

CernVM Users Workshop

Monday 06 June 2016 - Wednesday 08 June 2016



Book of Abstracts

Contents

Using CernVM-FS to deploy Euclid processing S/W on Computing Centres	1
CernVM-FS Operations in the CERN IT Storage Group	1
CVMFS for Data Federations	2
Status and Plans for CernVM	2
Performance Engineering the CernVM-FS Backend	2
Volunteer Computing	3
Container Virtualization	3
Experiment Data Namespace Distribution	3
Evolution of the HEP Content Distribution Network	3
Using CernVM-FS to deploy Euclid processing S/W on Computing Centres	3
Virtualizing High Performance Computing Workloads	3
Introduction to Mesos and the Datacenter Operating System (DC/OS)	4
IBM Technical Computing	4
CernVM-FS for Docker image distribution in Cloud Foundry	4
Performance Overhead of ARM Virtualization	5
HEP Software on Supercomputers	5
Setting up and Scaling up CernVM-FS Repositories	5
CernVM in the Cloud	5
Welcome	5
New Features and Development Plans in CernVM-FS	5
CernVM-FS at ASGC	6
CVMFS Stratum-1 Replica Daemon	6
EGI CernVM-FS Infrastructure / RAL Operations	6

Feedback from ALICE	6
Feedback from LHCb	6
Feedback from CMS	6
Feedback from ATLAS	7
CERN@School	7
HEP Software @ NERSC	7
softdrive.nl - CVMFS for the Rest of Us	7
CCP4 - STFC Use Case	7

Feedback and Experience from Experiments / 0**Using CernVM-FS to deploy Euclid processing S/W on Computing Centres**

Authors: Mark Holliman¹; Maurice Poncet²; Quentin LE BOULC'H³

¹ ROE

² CNES

³ CNRS

Corresponding Authors: maurice.poncet@cnes.fr, msh@roe.ac.uk, quentin.leboulch@cc.in2p3.fr

CernVM-FS is being tested by the System Team of the Science Ground Segment (SGS) within the Euclid Project (<http://www.euclid-ec.org/>) as the primary execution environment and software distribution tool for processing mission data. In particular, CernVM-FS is viewed as a possible solution to the problem of providing a homogenous execution environment across heterogeneous computing resources at 9 different Science Data Centres (SDCs) located across Europe and the US. It also is seen as an ideal solution for enabling a DevOps methodology of software development, testing, and distribution that does not require regular intervention by systems administrators at individual SDCs.

CernVM-FS is currently being tested at scale through a technical SGS IT Challenge involving existing SDC compute clusters and the Euclid Continuous integration platform (CODEEN). This presentation will focus on this challenge and its main outcomes.

Summary:

CernVM-FS is being tested by the System Team of the Science Ground Segment (SGS) within the Euclid Project (<http://www.euclid-ec.org/>) as the primary execution environment and software distribution tool for processing mission data. In particular, CernVM-FS is viewed as a possible solution to the problem of providing a homogenous execution environment across heterogeneous computing resources at 9 different Science Data Centres (SDCs) located across Europe and the US. It also is seen as an ideal solution for enabling a DevOps methodology of software development, testing, and distribution that does not require regular intervention by systems administrators at individual SDCs.

CernVM-FS is currently being tested at scale through a technical SGS IT Challenge involving existing SDC compute clusters and the Euclid Continuous integration platform (CODEEN). This presentation will focus on this challenge and its main outcomes.

Feedback and Experience from Infrastructure Providers II / 1**CernVM-FS Operations in the CERN IT Storage Group**

Author: Dan van der Ster¹

¹ CERN

Corresponding Author: daniel.vanderster@cern.ch

This talk will present the status and plans of the CERN IT's CVMFS Operations Service, recently reassigned to our IT Storage Group. The first part of the talk will outline the architecture of the stratum zero, CERN stratum one, local squids; the notable feature of our service is that during 2015 is has been (almost) entirely moved to CERN's OpenStack infrastructure, making extensive use of Ceph and ZFS which has proven to be quite flexible while also offering a disaster recovery solution.

