

CernVM Workshop 2022



Report of Contributions

Contribution ID: 1

Type: **Presentation**

Decarbonizing Scientific Computing

Scientific computing uses a tremendous amount of energy, and given the location of most HPC centers, results in a similarly large amount of CO₂ emissions. In the US, for example, in 2019, every MWh of generated power on average led to 0.7 metric tons of CO₂ emissions. To address the huge carbon footprint of computing Lancium Compute is building, low carbon, renewable-energy-driven data centers in the Great Plains of the United States.

Because the wind does not always blow, and the sun does not always shine, our data centers must be able to rapidly ramp our computing and electrical load up and down in order to balance the electrical grid. Not all applications are suitable for this sort of load management. However, many batch based scientific jobs, e.g., High Throughput Computing jobs, are ideal.

In 2021 we began to work with the Open Science Grid to support their HTC load. As part of that integration effort we added support for CVMFS for all containerized jobs. Later work with the US CMS and US ATLAS team led us to further deploy a hierarchical Squid architecture to support Frontier.

In this talk I will briefly present electrical grid basics, and explain how the characteristics of renewables make them difficult to integrate into the grid. I follow with a discussion on how controllable loads can solve these problems, and how computing can be an excellent controllable load. I will then describe our quality of service model and multi-site system architecture with CVMFS to support high throughput jobs single node jobs and low-degree parallel jobs at our clean compute campuses in Texas.

Primary author: GRIMSHAW, Andrew (Lancium Compute)

Presenter: GRIMSHAW, Andrew (Lancium Compute)

Contribution ID: 2

Type: **not specified**

Workshop Introduction

Monday, 12 September 2022 14:00 (15 minutes)

Presenters: HESTER, Mary; VAN DOK, Dennis; BLOMER, Jakob (CERN)

Contribution ID: 3

Type: **not specified**

CernVM-FS: Status and Plans

Monday, 12 September 2022 14:15 (20 minutes)

Presenter: BLOMER, Jakob (CERN)

Contribution ID: 4

Type: **not specified**

CernVM 5: A fully containerized CernVM

Monday, 12 September 2022 14:55 (20 minutes)

Presenter: EBERHARDT, Jakob Karl (University of Applied Sciences (DE))

Contribution ID: 5

Type: **not specified**

CernVM-FS publisher on Kubernetes

Monday, 12 September 2022 14:35 (20 minutes)

Presenter: VALENZUELA RAMIREZ, Andrea (CERN)

Contribution ID: 6

Type: **not specified**

Lazy Container Image Distribution With eStargz And P2P Image Sharing on IPFS

Tuesday, 13 September 2022 09:30 (20 minutes)

Pulling images is known as one of the time-consuming steps during starting up containers. In this talk, Kohei will share the approach of speeding up container startup using eStargz image format and recent works around eStargz.

eStargz allows container runtimes to start containers without waiting for the pull completion. This has been available on a variety of tools in the community, including Kubernetes, k3s, containerd, CRI-O, Podman, BuildKit, Kaniko, etc.

The talk will also cover IPFS-based P2P image distribution with containerd and its combination with eStargz-based lazy pulling.

About the speaker

Kohei Tokunaga is a software engineer at NTT Corporation, a Japan-based telecommunication company. He is a reviewer of CNCF containerd and a maintainer of BuildKit. He is also a maintainer of stargz-snapshotter subproject in containerd.

Presenter: TOKUNAGA, Kohei

Contribution ID: 7

Type: **not specified**

First experiences using CVMFS in Microsoft Azure

Monday, 12 September 2022 17:05 (20 minutes)

About the speaker

Principal Program Manager at Microsoft, helping customers discover, test and improve HPC technologies on Azure.

Presenter: MEILAND, Hugo (Microsoft)

Contribution ID: 8

Type: **not specified**

Exascale data processing with CVMFS

Monday, 12 September 2022 15:45 (20 minutes)

About the speaker:

Matt Harvey is a Linux production engineer at Jump Trading, working on high performance computing in a global, data-intensive environment. His background is in research computing services, in both the public sector, at Imperial College London, and the pharmaceutical industry with Acellera.

