

CernVM Workshop 2019



Report of Contributions

Contribution ID: 1

Type: **not specified**

Wrap-Up

Tuesday, June 4, 2019 5:45 PM (30 minutes)

Contribution ID: 2

Type: **not specified**

Exploring Cloudflare Workers

Tuesday, June 4, 2019 9:00 AM (45 minutes)

Cloudflare Workers is a serverless computing platform optimized to minimize latency to end users. Cloudflare runs every guest function in each server of its network's 175 points of presence, meaning code runs as close to the end user as possible, and is not confined to geographic regions. This requires overcoming a scalability challenge faced by container-based platforms: how to reduce each function's overhead enough to deploy them universally to a global server fleet.

The V8 JavaScript engine contains a solution: Isolates, a lightweight sandboxing technology. V8 Isolates allow Cloudflare Workers to use JavaScript and WebAssembly modules as serverless functions, minimizing overhead and providing a familiar language environment for web application developers. Workers reinforces this familiarity by implementing standardized JavaScript APIs found in browsers.

This talk examines the design of Cloudflare Workers, how it fits into the serverless computing landscape, and the problems Worker scripts can solve.

About the Author

Harris Hancock is a systems engineer who helps implement the Cloudflare Workers runtime environment, with a particular focus on the JavaScript API. He previously wrote communications middleware for an educational robotics startup, during which time he became a regular contributor to the Cap'n Proto RPC library. It was this interest in protocols and systems programming which lured him to Cloudflare in 2017.

Presenter: HANCOCK, Harris (Cloudflare)

Session Classification: Technology Outlook

Contribution ID: 3

Type: **not specified**

The evolution of the HPC facility at JSC

Tuesday, June 4, 2019 9:45 AM (45 minutes)

The Jülich Supercomputing Centre (JSC) operates a large-scale network-, data- and compute-infrastructure for scientific research covering a broad spectrum of use cases from scalability-focused large-scale simulations to community-specific services with high-throughput requirements. In this talk we will discuss the design of JSC's facility addressing such different use cases through a mixture of collocated and dedicated resources. Based on currently on-going, exemplary projects with different scientific communities, we will discuss JSC's future systems strategy in view of the on-going exascale race, the need for novel storage management capabilities and upcoming requirements for support of new system usage patterns.

About the Speaker

Dorian Krause leads the High Performance Systems division at the Jülich Supercomputing Centre at Forschungszentrum Jülich. His group is responsible for the operation of the two major supercomputers JURECA and JUWELS and the primary storage infrastructure JUST as well as the co-design and implementation of new compute and data services based on these systems.

Presenter: KRAUSE, Dorian (Jülich Supercomputing Centre)

Session Classification: Technology Outlook

Contribution ID: 4

Type: **not specified**

Using CVMFS for User Analysis Code Distribution

Monday, June 3, 2019 4:00 PM (20 minutes)

CVMFS is primarily used on the grid for distributing released experiment code, and it is expected that only a small number of people in each experiment manage that code. Fermilab has implemented a system for publishing temporary user analysis code in CVMFS, with repositories that are shared by many people and a server API that accepts tarballs from users in multiple experiments. Updates are expedited so they become available on worker nodes within a few minutes. The system is integrated with Fermilab's job submission system. This talk will describe the system and explain how it could be deployed at other sites and used by other job submission systems.

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

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

Session Classification: User Stories

Contribution ID: 5

Type: **not specified**

Welcome Note

Monday, June 3, 2019 9:00 AM (15 minutes)

Presenter: HEMMER, Frederic (CERN)

Session Classification: News from the Developers

Contribution ID: 6

Type: **not specified**

Workshop Introduction

Monday, June 3, 2019 9:15 AM (15 minutes)

Presenters: GANIS, Gerardo (CERN); BLOMER, Jakob (CERN); POPESCU, Radu (CERN)

Session Classification: News from the Developers

Contribution ID: 7

Type: **not specified**

Using CVMFS on a distributed Kubernetes cluster - The PRP experience

Monday, June 3, 2019 4:20 PM (20 minutes)

The Pacific Research Platform (PRP) is operating a Kubernetes cluster that manages over 2.5k CPU cores and 250 GPUs. The compute nodes are distributed over several locations, mostly in California, all connected with high speed networking of 10Gbps or higher.

In order to support OSG users, we were in need of providing CVMFS support in the Kubernetes cluster. Since user containers run unprivileged, CVMFS could not be used from the user containers themselves. The chosen path forward was to provide CVMFS as a mountable Kubernetes volume. We started with the CERN-provided CSI plugin, but had to make some changes to make it to work in our environment.

In this talk, we will detail the actual setup used, as well as the operational experience of the last 4 months. The perceived missing functionality will also be outlined.