The second part of the talk will look to the future. CERN IT's recent change of strategy with respect to OpenAFS will put increased demands on CVMFS. We'll outline some of these new use-cases

which are forthcoming for CVMFS, then present alternatives for the future scale-up and out of the system.

Summary:

CernVM Related Activities II / 2

CVMFS for Data Federations

Author: Brian Paul Bockelman¹

Co-authors: Dave Dykstra²; Derek John Weitzel¹

¹ *University of Nebraska (US)*

² *Fermi National Accelerator Lab. (US)*

Corresponding Authors: brian.bockelman@cern.ch, dwd@fnal.gov, derek.weitzel@cern.ch

A data federation is a cooperating set of storage resources transparently accessible across a wide area network via a common namespace. These are often implemented through a redirector hierarchy - clients query a centralized endpoint for a given file; this redirector locates an available storage resource, then redirects the client to the remote resource.

Data federations are an increasingly used as a way to distribute large-volumes of physics data. For example, the Compact Muon Solenoid (CMS) experiment has approximately 20PB of analysis data available through it's "Any Data, Any Time, Anywhere" (AAA) federation.

However, the namespace of AAA is extremely limited - it is equivalent to just a HTTP GET. There are no directory listings, authoritative size or checksum information - despite the fact this information is known to CMS and available in the underlying storage systems and across several services; it is user-hostile for data discovery.

In this presentation, we will discuss a series of improvements made to the CVMFS core to marry a user-friendly, CVMFS-based POSIX namespace with data federation. We will demonstrate a set of CVMFS repositories of increasing complexity that utilize these new CVMFS features. These repositories serve as frontends for data federations for OSG, LIGO, and CMS.

Finally, we will discuss plans to grow this work - in terms of scale (data volume), efficiency, and features used in production.

Summary:

An effort to utilize CVMFS's scalable namespace features to provide a POSIX interface for data federations.

Status and Roadmap of the CernVM Ecosystem / 3

Status and Plans for CernVM

Corresponding Author: gerardo.ganis@cern.ch

Status and Roadmap of the CernVM Ecosystem / 5

Performance Engineering the CernVM-FS Backend

Corresponding Author: rene.meusel@cern.ch

CernVM Related Activities I / 6

Volunteer Computing

Corresponding Author: ivan.reid@cern.ch

CernVM Related Activities I / 7

Container Virtualization

CernVM Related Activities I / 8

Experiment Data Namespace Distribution

Corresponding Author: brian.bockelman@cern.ch

CernVM Related Activities I / 9

Evolution of the HEP Content Distribution Network

Corresponding Author: dwd@fnal.gov

Feedback and Experience from Experiments / 10

Using CernVM-FS to deploy Euclid processing S/W on Computing Centres

Corresponding Author: quentin.leboulch@cc.in2p3.fr

Technology Outlook II / 13

Virtualizing High Performance Computing Workloads

A short presentation on how and why HPC users are moving beyond public cloud for HPC workloads, including a discussion of achievable performance for both throughput and (briefly) MPI applications. We will also touch on compute accelerators, file system access, and containers, and provide some performance tuning tips for virtualizing HPC applications with high performance.

About the Speaker

With over 20 years of experience in High Performance and Technical Computing, Josh currently leads an effort at VMware to bring the value of virtualization to Research, Engineering, and Science environments. Previously, he was a Distinguished Engineer at Sun Microsystems with broad responsibilities for HPC direction and strategy. He joined Sun in 1996 from Thinking Machines Corporation, a pioneering company in the area of Massively Parallel Processors (MPPs), where he held a variety of technical positions. Josh has worked on developer tools for distributed parallel computing, including language and compiler design, scalable parallel debugger design and development, and MPI. He has also worked in the areas of 3D graphics, image processing, and real-time device control. Josh has an undergraduate degree in Engineering from Harvard College and a Masters in Computer Science from Harvard University. He has served as a member of the Board of Directors of OpenMP since 2002.

Summary:**Technology Outlook I / 14****Introduction to Mesos and the Datacenter Operating System (DC/OS)**

Corresponding Author: artem@mesosphere.io

Existing research has shown the benefits of running multi-level schedulers, either for single node parallel computation or multi-node distributed computation. Apache Mesos is a 2-level distributed scheduling system that has been used in organizations such as Twitter, PayPal, and Apple.

