Contribution ID: **52** Type: **Oral**

Managing the CERN Batch System with Kubernetes

Tuesday, 5 November 2019 17:15 (15 minutes)

The CERN Batch Service faces many challenges in order to get ready for the computing demands of future LHC runs. These challenges require that we look at all potential resources, assessing how efficiently we use them and that we explore different alternatives to exploit opportunistic resources in our infrastructure as well as outside of the CERN computing centre.

Several projects, like BEER, Helix Nebula Science Cloud and the new OCRE project, have proven our ability to run batch workloads on a wide range of non-traditional resources. However, the challenge is not only to obtain the raw compute resources needed but how to define an operational model that is cost and time efficient, scalable and flexible enough to adapt to a heterogeneous infrastructure.

In order to tackle both the provisioning and operational challenges it was decided to use Kubernetes. By using Kubernetes we benefit from a de-facto standard in containerised environments, available in nearly all cloud providers and surrounded by a vibrant ecosystem of open-source projects. Leveraging Kubernetes'built-in functionality, and other open-source tools such as Helm, Terraform and GitLab CI, we have deployed the first cluster prototype which we discuss in detail. The effort has simplified many of the existing operational procedures we currently have, but has also made us rethink many established procedures and assumptions that were only valid in a VM based cloud environment.

This contribution presents how we have adopted Kubernetes in the CERN Batch Service, the impact its adoption has had in daily operations, a comparison on resource usage efficiency and the experience so far evolving our infrastructure towards this model.

Consider for promotion

No

Authors: FERNANDEZ ALVAREZ, Luis (CERN); DATSKOVA, Olga Vladimirovna (University of Houston (US)); JONES, Ben (CERN); MCCANCE, Gavin (CERN)

Presenter: FERNANDEZ ALVAREZ, Luis (CERN)

Session Classification: Track 7 –Facilities, Clouds and Containers

Track Classification: Track 7 –Facilities, Clouds and Containers