

Preparing for HL-LHC: Increasing the LHCb software publication rate to CVMFS by an order of magnitude

Wednesday, 19 May 2021 11:42 (13 minutes)

In the HEP community, software plays a central role in the operation of experiments' facilities and for reconstruction jobs, with CVMFS being the service enabling the distribution of software at scale. In view of High Luminosity LHC, CVMFS developers investigated how to improve the publication workflow to support the most demanding use cases. This paper reports about recent CVMFS developments and infrastructural updates that enable faster publication into existing repositories. A new CVMFS component, the CVMFS Gateway, allows for concurrent transactions and the use of multiple publishers, increasing the overall publication rate on a single repository. Also, the repository data has been migrated to Ceph-based S3 object storage, which brings a relevant performance enhancement over the previously-used Cinder volumes. We demonstrate how recent improvements allow for faster publication of software releases in CVMFS repositories by focusing on the LHCb nightly builds use case, which is currently by far the most demanding one for the CVMFS infrastructure at CERN. The publication of nightly builds is characterized by a high churn rate, needs regular garbage collection, and requires the ability to ingest a huge amount of software files over a limited period of time.

Primary authors: BOCCHI, Enrico (CERN); BLOMER, Jakob (CERN); COUTURIER, Ben (CERN); BURR, Chris (CERN); VAN DER STER, Dan (CERN)

Presenter: BOCCHI, Enrico (CERN)

Session Classification: Storage

Track Classification: Distributed Computing, Data Management and Facilities