In this talk we'll describe the 2-level Mesos architecture in detail. We'll also discuss Mesos features that have been added over the years (or are being added today) driven by practical needs, from weights, to reservations, to quotas, to optimistic allocations, and deallocation. In the second part we will focus on some of the operational characteristics of Mesos, as well as present the open-source Datacenter Operating System that leverages Mesos while providing capabilities such as service discovery, load balancing, privacy and ease of installation.

About the Speaker

Artem Harutyunyan is a Director of Engineering at Mesosphere where he works on Apache Mesos project. Prior to joining Mesosphere he spent several years at Qualys where he led the design and development of a distributed system for storing and analyzing large volumes of data. Prior to Qualys, Artem worked at CERN, first as student in ALICE Offline, and later as a fellow in CernVM group.

Summary:**Technology Outlook III / 15****IBM Technical Computing**

Corresponding Author: oliver.oberst@de.ibm.com

Technology Outlook III / 16

CernVM-FS for Docker image distribution in Cloud Foundry

Cloud foundry is a PaaS that facilitates deployment and scale of web applications. It aims at providing a good user experience while ensuring that deployed application follow the best practices. Additionally, It provides a full set of features that enable debugging and monitoring of production systems. However, Docker is disrupting more traditional PaaS by making containers a prominent term in the cloud terminology. George, who works for Pivotal implementing Cloud Foundry's container runtime, will introduce Docker and Cloud Foundry, and discuss their differences. He will present the Docker image support in Cloud Foundry and eventually, how CernVM-FS could fit in the area of Docker image distribution.

Summary:

Technology Outlook II / 17

Performance Overhead of ARM Virtualization

Summary:

CernVM Related Activities I / 18

HEP Software on Supercomputers

Corresponding Author: rodney.walker@physik.uni-muenchen.de

Hands-on session / 19

Setting up and Scaling up CernVM-FS Repositories

Hands-on session / 20

CernVM in the Cloud

Status and Roadmap of the CernVM Ecosystem / 22

Welcome

Corresponding Authors: catalin.condurache@cern.ch, jakob.blomer@cern.ch, gerardo.ganis@cern.ch

Status and Roadmap of the CernVM Ecosystem / 23

New Features and Development Plans in CernVM-FS

Corresponding Author: jakob.blomer@cern.ch

Feedback and Experience from Infrastructure Providers I / 25

CernVM-FS at ASGC

Author: Wei-Jen Chang¹

Co-author: Felix.hung-te Lee ¹

¹ *Academia Sinica (TW)*

Corresponding Authors: felix@twgrid.org, weijen.chang@twgrid.org

CernVM Related Activities II / 26

CVMFS Stratum-1 Replica Daemon

Corresponding Author: jose.caballero@cern.ch

Feedback and Experience from Infrastructure Providers II / 27

EGI CernVM-FS Infrastructure / RAL Operations

Author: Catalin Condurache¹

¹ *STFC - Rutherford Appleton Lab. (GB)*

Corresponding Authors: catalin.condurache@cern.ch, john.kelly@stfc.ac.uk

Feedback and Experience from Experiments / 28

Feedback from ALICE

Corresponding Author: dario.berzano@cern.ch

Feedback and Experience from Experiments / 29

Feedback from LHCb

Corresponding Author: ben.couturier@cern.ch

Feedback and Experience from Experiments / 30

Feedback from CMS

Corresponding Author: andrew.david.lahiff@cern.ch

Feedback and Experience from Experiments / 31

Feedback from ATLAS

Corresponding Author: alastair.dewhurst@cern.ch

Summary:

CernVM Related Activities II / 32

CERN@School

Corresponding Author: tom.whyntie@cern.ch

CernVM Related Activities II / 33

HEP Software @ NERSC

Corresponding Author: markus.fasel@cern.ch

Summary:

Feedback and Experience from Infrastructure Providers II / 34

softdrive.nl - CVMFS for the Rest of Us

Corresponding Author: dennisvd@nikhef.nl

NIKHEF use case

Summary:

CernVM Related Activities II / 35

CCP4 - STFC Use Case

Summary: