command allows the serving of already-untarred gridpacks, reducing the load of unpacking them on computing sites. Secondly, the use of opportunistic HPC resources using the `singcvmfs` command in the `cvmfsexec` package. The CMS Offline infrastructure already makes use of AMD GPUs from the LUMI HPC in Finland, thanks to FUSE pre-mounting. Finally, we are currently building the CMSSW software stack for the RISC-V architecture. To include this new architecture in the CMSSW IBs, deployment to CernVM-FS using emulation on the publisher nodes has been implemented.

These new use cases at CMS highlight the importance of CernVM-FS in streamlining workflows and improving the performance of the CMS Offline infrastructure.

**Primary author:** VALENZUELA RAMIREZ, Andrea (CERN)

**Co-author:** MUZAFFAR, Malik Shahzad (CERN)

**Presenter:** VALENZUELA RAMIREZ, Andrea (CERN)

**Session Classification:** Tuesday afternoon: experiments and sites

Contribution ID: 43

Type: **not specified**

# Optimizing CERN SFT Nightlies Publication through Parallelized CVMFS Gateway Installation

*Monday 16 September 2024 17:00 (20 minutes)*

The EP-SFT Stacks team provides LCG software stack nightly builds for 50 different combinations of STACK and PLATFORM.

Historically, the publication of nightlies on CVMFS utilized a single machine, resulting in extended processing times that lasted until the end of the working day. To speed up the publication, and enable the distribution of a larger set of platforms every day, the LCG stack publication was moved to distributing the CVMFS gateway publication across multiple machines, thus reducing the CVMFS installation bottleneck.

Because we are constrained by the requirement of keeping the historic file system hierarchy, not all parts of the publication can be parallelized.

This talk will describe the development of a Jenkins pipeline to parallelize the publication of the LCG software stacks. It will explain how we use the Jenkins locking mechanism in combination with gateway leases to manage parallelization and mitigate potential file system conflicts. It will show in detail which parts of our publication stages, such as the installation of packages, remain partially serialized due to their structure in the file system, while other stages, such as the creation and installation of platform-specific views, are now executed in parallel.

**Primary author:** Mr EHMANN, Tim (CERN EP-SFT)

**Presenter:** Mr EHMANN, Tim (CERN EP-SFT)

**Session Classification:** Monday afternoon: CERN team presentations

Contribution ID: 44

Type: **Presentation**

## Efficient and fast container execution using image snapshotters

*Monday 16 September 2024 16:40 (20 minutes)*

A large fraction of computing workloads in high-energy and nuclear physics is executed using software containers. For physics analysis use, such container images often have sizes of several gigabytes. Executing a large number of such jobs in parallel on different compute nodes efficiently, demands the availability and use of caching mechanisms and image loading techniques to prevent network saturation and significantly reduce startup time. Using the industry-standard containerd container runtime for pulling and running containers, enables the use of various so-called snapshotter plugins that “lazily” load container images. We present a quantitative comparison of the performance of the CVMFS, SOCI, and Stargz snapshotter plugins. Furthermore, we also evaluate the user-friendliness of such approaches and discuss how such seamlessly containerised workloads contribute to the reusability and reproducibility of physics analyses.

**Primary authors:** THUNDIYIL, Amal Santosh (CERN); LANGE, Clemens (Paul Scherrer Institute (CH)); FEICHTINGER, Derek (Paul Scherrer Institute (CH)); BLOMER, Jakob (CERN); FATOUROS, Max (Paul Scherrer Institute (CH)); VOLKL, Valentin (CERN)

**Presenter:** FATOUROS, Max (Paul Scherrer Institute (CH))

**Session Classification:** Monday afternoon: CERN team presentations

Contribution ID: 45

Type: **not specified**

## **CernVM-FS: Status and Plans**

*Monday 16 September 2024 15:00 (30 minutes)*

**Presenter:** VOLKL, Valentin (CERN)

**Session Classification:** Monday afternoon: CERN team presentations

Contribution ID: 46

Type: **not specified**

## FUSE Kernel Developments: Passthrough mode

*Tuesday 17 September 2024 09:40 (20 minutes)*

**Presenter:** SZEREDI, Miklos

**Session Classification:** Tuesday morning: FUSE and Varnish guest speakers

Contribution ID: 47

Type: **not specified**

## ALICE installations on CVMFS

*Tuesday 17 September 2024 17:05 (20 minutes)*

**Presenter:** EULISSE, Giulio (CERN)

**Session Classification:** Tuesday afternoon: experiments and sites



Contribution ID: 48

Type: **not specified**

## Containerizing CVMFS publishing

*Monday 16 September 2024 17:20 (10 minutes)*

**Presenter:** THUNDIYIL, Amal Santosh (CERN)

**Session Classification:** Monday afternoon: CERN team presentations

Contribution ID: 49

Type: **not specified**

## Parallel Publishing for sft.cern.ch

**Presenter:** EHMANN, Tim Nicolai (University of Applied Sciences (DE))