Presenter: HARVEY, Matt (Jump Trading)

Contribution ID: 9

Type: **not specified**

ATLAS Installations on CVMFS

Tuesday, 13 September 2022 12:00 (20 minutes)

Presenter: VICKEY BOERIU, Oana (University of Sheffield (UK))

Contribution ID: **10**

Type: **not specified**

CMS deployments on CernVM-FS

Tuesday, 13 September 2022 11:40 (20 minutes)

Presenter: VALENZUELA RAMIREZ, Andrea (CERN)

Contribution ID: **11**

Type: **not specified**

WLCG Stratum 1 Status and Plans

Monday, 12 September 2022 16:45 (20 minutes)

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

Contribution ID: 12

Type: **not specified**

CVMFS usage and performance for LHCb

Monday, 12 September 2022 16:25 (20 minutes)

Presenter: COUTURIER, Ben (CERN)

Contribution ID: 13

Type: **not specified**

ALICE's software infrastructure

Monday, 12 September 2022 16:05 (20 minutes)

Presenter: Mr WILKEN, Timo (CERN)

Contribution ID: **14**

Type: **not specified**

Cvmfs at GSI

Tuesday, 13 September 2022 10:30 (20 minutes)

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)

Tuesday, 13 September 2022 14:00 (20 minutes)

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: HOSTE, Kenneth

Contribution ID: 16

Type: **not specified**

CVMFS in Canadian Advanced Research Computing

Tuesday, 13 September 2022 16:40 (20 minutes)

About the speaker

Pursuing his passion to figure out the inner workings of complex systems, Ryan Taylor completed a M.Sc. in particle physics with the ATLAS experiment at the Large Hadron Collider at CERN in 2009, and joined the University of Victoria Research Computing Services team shortly thereafter, where he is now a senior Advanced Research Computing (ARC) specialist. He has expertise in areas including grid computing and storage, cloud computing, content distribution technologies such as CVMFS, and container orchestration technologies in the Kubernetes ecosystem. As the leader of the Canadian CVMFS National Team, he guides the architecture and deployment strategy of the infrastructure which provides the Compute Canada ARC software stack and other content to sites and users across Canada and the global research community.

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

Contribution ID: 17

Type: **not specified**

Data distribution in LIGO/VIRGO

Contribution ID: **18**

Type: **not specified**

Workshop Closing

Tuesday, 13 September 2022 17:00 (20 minutes)

Contribution ID: **19**

Type: **not specified**

Infrastructure support for NL researchers

Contribution ID: 20

Type: **not specified**

Decarbonizing Scientific Computing

Tuesday, 13 September 2022 11:20 (20 minutes)

Scientific computing uses a tremendous amount of energy, and given the location of most HPC centers, results in a similarly large amount of CO₂ emissions. In the US, for example, in 2019, every MWh of generated power on average led to 0.7 metric tons of CO₂ emissions. To address the huge carbon footprint of computing Lancium Compute is building, low carbon, renewable-energy-driven data centers in the Great Plains of the United States.

Because the wind does not always blow, and the sun does not always shine, our data centers must be able to rapidly ramp our computing and electrical load up and down in order to balance the electrical grid. Not all applications are suitable for this sort of load management. However, many batch based scientific jobs, e.g., High Throughput Computing jobs, are ideal.

In 2021 we began to work with the Open Science Grid to support their HTC load. As part of that integration effort we added support for CVMFS for all containerized jobs. Later work with the US CMS and US ATLAS team led us to further deploy a hierarchical Squid architecture to support Frontier.

In this talk I will briefly present electrical grid basics, and explain how the characteristics of renewables make them difficult to integrate into the grid. I follow with a discussion on how controllable loads can solve these problems, and how computing can be an excellent controllable load. I will then describe our quality of service model and multi-site system architecture with CVMFS to support high throughput jobs single node jobs and low-degree parallel jobs at our clean compute campuses in Texas.

Presenter: GRIMSHAW, Andrew (Lancium Compute)

Contribution ID: 21

Type: **not specified**

Software infrastructure of XENONnT

Tuesday, 13 September 2022 14:20 (20 minutes)

Presenter: Mr ANGEVAARE, Joran (Nikhef - University of Amsterdam (the Netherlands))

Contribution ID: 22

Type: **not specified**

CVMFS for KM3NeT

Tuesday, 13 September 2022 14:40 (20 minutes)

Presenter: BOUWHUIS, Mieke (NIKHEF)

Contribution ID: 23

Type: **not specified**

ATLAS Cloud R&D, Kubernetes and CVMFS

Tuesday, 13 September 2022 16:00 (20 minutes)

Different groups, sites and experiments in the WLCG community have started using Kubernetes to manage services, implement novel analysis facilities or run batch services . Despite being in a native containerised environment, many of these use cases depend on CVMFS to stay compatible with the existing Grid model or to benefit from a well established software distribution model. Multiple of these groups therefore implemented their own helm charts and images to install the CVMFS client on Kubernetes.

