



Contribution ID: 40

Type: **Presentation**

Improving swift as Owncloud backend to scale

Wednesday, January 31, 2018 2:00 PM (20 minutes)

The National Education and Research Network (RNP) is an organization that plans, designs, implements and operates the national network infrastructure under contract with the Ministry of Science, Technology, Innovation and Communications (MCTIC). A current government program includes five ministries - MCTI, Education (MEC), Culture (MinC), Health (MS) and Defense (MD), and annually define the objectives of the contract and its plan.

The increasing production of scientific data (eg, environmental monitoring, biodiversity databases, a variety of simulation and visualization systems such as climate forecasting, high physical energy data collection, astronomy and cosmology), cultural datasets and others, brings the need for a scalable, sustainable and high availability IT infrastructure to support these demands, and these facilities must be located in a distributed manner and in locations offering telecommunications, energy and safety services, as well as appropriate physical space and infrastructure.

In addition to the question of where these data are stored and processed, there is a need to create different institutions in research programs, such as the Large-scale Biosphere-Atmosphere Experiment in the Amazon (LBA), which is composed of 280 institutional and international, with about 1400 Brazilian scientists, 900 researchers from Amazonian countries and 8 European nations and from American institutions, aiming to study and understand the climate and environmental changes in the Amazon. In this type of communities sharing information securely and in compliance with legislation, whether from data collected from sensors, or in the production of articles or articles, is vital, and a cloud file synchronization and sharing solution meets this demand.

edudrive@RNP is cloud file synchronization and sharing offered by RNP to your community, it allows users to sync your data between your desktops, notebooks, smartphones and tablets and share this data with others users of the service or not, supporting researchers, teachers and students in research projects and during his academic studies.

The service is developed by RNP, in partnership with Anolis IT and is based on Owncloud software, which acts as the frontend of the service offering a web portal, desktop and mobile clients that users uses to synchronize and share their files. In addition, Openstack Swift is also used as a multi-tenant object storage backend, which provides high scalability, cost savings and significant resiliency, as a Software Defined Storage (SDS). Additionally, one of the key service requirements it's the integration with Shibboleth-based SAML federation for user authentication and authorization.

During the development of the service was identified a problem in the way that Owncloud use Openstack Swift as a storage backend, which caused a great slowness in file upload, download and delete operations when the number of files in the service grows.

This problem occurs in the connection between Openstack Swift and Owncloud, because when configuring Openstack Swift as the Owncloud storage backend it is mapped a single tenant and a single container from Openstack Swift where all the data are stored, it do not like a problem with a few number of files, but when the number of files grows, the search and replication activities of this data become slow, this occurs because the growth of sqlite metadata databases, which are synchronized between the storage nodes in each upload and delete operations, and this performance issue increase when you have a geographically separated infrastructure, negatively impacting the use of the service.

In addition, another important issue identified is the security of the information stored in the service, because all data of all users are stored in the same tenant and container, and this do not guarantee a logical segregation of this data, and in case of a leak of the credentials of the tenant and container used in the backend all the data stored on this can be exposed.

Based on this performance issue, it was necessary a development effort, in order to balance that storage backend load and use all features that is offered by this type of storage backend properly. To accomplish this, for each institution (university, research center, etc) was mapped as a tenant inside the storage backend, and each user was mapped as a single container inside your institution (tenant).

With this new mapping, the performance problems were solved, in addition, the data of the institutions and users of the service were segregated in a more appropriate way, bringing more security to the service. All changes were made to OwnCloud's OpenStack Object Storage plug-in, more specifically in the swift class, with changes to the default flow of user access to the service, and the execution of operations such as uploading or downloading of files.

All this new flow, which originally was controlled by OwnCloud, is now controlled by the Federated Access Control System (FACS), managing all the institutions which are authorized to access the service in an identity federation and creating the tenant of each institution based on its identifier in the identity federation and the user container during his first access, based on its unique identifier in the identity federation (EPPN - EduPersonPrincipalName). After that, all user data is saved in its own container within the tenant of your institution.

Authors: MAKINO, Ricardo (RNP); DINIZ, Thomas (Anolis); CASTRO, André (Anolis); MALUF, Guilherme (Anolis); AZEVEDO, Rodrigo (RNP)

Presenters: MAKINO, Ricardo (RNP); AZEVEDO, Rodrigo (RNP)

Session Classification: Scalable Storage Backends for Cloud and HPC: Integration