

Next generation of large-scale storage services at CERN

Wednesday, May 16, 2018 9:40 AM (20 minutes)

CERN IT Storage (IT/ST) group leads the development and operation of large-scale services based on EOS for the full spectrum of use-cases at CERN and in the HEP community. IT/ST group also provides storage for other internal services, such as Open Stack, using a solution based on Ceph. In this talk we present current operational status, ongoing development work and future architecture outlook for next generation storage services for the users based on EOS – a technology developed and integrated at CERN.

EOS is the home for all physics data-stores for LHC and non-LHC experiments (at present 250PB storage capacity). It is designed to operate at high data rates for experiment data-taking while running concurrent complex production work-loads. EOS also provides a flexible distributed storage back-end and architecture with plugins for tape archival (CTA - evolution and replacement for CASTOR), synchronization&sharing services (CERNBox) and general-purpose filesystem access for home directories (FUSE for Linux and SMB Gateways for Windows and Mac).

CERNBox is the cloud storage front-end for desktop, mobile and web access focused on personal user files, general-purpose project spaces and smaller physics datasets (at present 12K user accounts and 500M files). CERNBox provides simple and uniform access to storage on all modern devices and operating systems. CERN-Box is also hub for integration with other services: Collaborative editing – MS Office365 and alternatives: Collabora and OnlyOffice; Web-based analysis – SWAN Jupyter Notebooks with access to computational resources via Spark and Batch; and software distribution via CVMFS.

This storage service ecosystem is designed to provide “total data access”: from end-user devices to geo-aware data lakes for WLCG and beyond. It also provides a foundation for strategic partnerships (AARNet, JRC, ...), new communities such as CS3 (Cloud Storage and Synchronization Services) and new application projects such as Up2University (cloud storage ecosystem for education). CERN Storage technology has been show-cased to work with commercial cloud providers such as Amazon, T-Systems (Helix Nebula) or COMTRADE (Openlab) and there is an increasing number of external sites testing the CERN storage service stack in their local computing centers.

This strategy proves very successful with the users and as a result storage services at CERN see exponential growth: CERNBox alone has grown by 450% in 2017. Growing overall demand drive the evolution of the service design and implementation of the full ecosystem: EOS core storage as well as CERNBox and SWAN. Recent EOS improvements include new distributed namespace to provide scaling and high-availability; new robust FUSE module providing client-side caching, lower latency and more IOPs; new workflow engine and many more. CERNBox is moving to micro-service oriented architecture and SWAN is tested with Kubernetes container orchestration.

New developments come together with a constant effort to streamline QA, testing and documentation as well as reduce manual configuration and operational effort for managing large-scale storage services.

Desired length

20

Primary author: MOSCICKI, Jakub (CERN)

Co-authors: BOCCHI, Enrico (CERN); CONTESCU, Cristian (CERN); GONZALEZ LABRADOR, Hugo (CERN); LAMANNA, Massimo (CERN); MASCETTI, Luca (CERN); PETERS, Andreas Joachim (CERN); LO PRESTI, Giuseppe (CERN); ROUSSEAU, Herve (CERN); VALVERDE CAMESELLE, Roberto (Universidad de Oviedo (ES))

Presenter: MOSCICKI, Jakub (CERN)

Session Classification: Storage and file systems

Track Classification: Storage & Filesystems