Contribution ID: **12** Type: **Oral Presentation**

Distributed filesystems (GPFS, CephFS and Lustre-ZFS) deployment on Kubernetes/Docker clusters

Friday, 26 March 2021 12:00 (30 minutes)

Nowadays Kubernetes has become a powerful tool to deploy and manage containerized applications.
Modern datacenters need distributed filesystems to provide user applications with access to stored data on a large number of nodes.
The possibility to mount a distributed filesystem and exploit its native application programming interfaces in a Docker container,
combined with the advanced orchestration features provided by Kubernetes, may enhance flexibility in data management and transfer services' installation, running and monitoring,
allowing the execution of dedicated services on different nodes, in isolated and automatically replicable environments, this way improving deployment efficiency and fail-safeness.
The goal of this work is to demonstrate the feasibility of using Kubernetes and Docker to setup clients capable to access a distributed filesystem from existing clusters and to create clusters based on containerized servers.
Although this is just a proof of concept, the effort has shown the possibility of using different types of distributed filesystems (GPFS, CephFS, Lustre-ZFS) with equally positive results. Read/write performances with these filesystem setups have been tested and compared to each other.

Summary

The goal of this work is to demonstrate the feasibility of using Kubernetes and Docker to setup clients capable to access a distributed filesystem from existing clusters and to create clusters based on containerized servers.

Primary author: Dr FORNARI, Federico (INFN-CNAF)

Presenter: Dr FORNARI, Federico (INFN-CNAF)

Session Classification: Infrastructure Clouds and Virtualisation Session

Track Classification: Infrastructure Clouds and Virtualisation