Primary authors: SFILIGOI, Igor (UCSD); MISHIN, Dima (UCSD); DYKSTRA, Dave (Fermi National Accelerator Lab. (US))

Presenter: SFILIGOI, Igor (UCSD)

Session Classification: User Stories

Contribution ID: 8

Type: **not specified**

The application of CVMFS at IHEP

Monday, June 3, 2019 2:50 PM (20 minutes)

IHEP has been using CVMFS since 2017 and cvmfs-stratum-one.ihep.ac server provides software repository duplicate services for cern.ch,opensciencegrid.org,egi.eu,ihep.ac.cn .Part of this report will introduce the status and next plan of CVMFS-stratum-one.ihep.ac.cn.

China's Large High Altitude Air Shower Observatory (LHAASO), a cosmic ray detection facility located in the high mountains of Sichuan province. The experimental data of LHAASO captured by the detector will be processed in the computer room which located in the high altitude and poor natural environment , and then transmitted to IHEP. The other part of this report will introduce the application of CVMFS in LHAASO experiments, and share some problems encountered in the use of CVMFS.

Expect to use for 10 minutes.

Primary authors: Mr HU, Qingbao (IHEP); Dr ZHANG, xiaomei (IHEP,Beijing); CHENG, Yaodong (IHEP)

Presenter: Mr HU, Qingbao (IHEP)

Session Classification: User Stories

Contribution ID: 9

Type: **not specified**

Using CVMFS with Spack at the FCC experiment

Monday, June 3, 2019 2:30 PM (20 minutes)

In preparation for the post-LHC era, the Future Circular Collider (FCC) Collaboration is undertaking design studies for multiple accelerator projects with emphasis on proton-proton and electron-positron high-energy frontier machines. From the beginning of the collaboration, the development of a software stack with common and interchangeable packages plays an important role in simulation, reconstruction or analysis studies. Similarly to the existing LHC experiments, these packages need to be built and deployed for different platforms and compilers. For these tasks FCC relies on Spack, a new package manager tool recently adopted by other experiments inside the High-Energy Physics (HEP) community. Despite its warm adoption, the integration of Spack with CVMFS for software distribution, while already possible, exposes some limitations. This talk provides an overview of these difficulties in the context of the FCC Software.

Primary author: CERVANTES VILLANUEVA, Javier (CERN)

Presenter: CERVANTES VILLANUEVA, Javier (CERN)

Session Classification: User Stories

Contribution ID: **10**

Type: **not specified**

A Full-Featured CernVM-FS Publish Deployment

Monday, June 3, 2019 9:30 AM (25 minutes)

Presenter: POPESCU, Radu (CERN)

Session Classification: News from the Developers

Contribution ID: 11

Type: **not specified**

CernVM-FS Status and Plans

Monday, June 3, 2019 9:55 AM (25 minutes)

Presenter: BLOMER, Jakob (CERN)

Session Classification: News from the Developers

Contribution ID: 12

Type: **not specified**

Shrinkwrap: Creating HPC Containers

Monday, June 3, 2019 11:30 AM (30 minutes)

Presenter: SHAFFER, Timothy

Session Classification: News from the Developers

Contribution ID: 13

Type: **not specified**

Evolution of the Gateway Services

Monday, June 3, 2019 11:00 AM (20 minutes)

Presenter: POPESCU, Radu (CERN)

Session Classification: News from the Developers

Contribution ID: 14

Type: **not specified**

CernVM Appliance: Status and Plans

Monday, June 3, 2019 11:20 AM (10 minutes)

Presenter: BLOMER, Jakob (CERN)

Session Classification: News from the Developers

Contribution ID: 15

Type: **not specified**

Containes for ATLAS User Analyses

Tuesday, June 4, 2019 2:20 PM (20 minutes)

Presenter: HEINRICH, Lukas Alexander (CERN)

Session Classification: User Stories

Contribution ID: **16**

Type: **not specified**

LHCb Nightly Build Use Case

Monday, June 3, 2019 2:10 PM (20 minutes)

Presenter: COUTURIER, Ben (CERN)

Session Classification: User Stories

Contribution ID: 17

Type: **not specified**

Evolution of CernVM-FS Infrastructure at CERN

Monday, June 3, 2019 4:40 PM (20 minutes)

The Storage group of the CERN IT department operates the CVMFS release managers for repositories hosted at CERN (also known as stratum zeros), the CERN replica servers (stratum ones), and the local squid caches. This talk describes the current architecture of the CVMFS service at CERN and reports on the introduction of S3 as the storage backend for stratum zeros, the upgrade of selected release managers to CERN CentOS 7, the deployment of the Gateway component, and the instantiation of dedicated squid proxies for ATLAS repositories.

Presenter: BOCCHI, Enrico (CERN)

Session Classification: User Stories

Contribution ID: 18

Type: **not specified**

Using CVMFS to distribute LCG Releases

Monday, June 3, 2019 5:00 PM (20 minutes)

This talk outlines the usage of the CernVM file system in the EP-SFT-SPI section at CERN. Its task is to distribute and continuously update a software stack (called LCG Releases) that is used by ATLAS, LHCb, SWAN and several more.

