



Contribution ID: 341

Type: poster presentation

Local storage federation through XRootD architecture for interactive distributed analysis

A cloud-based Virtual Analysis Facility (VAF) for the ALICE experiment at the LHC has been developed in Bari. Similar facilities are currently running in other Italian sites with the aim to create a federation of interoperating farms able to provide their computing resources for interactive distributed analysis. The facility consists in a PROOF cluster of virtual machines dynamically deployed by the Openstack cloud infrastructure built in Bari for the PRISMA project.

The use of cloud technology, along with elastic provisioning of computing resources as an alternative to the grid for running data intensive analyses, is the main challenge of the facility. One of the crucial aspects of the user-driven analysis execution is the data access. A local storage facility has the disadvantage that the stored data can be accessed by PROOF only locally, i.e. from within the single VAF. To overcome such a limitation a federated infrastructure, which provides full access to all the data belonging to the federation independently from the site where they are stored, has been set up. The federation architecture exploits both cloud computing (block storage, posix storage and object storage) and XRootD technologies, in order to provide a dynamic, easy-to-use and well performing solution for data handling. It allows the users to store the files and efficiently retrieve the data, since it implements a dynamic distributed cache among many datacenters in Italy connected to one another through the high-bandwidth national network.

In this contribution we will show the technical solution, based on CEPH, GlusterFS or Swift together with XRootD, which has been implemented to build this distributed storage infrastructure. This solution will be able either to store data for long time periods or just cache data locally, so that they can be processed by the PROOF facilities at any given site. Details on the architecture implementation and the achieved performance in the first prototype tests will be also discussed.

The present work is supported by the Istituto Nazionale di Fisica Nucleare (INFN) of Italy and is partially funded under contract 20108T4XTM of Programmi di Ricerca Scientifica di Rilevante Interesse Nazionale (STOA-LHC PRIN, Italy).

Authors: Dr FRANCO, Antonio (Universita e INFN (IT)); COLELLA, Domenico (Universita e INFN (IT)); ELIA, Domenico (INFN Bari); Dr DONVITO, Giacinto (INFN-Bari); MINIELLO, Giorgia (Universita e INFN (IT)); MAGGI, Giorgio (Universita e INFN (IT))

Presenters: ELIA, Domenico (INFN Bari); MINIELLO, Giorgia (Universita e INFN (IT))

Track Classification: Track3: Data store and access