CernVM Virtual Workshop 2021

Report of Contributions

https://indico.cern.ch/e/cvm21

Contribution ID: **1** Type: **not specified**

Infrastructure Perspective

Contribution ID: **2** Type: **not specified**

CernVM-FS: Status and Roadmap

Monday, 1 February 2021 15:10 (35 minutes)

Presenter: BLOMER, Jakob (CERN)

Session Classification: CernVM Ecosystem: Roadmap and Latest Developments

Contribution ID: **3** Type: **not specified**

CernVM-FS and Containers

Monday, 1 February 2021 15:55 (20 minutes)

Presenter: MOSCIATTI, Simone (CERN)

Session Classification: CernVM Ecosystem: Roadmap and Latest Developments

Contribution ID: **4** Type: **not specified**

CernVM: Containerized Compatibility Layer

Monday, 1 February 2021 16:15 (15 minutes)

Presenter: BLOMER, Jakob (CERN)

Session Classification: CernVM Ecosystem: Roadmap and Latest Developments

Contribution ID: **5** Type: **not specified**

Workshop Introduction

Monday, 1 February 2021 15:00 (10 minutes)

Contribution ID: **6** Type: **not specified**

EESSI: A software stack for HPC

Tuesday, 2 February 2021 15:00 (30 minutes)

Presenter: DRÖGE, Bob (University of Groningen)

Session Classification: Technology Outlook

Contribution ID: **7** Type: **not specified**

Experience with CernVM-FS and Ceph/S3 at CERN

Tuesday, 2 February 2021 17:00 (20 minutes)

Presenter: BOCCHI, Enrico (CERN)

Session Classification: Use Cases, Requirements, and Feedback

Contribution ID: **8** Type: **not specified**

OSG Storage Repositories

Tuesday, 2 February 2021 17:20 (20 minutes)

Presenter: WEITZEL, Derek (University of Nebraska Lincoln (US))

Session Classification: Use Cases, Requirements, and Feedback

Contribution ID: **9** Type: **not specified**

CernVM-FS at NERSC

Tuesday, 2 February 2021 17:40 (20 minutes)

Presenter: BHIMJI, Wahid

Session Classification: Use Cases, Requirements, and Feedback

Contribution ID: **10** Type: **not specified**

Use Cases and Applications

Session Classification: Use Cases, Requirements, and Feedback

Contribution ID: **11** Type: **not specified**

TBA

Session Classification: Technology Outlook

Contribution ID: **12** Type: **not specified**

TBA

Session Classification: Technology Outlook

Contribution ID: **13** Type: **not specified**

The future of Volunteer Computing

Session Classification: Use Cases, Requirements, and Feedback

Contribution ID: **14** Type: **not specified**

CernVM-FS @ NIKHEF

Tuesday, 2 February 2021 15:50 (20 minutes)

Presenter: VAN DOK, Dennis

Session Classification: Use Cases, Requirements, and Feedback

Contribution ID: **15** Type: **not specified**

Performance of Distributed Publishing @ LHCb Nightlies

Tuesday, 2 February 2021 16:40 (20 minutes)

Presenter: BURR, Chris (CERN)

Session Classification: Use Cases, Requirements, and Feedback

Contribution ID: **16** Type: **not specified**

Unprivileged /cvmfs access with cvmfsexec

Monday, 1 February 2021 17:20 (20 minutes)

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

Session Classification: Use Cases, Requirements, and Feedback

Contribution ID: **17** Type: **not specified**

Powering k8s Clusters with the CernVM-FS Containerd Snapshotter

Monday, 1 February 2021 17:00 (20 minutes)

Presenter: LANGE, Clemens (CERN)

Session Classification: Use Cases, Requirements, and Feedback

Contribution ID: **18** Type: **not specified**

Increasing the execution speed of containerised analysis workflows using a Kubernetes image snapshotter in combination with CVMFS

The past years have shown a revolution in the way scientific workloads are being executed thanks to the wide adoption of software containers. These containers run largely isolated from the host system, ensuring that the development and execution environments are the same everywhere. This enables full reproducibility of the workloads and therefore also the associated scientific analysis performed. However, as the research software used becomes increasingly complex, the software images grow easily to sizes of multiple gigabytes. Downloading the full image onto every single compute node on which the containers are executed becomes unpractical. In our presentation, we describe a novel way of distributing software images on the kubernetes platform, with which the container can start before the entire image contents become available locally (so-called "lazy pulling"). Each file required for the execution is fetched individually and subsequently cached on-demand using CVMFS, enabling the execution of very large software images on potentially thousands of Kubernetes nodes with very little overhead. We present several performance benchmarks making use of typical analysis workloads performed by the CMS Experiment at the CERN Large Hadron Collider.

Primary authors: LANGE, Clemens (CERN); MOSCIATTI, Simone (CERN); BLOMER, Jakob (CERN)

Presenters: LANGE, Clemens (CERN); MOSCIATTI, Simone (CERN)

Contribution ID: **19** Type: **not specified**

GlideinWMS' use of cvmfsexec

Monday, 1 February 2021 17:40 (20 minutes)

Presenter: MAMBELLI, Marco (Fermilab (US))

Session Classification: Use Cases, Requirements, and Feedback

Contribution ID: **20** Type: **not specified**

CernVM at the ATLAS HLT

Tuesday, 2 February 2021 15:30 (20 minutes)

Presenter: Mr LEE, Chris (Stony Brook University (US))

Session Classification: Use Cases, Requirements, and Feedback

Contribution ID: **21** Type: **not specified**

CernVM-FS Ephemeral Writable Shell Demo

Monday, 1 February 2021 15:45 (10 minutes)

Presenter: VALENZUELA RAMIREZ, Andrea (Universitat Oberta de Catalunya (ES))

Session Classification: CernVM Ecosystem: Roadmap and Latest Developments