The stack contains HEP specific software as well as external packages. There are usually two major releases per year as well as nightly builds with the latest updates. For these nightly builds we have two main challenges to solve in the future: Increase publication speed and improve the control of the build environments.

We present our approach to solving these challenges by using recent CVMFS features such as parallel publications with a gateway machine. On the other hand we want to outline our future usage of Kubernetes orchestration to cope with the growing number of builds each night.

Presenter: HEINZ, Johannes Martin

Session Classification: User Stories

Contribution ID: 19

Type: **not specified**

Solving Problems in HPC with Singularity

Tuesday, June 4, 2019 11:00 AM (45 minutes)

The Singularity container runtime has become widely adopted as the de facto standard container platform for HPC workloads. At the beginning of 2018, Sylabs was founded to further HPC innovation by driving Singularity development. This talk will explore some of the ways in which the Singularity community and Sylabs are helping to solve problems in the HPC space, with a focus on the efforts of streamlining CVMFS usage with containers.

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 (Sylabs)

Session Classification: Technology Outlook

Contribution ID: 20

Type: **not specified**

A Tale of Two Clusters: CernVM-FS and CephFS in Context

Tuesday, June 4, 2019 3:00 PM (45 minutes)

This talk explores the underlying purposes of CernVM-FS and CephFS, providing intuition for the general goals of each project, indicating recent directions and highlighting use cases for which they are each well-suited.

About the Speaker

Jesse Williamson is a software engineer with experience in a wide variety of environments. He has contributed to Ceph, CernVM-FS, Riak, and the Boost C++ libraries. A long-time user of C++, he is an organizer of the Portland C++ User's Group, a member of several committees at CppCon, and has participated in WG21.

Presenter: WILLIAMSON, Jesse (Canonical)

Session Classification: Technology Outlook

Contribution ID: 21

Type: **not specified**

Scalable Applications

Tuesday, June 4, 2019 4:15 PM (45 minutes)

About the Speaker

Douglas Thain is an Professor and Associate Chair of Computer Science and Engineering at the University of Notre Dame. He received the Ph.D. in Computer Sciences from the University of Wisconsin and the B.S. in Physics from the University of Minnesota. His research is focused on the design of large scale computing systems for science and engineering, including scientific workflows, distributed filesystems, and high throughput computing.

Presenter: THAIN, Douglas (University of Notre Dame)

Session Classification: Technology Outlook

Contribution ID: 22

Type: **not specified**

Rootless containers with Podman and fuse-overlayfs

Tuesday, June 4, 2019 5:00 PM (45 minutes)

During my talk, I will show how it is possible to run OCI containers with Podman without requiring root privileges on the host.

In the second part of the talk, I'll focus on how we use fuse-overlayfs to replace the kernel overlay file system implementation, and the fuse-overlayfs built-in shifts-like capabilities.

About the Speaker

Giuseppe is a Principal Software Engineer at Red Hat. He is part of the Container Runtime Team, where he mainly works on Podman, Buildah, CRI-O.

Presenter: SCRIVANO, Giuseppe (Red Hat)

Session Classification: Technology Outlook

Contribution ID: 23

Type: **not specified**

Cvmfs on volunteer machines

Tuesday, June 4, 2019 2:00 PM (20 minutes)

Presenter: FIELD, Laurence (CERN)

Session Classification: User Stories

Contribution ID: 24

Type: **not specified**

Container Publishing Service

Monday, June 3, 2019 10:20 AM (10 minutes)

Presenter: MOSCIATTI, Simone (Politecnico di Milano (IT))

Session Classification: News from the Developers

Contribution ID: 25

Type: **not specified**

CVMFS at Compute Canada

Monday, June 3, 2019 3:10 PM (20 minutes)

This talk will present the activities and developments at Compute Canada related to CVMFS and publication and management of research software and data, including topics such as availability and resiliency considerations for globally-accessible repositories, collaboration and coordination challenges that can arise when different organizations use each other's repositories, and container image distribution in an experimental Kubernetes cluster at the University of Victoria.

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

Session Classification: User Stories

Contribution ID: 26

Type: **not specified**

Registration

Monday, June 3, 2019 8:30 AM (30 minutes)

Contribution ID: 27

Type: **not specified**

Announcements

Tuesday, June 4, 2019 8:45 AM (15 minutes)

Contribution ID: **28**

Type: **not specified**

Session Introduction

Monday, June 3, 2019 2:00 PM (10 minutes)

Contribution ID: 29

Type: **not specified**

CloudScheduler V2

Tuesday, June 4, 2019 2:40 PM (20 minutes)

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

Session Classification: User Stories