

Contribution ID: 74

Type: Presentazione orale

INFN Cloud Object Storage Service: Gateway to a Multisite Infrastructure

Tuesday, 21 May 2024 15:30 (20 minutes)

INFN Cloud has been offering its users an S3-based Object Storage service for data archiving that is accessible via a web user interface or programmatically. Key features of the S3 service are, among others, the authentication via the Indigo DataCloud IAM service, the use of Open Policy Agent for fine grained authorization, the full integration with other INFN Cloud services, and data replication over two data centers. Such service has been provided via the Minio Gateway on top of a distributed OpenStack Swift cluster.

The end of support of Minio Gateway forced us to find an alternative solution. As new implementation we considered the use of a multisite CEPH Rados Gateway (RGW) infrastructure. The challenge of replacing the MinIO-based architecture with the Ceph RGW sytem was eased by the already existing CEPH infrastructure that supports the two sites of the INFN Cloud Backbone.

CEPH RGW was configured to allows access via OIDC Identity Provider and cross-account operations by offering Security Token Service (STS). In addition, the integration with OPA allows for a granular an authorization policy reinforcement: while CEPH administrator creates a role to define both location and type of storage resources available for the identity provider's users, the decision making for buckets and objects management is offloaded to OPA. In addition, policy decoupling allows dynamic modification of accessibility rules and supports scalability that can accommodate increasing resource demand on geographically distributed infrastructures. Thanks to its multi-site functionality, CEPH RGW can be configured to enable object storage synchronization among different sites, ensuring geographically distributed data replication.

The S3 compatible storage service that is available with the new solution is also accessible via a wide range of client applications such as SDK (Boto3), CLI that manages data using subcommands (s3cmd) or mounting buckets to filesystem using FUSE(sts-wire), and GUI with the in-house designed Web Application that is based on ReactJS and the official AWS SDK for JavaScript.

Primary authors: ALKHANSA, Ahmad (Istituto Nazionale di Fisica Nucleare); COSTANTINI, Alessandro (Istituto Nazionale di Fisica Nucleare); MARTELLI, Barbara (Istituto Nazionale di Fisica Nucleare); Dr MICH-ELOTTO, Diego (CNAF); MALATESTA, Giada (Istituto Nazionale di Fisica Nucleare); GASPARETTO, Jacopo (Istituto Nazionale di Fisica Nucleare); SGARAVATTO, Massimo (Istituto Nazionale di Fisica Nucleare); STALIO, Stefano (Istituto Nazionale di Fisica Nucleare)

Presenter: ALKHANSA, Ahmad (Istituto Nazionale di Fisica Nucleare)

Session Classification: Sessione "Infrastrutture ICT e calcolo distribuito"

Track Classification: Infrastrutture ICT e Calcolo Distribuito