



Contribution ID: 332

Type: **Oral presentation to parallel session**

Testing of several distributed file-system (HadoopFS, CEPH and GlusterFS) for supporting the HEP experiments analysis.

Tuesday 15 October 2013 14:16 (20 minutes)

In this work we will show the testing activity carried on several distributed file-system in order to check the capability of supporting the HEP data analysis

In particular, we focused our attention and our test on HadoopFS, CEPH, and GlusterFS.

All are Open Source software.

HadoopFS is an Apache foundation software and is part of a more general framework, that contains: task scheduler, a NOSQL database, a data warehouse system, etc. It is used by several big company and institution (Facebook, Yahoo, LinkedIn, etc).

CEPH is a quite young file-system that has very good design in order to guarantee great scalability, performance and very good high availability features. It is also the unique file-system that is able to provide three interface to storage: posix file-system, REST object storage and device storage. The support for CEPH was introduced as a native in the last release of the kernel.

GlusterFS is recently acquired by RedHat and this will ensure the long term support of the code. It has indeed a large user base both in HPC computing farms, and in several Cloud computing facilities. Indeed it support access to storage both in terms of posix file-system and via a REST gateway for object storage support.

All those file-system are capable of supporting high availability of the data and metadata in order to build a distributed file-system that could provide resilience to the hardware and/or software failure of one or more data server in the cluster.

We will describe each file-system in details providing the technical specification and reporting about the testing of the most interesting functionalities of each of the softwares.

We will focus our attention on the capabilities of recover from failures of both hardware and software and on how each software is able to provide those capabilities and describing the test carried on to prove them.

We will show also performance test carried on using data analysis application that reads data in standard ROOT format in order to better compare those software from a point of view of the HEP community.

In this work we will also present the results of tests that will highlight the scalability of each of those file systems.

We will show also the development that we have done to provide more powerful monitoring capabilities for HadoopFS. We have developed a web based monitoring system that is capable to show in details the information about the status of the data nodes or the status and the historical information about the location of each block.

We will also provide detailed information on automatic procedures and script developed in order to easily manage a big datacenter composed of hundreds of data node installed with HadoopFS

In this work we will also focus on the test executed in order to exploit the GlusterFS and CEPH file-system within an IaaS Cloud Infrastructure based on OpenStack thanks to the interfaces available in those storage technologies

Primary authors: Mr DIACONO, Domenico (INFN-Bari); DONVITO, Giacinto (Universita e INFN (IT)); Mr MARZULLI, Giovanni (GARR INFN)

Presenter: DONVITO, Giacinto (Universita e INFN (IT))

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

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