



Contribution ID: 61

Type: **not specified**

Too much BEER and the Kubernetes dance

Wednesday 10 October 2018 17:15 (25 minutes)

With the demands of LHC computing, coupled with pressure on the traditional resources available, we need to find new sources of compute power. We have described, at HEPiX and elsewhere how we have started to explore running batch workloads on storage servers at CERN, and on public cloud resources. Since the summer of 2018, ATLAS & LHCb have started to use a pre-production service on storage services (BEER) on 3k cores. We have run cloud services with Helix Nebula, Oracle and others at the scale of 10k cores, till now using fat VMs configured using traditional configuration management but with containers for the payload. The kubernetes ecosystem would seem to promise a simplification of public cloud deployments, as well as internal use cases.

The talk will therefore report on the gained experience of running BEER services at some scale, and at the progress and motivation of putting the condor worker itself in a container, managed with kubernetes.

Desired length

20

Authors: JONES, Ben (CERN); FERNANDEZ ALVAREZ, Luis (CERN)

Presenter: JONES, Ben (CERN)

Session Classification: Computing & Batch Services

Track Classification: Computing & Batch Services