

Reproducible and Scalable workflows for SkyhookDM experimentation on Kubernetes

Tuesday 1 December 2020 11:20 (20 minutes)

Preparing a Systems experiment environment requires setting up infrastructure, baselining the infrastructure, installing dependencies and tools, running experiments, and manually plotting results, which if done manually, is cumbersome and error-prone. This same scenario applies to researchers starting to experiment with Ceph or SkyhookDM, which is an extension for Ceph to run queries on tabular datasets stored as objects. To address this issue, we used Popper, the container-native workflow engine, to build scalable and reproducible workflows for automating an end-to-end pipeline for experimenting with Ceph and SkyhookDM deployed on Kubernetes via Rook.

Author: CHAKRABORTY, Jayjeet (University of California, Santa Cruz)

Presenter: CHAKRABORTY, Jayjeet (University of California, Santa Cruz)

Session Classification: Overflow & Open Discussion