



Contribution ID: 468

Type: Oral

Evolution of the S3 service at CERN as a storage backend for infrastructure services and software repositories

Thursday 7 November 2019 14:00 (15 minutes)

The S3 service at CERN (S3.CERN.CH) is a horizontally scalable object storage system built with a flexible number of virtual RADOS Gateways on top of a conventional Ceph cluster. A Traefik load balancing frontend (operated via Nomad and Consul) redirects HTTP traffic to the RGW backends, and LogStash publishes to Elasticsearch for monitoring the user traffic. User and quota management is delegated to OpenStack; a novel synchronization mechanism was developed to maintain consistency between Keystone and Ceph.

The RADOS backend went through some notable operational exercises, including massive data rebalancing, a FileStore to BlueStore migration, and the addition of SSDs for bucket indexes; these operations gave us increased capacity and up to 50x speedup for some performance metrics.

The usage of S3 at CERN has evolved in the recent years. The earliest adopters of S3 were the ATLAS Event Service and CMS BOINC use-cases; the latter takes advantage of S3 pre-signed URLs to offer a secure storage to untrusted remote user PCs. More recently our IT applications demanded S3 object storage: Gitlab and Nexus artifacts, Elasticsearch backups, Kubernetes and others. And our WLCG-specific software is also increasingly taking advantage of the service: CVMFS stratum zero data is a natural fit, and S3 combined with Restic offers a compelling \$HOME backup service. Lastly, with the decommissioning of the CERN data centre in Hungary, disaster recovery requirements have motivated a second S3 region on the CERN Preveessin site, with async RGW replication for important buckets.

Consider for promotion

No

Authors: CASTRO LEON, Jose (CERN); COLLET, Julien (CERN); MOURATIDIS, Theofilos (National and Kapodistrian University of Athens (GR)); ROUSSEAU, Herve (CERN); VALVERDE CAMESELLE, Roberto (Universidad de Oviedo (ES)); VAN DER STER, Dan (CERN); BOCCHI, Enrico (CERN)

Presenter: BOCCHI, Enrico (CERN)

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

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