



Contribution ID: 491

Type: Poster

Migration of user and project spaces with EOS\CERNBox: experience on scaling and large-scale operations

Thursday 7 November 2019 16:15 (15 minutes)

EOS is the key component of the CERN Storage strategy and is behind the success of CERNBox, the CERN cloud synchronisation service which allows syncing and sharing files on all major mobile and desktop platforms aiming to provide offline availability to any data stored in the infrastructure.

CERNBox faced an enormous success within the CERN users' community thanks to its always increasing popularity and to its integration with a multitude of other CERN services such as Batch, SWAN, Microsoft Office to mention a few.

This success directly translated into an exponential growth in the last couple of years in terms of files and data stored which were leading to an explosion of the in-memory namespace technology used for the main catalogue.

The original deployment from 2014 has proven effective and reliable, however the infrastructure was very quickly approaching the limits of its scalability and its capability of coping with the exponential growth. The sense of urgency was real: if the memory limit of the namespace server was reached, the service would simply stop. To counter this threat the storage team re-designed the EOS backend service architecture for end-user and project spaces.

In this paper we present our new EOS deployment which entered in production during summer 2018 and we show how we successfully sharded the system into smaller failure domains and transform the original scale-up namespace deployment into a less expensive and more robust scale-out solution. We report about the experience gained so far in the simplification of the cluster upgrade procedures, now completely transparent, and the effort behind a transparent migration of 16k users from the old architecture into the new EOS architecture.

Consider for promotion

Yes

Author: MASCETTI, Luca (CERN)

Co-authors: LAMANNA, Massimo (CERN); KARAVAKIS, Edward (CERN); MOSCICKI, Jakub (CERN); GONZALEZ LABRADOR, Hugo (CERN); PETERS, Andreas Joachim (CERN)

Presenter: MASCETTI, Luca (CERN)

Session Classification: Posters

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