



Contribution ID: 470

Type: Oral

CERN Disk Storage Services: report from last data taking, evolution and future outlook towards Exabyte-scale storage

Thursday, November 7, 2019 12:15 PM (15 minutes)

The CERN IT Storage group operates multiple distributed storage systems to support all CERN data storage requirements: the physics data generated by LHC and non-LHC experiments; object and file storage for infrastructure services; block storage for the CERN cloud system; filesystems for general use and specialized HPC clusters; content distribution filesystem for software distribution and condition databases; and sync&share cloud storage for end-user files. The total integrated capacity of these systems exceeds 0.6 Exabyte.

Large-scale experiment data taking has been supported by EOS and CASTOR for the last 10+ years. Particular highlights for 2018 include the special Heavy-Ion run which was the last part of the LHC Run2 Programme: the IT storage systems sustained over 10GB/s to flawlessly collect and archive more than 13 PB of data in a single month. While the tape archival continues to be handled by CASTOR, the effort to migrate the current experiment workflows to the new CERN Tape Archive system (CTA) is underway.

Ceph infrastructure has operated for more than 5 years to provide block storage to CERN IT private OpenStack cloud, a shared filesystem (CephFS) to HPC clusters and NFS storage to replace commercial Filers. S3 service was introduced in 2018, following increased user requirements for S3-compatible object storage from physics experiments and IT use-cases.

Since its introduction in 2013, CERNBox has become a ubiquitous cloud storage interface for all CERN user groups: physicists, engineers and administration. CERNBox provides easy access to multi-petabyte data stores from a multitude of mobile and desktop devices and all mainstream, modern operating systems (Linux, Windows, macOS, Android, iOS). CERNBox provides synchronized storage for end-user's devices as well as easy sharing for individual users and e-groups. CERNBox has also become a storage platform to host online applications to process the data such as SWAN (Service for Web-based Analysis) as well as file editors such as Collabora Online, Only Office, Draw.IO and more. An increasing number of online applications in the Windows infrastructure uses CIFS/SMB access to CERNBox files.

CVMFS provides software repositories for all experiments across the WLCG infrastructure and has recently been optimized to efficiently handle nightly-builds. While AFS continues to provide general-purpose filesystem for internal CERN users, especially as \$HOME login area on central computing infrastructure, the migration of project and web spaces has significantly advanced.

In this paper, we report on the experiences from the last year of LHC RUN2 data taking and evolution of our services in the past year. We will highlight upcoming changes and future improvements and challenges.

Consider for promotion

Yes

Primary authors: MOSCICKI, Jakub (CERN); MASCETTI, Luca (CERN); VAN DER STER, Dan (CERN); IVEN,

Jan (CERN); COLLET, Julien (CERN); PELLETIER, Remy (CERN); MOURATIDIS, Theofilos (National and Kapodistrian University of Athens (GR)); MUSSET, Paul (CERN); CASTRO, Diogo (CERN); CALADO VICENTE, Joao (Universidade de Lisboa (PT)); VALVERDE CAMESELLE, Roberto (Universidad de Oviedo (ES)); ARSUAGA RIOS, Maria (CERN); BOCCHI, Enrico (CERN); LAMANNA, Massimo (CERN); CHAN KWOK CHEONG, Belinda (CERN); GONZALEZ LABRADOR, Hugo (CERN); CONTESCU, Cristian (CERN); LO PRESTI, Giuseppe (CERN)

Presenter: MASCETTI, Luca (CERN)

Session Classification: Track 4 – Data Organisation, Management and Access

Track Classification: Track 4 – Data Organisation, Management and Access