In the case of the ATLAS Cloud R&D we use Kubernetes native batch capabilities to integrate multiple research and commercial clouds with the PanDA ecosystem. The clusters can be quickly scaled up and down, and the nodes are usually fully exploited. The CVMFS client needs to run like a clock to avoid expensive job failures due to hanging clients. We spent a significant amount of time comparing and customising some of the existing CVMFS clients until we reached a satisfactory situation.

This contribution will give an overview of some CVMFS clients, describe the most common issues and propose where the expertise of the CVMFS team and an official Kubernetes client would make a difference for the WLCG community.

Presenter: BARREIRO MEGINO, Fernando Harald (University of Texas at Arlington)

Contribution ID: 24

Type: **not specified**

Harbor Registry at CERN: Status and Enhancements

Tuesday, 13 September 2022 09:50 (20 minutes)

CERN offers a centralized OCI registry at registry.cern.ch, based on the Harbor project and available to the whole community. In addition to the standard container registry functionality, Harbor adds support for any kind of OCI artifact (Helm Charts, ML Models, etc) as well as support for proxy caches to external registries, automated replication between multiple registry instances, vulnerability checks, image signing, webhook integration, among many others.

This talk will summarize the status of our deployment, as well as giving more details around recently added functionality including multiple vulnerability checks using different CVE sources and tools and support for sigstore for image signing and verification.

Presenter: ROCHA, Ricardo (CERN)

Contribution ID: 25

Type: **not specified**

CSI Driver for CernVM-FS

Tuesday, 13 September 2022 10:10 (20 minutes)

Presenter: VASEK, Robert (CERN)

Contribution ID: 26

Type: **Presentation**

Lazy Container Image Distribution With eStargz And P2P Image Sharing on IPFS

Pulling images is known as one of the time-consuming steps during starting up containers. In this talk, Kohei will share the approach of speeding up container startup using eStargz image format and recent works around eStargz.

eStargz allows container runtimes to start containers without waiting for the pull completion. This has been available on a variety of tools in the community, including Kubernetes, k3s, containerd, CRI-O, Podman, BuildKit, Kaniko, etc.

The talk will also cover IPFS-based P2P image distribution with containerd and its combination with eStargz-based lazy pulling.

Primary author: TOKUNAGA, Kohei

Presenter: TOKUNAGA, Kohei

Contribution ID: 27

Type: **Presentation**

A new UK StashCache at Edinburgh for DUNE

Tuesday, 13 September 2022 16:20 (20 minutes)

The Deep Underground Neutrino Experiment (DUNE) utilizes CVMFS and CVMFS StashCache to distribute both its software stack and reference files for distributed computing workflows. DUNE utilizes CVMFS as it provides a read-only POSIX interface to StashCache, with the redundant features, e.g. built-in GeoIP locating, rate monitoring, and fallback in failures. During Dec-'21 to Jan-'22, it was found that DUNE HTC grid jobs were suffering from low CPU efficiency in the UK, due to slow access of data from a StashCache instance at Cardiff, which was intended only for LIGO jobs. A new UK StashCache instance at Edinburgh was implemented to solve this problem for DUNE as well as the other OSG experiments except LIGO in the UK. This talk will be from the perspective of the DUNE use case, to introduce the motivation, diagnosis of the inefficiency, deployment process, and current status/performance of this StashCache instance at Edinburgh.

Primary author: YUAN, Wenlong (Edinburgh University)

Presenter: YUAN, Wenlong (Edinburgh University)

Contribution ID: **28**

Type: **not specified**

CernVM-FS Hacking Session

Presenters: BLOMER, Jakob (CERN); EBERHARDT, Jakob Karl (University of Applied Sciences (DE)); PROMBERGER, Laura (CERN)

Contribution ID: 29

Type: **Presentation**

CVMFS mix at CERN

Tuesday, 13 September 2022 15:00 (20 minutes)

The CVMFS deployment at CERN is constant evolution and adjustment to follow user's needs. This talk briefly describes some of the main points of the deployment of CVMFS at CERN, and the challenges that are involved.

Primary author: FURANO, Fabrizio (CERN)

Presenter: FURANO, Fabrizio (CERN)

Contribution ID: **30**

Type: **not specified**

CernVM-FS Hacking Session

Monday, 12 September 2022 09:00 (3 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); EBERHARDT, Jakob Karl (University of Applied Sciences (DE)); PROMBERGER, Laura (CERN)