

Published on *CERN openlab* (<http://test-static-05.web.cern.ch>)

[Home](#) > The Implementation of OpenStack Cinder and Integration with NetApp and Ceph

The Implementation of OpenStack Cinder and Integration with NetApp and Ceph ^[1]

Date published:

Sunday, 1 September, 2013

Document type:

Summer student report

Author(s):

G. McGilvary

Project Specification: CERN is establishing a large scale private cloud based on OpenStack as part of the expansion of the computing infrastructure for the Large Hadron Collider (LHC). Depending on the application running on the cloud, some virtual machines require large disk capacities or high reliability/performance volumes. This project involves the configuration, deployment and testing of OpenStack Cinder as well as the integration with other block storage alternatives such as NetApp and Ceph. A performance analysis and comparison between these storage mechanisms will be undertaken to determine the most suitable for use at CERN. Furthermore, modifications will also be made to OpenStack to allow user-specified Ceph data striping values to be set during volume creation as well as the Ceph/QEMU caching method upon volume attachment to an instance.

Abstract: With the ever increasing amount of data produced from Large Hadron Collider (LHC) experiments, new ways are sought to help analyze and store this data as well as help researchers perform their own experiments. To help offer solutions to such problems, CERN has employed the use of cloud computing and in particular OpenStack; an open source and scalable platform for building public and private clouds. The OpenStack project contains many components such as Cinder used to create block storage that can be attached to virtual machines and in turn help increase performance. However instead of creating volumes locally with OpenStack, others remote storage clusters exist offering block based storage with features not present in the current OpenStack implementation; two popular solutions are NetApp and Ceph. Two features Ceph offers is the ability to stripe data stored within volumes over the distributed cluster as well as locally cache this data, both with the aim of improving performance. When in use with OpenStack, Ceph performs default data striping where the number and size of stripes is fixed and cannot be changed dependent on the volume to be created. Similarly, Ceph does not perform data caching when integrated with OpenStack. In this project we outline and document the integration of NetApp and Ceph with OpenStack as well as benchmark the

performance of the NetApp and Ceph clusters already present at CERN. To allow Ceph data striping, we modify OpenStack to take the number and size of stripes input via the user to create volumes whose data is then striped according to the values they specify. Similarly, we also modify OpenStack to enable Ceph caching and allow users to select the caching policy they require per-volume. In this report, we describe how these features are implemented.

Report on ZENODO:

[Document on ZENODO](#) [2]

- [Visit Us](#)
- [RSS Feeds](#)

DISCLAIMER: This Web page contains pointers to material related to the management of CERN openlab in the Information Technology Department at the European Organization for Nuclear Research (CERN). Their use and distribution are regulated by the [CERN copyright notice](#).



Source URL: http://test-static-05.web.cern.ch/publications/technical_documents/implementation-openstack-cinder-and-integration-netapp-and-ceph

Links

[1] http://test-static-05.web.cern.ch/publications/technical_documents/implementation-openstack-cinder-and-integration-netapp-and-ceph

[2] <https://zenodo.org/record/7570?ln=en>