Contribution ID: **172** Type: **Poster**

New Directions in the CernVM File System

Thursday 13 October 2016 16:30 (15 minutes)

The CernVM File System today is commonly used to host and distribute application software stacks. In addition to this core task, recent developments expand the scope of the file system into two new areas. Firstly, CernVM-FS emerges as a good match for container engines to distribute the container image contents. Compared to native container image distribution (e.g. through the "Docker registry"), CernVM-FS massively reduces the network traffic for image distribution. This has been shown, for instance, by a prototype integration of CernVM-FS into Mesos developed by Mesosphere, Inc. We present possible paths for a smooth integration in Docker and necessary changes to the CernVM-FS server to support the typical container image work flow and lifecycle.

Secondly, CernVM-FS recently raised interest as an option for the distribution of experiment data file catalogs. While very powerful tools are in use for accessing data files in a distributed and scalable manner, finding the file names is typically done by a central, experiment specific SQL database. A name space on CernVM-FS can particularly benefit from an existing, scalable infrastructure, from the POSIX interface and the end-to-end content verification. For this use case, we outline necessary modifications to the CernVM-FS server in order to provide a generic, distributed namespace that supports billions of files and thousands of writes per second.

Tertiary Keyword (Optional)

Databases

Secondary Keyword (Optional)

Virtualization

Primary Keyword (Mandatory)

Storage systems

Primary author: BLOMER, Jakob (CERN)

Co-authors: GANIS, Gerardo (CERN); BUNCIC, Predrag (CERN); MEUSEL, Rene (CERN)

Presenter: BLOMER, Jakob (CERN)

Session Classification: Posters B / Break

Track Classification: Track 4: Data Handling