

Status and Roadmap of CernVM

Primary authors: J.Blomer, D.Berzano, P.Buncic, I.Charalampidis, G.Ganis, R.Meusel

Cloud resources nowadays contribute an essential share of resources for computing in high-energy physics. Such resources can be either provided by private or public IaaS clouds (e.g. OpenStack, Amazon EC2, Google Compute Engine) or by volunteers' computers (e.g. LHC@Home 2.0). In any case, experiments need to prepare a virtual machine image that provides the execution environment for the physics application at hand.

The CernVM virtual machine since version 3 is a minimal and versatile virtual machine image capable of booting different operating systems. The virtual machine image is less than 20 megabyte in size. The actual operating system is delivered on demand by the CernVM File System.

CernVM 3 has matured from a prototype to a production environment. It is used, for instance, to run LHC applications in the cloud, to tune event generators using a network of volunteer computers, and as a container for the historic Scientific Linux 5 and Scientific Linux 4 based software environments in the course of long-term data preservation efforts of the ALICE, CMS, and ALEPH experiments. We present experience and lessons learned from the use of CernVM at scale. We also provide an outlook on the upcoming developments. These developments include adding support for Scientific Linux 7, the use of container virtualization, such as provided by Docker, and the streamlining of virtual machine contextualization towards the cloud-init industry standard.