

Storage services at CERN

Wednesday 27 March 2019 11:10 (25 minutes)

The Storage group of the CERN IT department is responsible for the development and the operation of petabyte-scale services needed to accommodate the diverse requirements for storing physics data generated by LHC and non-LHC experiments as well as supporting users of the laboratory in their day-by-day activities.

This contribution presents the current operational status of the main storage services at CERN, summarizes our experience in operating largely distributed systems and highlights the ongoing efforts for the evolution of the storage infrastructure.

It reports about EOS, the high-performance distributed filesystem developed at CERN designed to store all the physics data and to operate at the high rates demanded by experiments data taking. EOS is also used as the storage backend for CERNBox, the cloud storage synchronization and sharing service for users' personal files. CERNBox provides uniform access to storage from all modern devices and represents the data hub for integration with various applications ranging from office suites (Microsoft Office 365, OnlyOffice, Draw.io) to specialized tools for data analysis (SWAN).

Besides storage for physics data and personal files, the Storage group runs multiple large Ceph clusters to provide the storage backbones for the OpenStack infrastructure and the HPC facility, and to offer an S3 service and a CephFS/Manila shares for other internal IT services. Also, the Storage group operates the release managers, replica servers and caches of CVMFS (a fundamental WLCG service used for software distribution) in collaboration with the SoFTware Development for Experiments (CERN EP-SFT) department.

Primary author: BOCCHI, Enrico (CERN)

Presenter: BOCCHI, Enrico (CERN)

Session Classification: Storage & Filesystems

Track Classification: Storage & Filesystems