



Contribution ID: 229

Type: Oral presentation to parallel session

Integration of S3-based cloud storage in BES III computing environment

Monday, 14 October 2013 16:31 (20 minutes)

Object storage systems based on Amazon's Simple Storage Service (S3) have substantially developed in the last few years. The scalability, durability and elasticity characteristics of those systems make them well suited for a range of use cases where data is written, seldom updated and frequently read. Storage of images, static web sites and backup systems are some of the use cases where S3 systems have proven effective. Experimental data for high-energy physics research can also benefit from storage systems optimized for write-once read-many operational models.

The BES III experiment studies physics in the tau-charm energy region from 2GeV to 4.6 GeV, at the Institute of High Energy Physics (IHEP) in Beijing, China. Since spring 2009, BES III has been recording and accumulating a significant amount of experimental data, in the order of 1 PB per year. Organized around the central data repository operated by IHEP's computing center, the experiment's computing environment is composed of sites located in several countries.

In this contribution we present an ongoing work, which aims to evaluate the suitability of S3-based cloud storage as a supplement to the Lustre file system for storing experimental data for BES III. In particular, we discuss our findings regarding the integration of S3-based storage in the software stack of the experiment. We report on our development work that improves the support of CERN's ROOT data analysis framework and allows efficient remote access to data through the S3 protocol. We also discuss our results providing the experiment with efficient command line tools for interacting with S3-based data repositories from interactive sessions and grid jobs. The FUSE-based file system interface for a S3 storage backend that we developed is also presented and our efforts for providing tools for easily navigating the experiment's data repository and making it seamlessly accessible in particular from the researcher's personal computer.

This work is being validated through real use cases of production BES III jobs by using two different storage backends: a hardware-based solution around Huawei UDS appliance and a software-based solution around OpenStack Swift. We compare the performance of those systems with the Lustre file system for local and grid jobs and also for transferring data from and to remote sites participating in the BES collaboration.

Summary

Primary authors: Mr HERNANDEZ, Fabio (IN2P3/CNRS and Institute of High Energy Physics,CAS); Dr LU, Wang (Institute of High Energy Physics,CAS); Dr DENG, Ziyang (Institute of High Energy Physics,CAS)

Presenter: Dr LU, Wang (Institute of High Energy Physics,CAS)

Session Classification: Data Stores, Data Bases, and Storage Systems

Track Classification: Data Stores, Data Bases, and Storage Systems