

Second K8s-HEP Meetup

Report of Contributions

Contribution ID: 1

Type: **not specified**

Welcome! What's this?

Tuesday 1 December 2020 09:00 (10 minutes)

Presenters: BRYANT, Lincoln (University of Chicago (US)); GARDNER JR, Robert William (University of Chicago (US))

Session Classification: block 1 - presentations

Contribution ID: 2

Type: **not specified**

K8s autoscaling based on custom metrics. Two examples of application: CMSWEB and HTCondor in the CMS Analysis Facility@INFN

Tuesday 1 December 2020 09:50 (20 minutes)

At the moment, Kubernetes only supports horizontal pod autoscaling based on predefined pod metrics (CPU and memory usage). Therefore, in order to achieve an actually green elastic cloud model (optimizing resource usage) a key point is to integrate this tool with autoscaling solutions based on custom metrics, and this requires the usage of third-party elements.

In this work we show the horizontal pod autoscaling based on custom metrics: in this workflow metrics are collected by a Prometheus server, and are then manipulated and made available to k8s-native Horizontal Pod Autoscaler (HPA) resources.

We show how we apply the presented feature to two HEP-related use cases: in the first one this solution is applied to CMSWEB (i.e. CMS web services) infrastructure, in the second one it is used to enhance elasticity of an analysis facility prototype on INFN-Cloud, with the automatic scaling of HTCondor instances.

Authors: SPIGA, Daniele (Universita e INFN, Perugia (IT)); CIANGOTTINI, Diego (INFN, Perugia (IT)); TEDESCHI, Tommaso (Universita e INFN, Perugia (IT)); KUZNETSOV, Valentin Y (Cornell University (US))

Presenter: TEDESCHI, Tommaso (Universita e INFN, Perugia (IT))

Session Classification: block 1 - presentations

Contribution ID: 3

Type: **not specified**

Lightweight integration of Kubernetes clusters for ATLAS batch processing

Tuesday 1 December 2020 10:10 (20 minutes)

The PanDA team has evaluated the possibility of native Kubernetes job submission in order to process ATLAS workloads and offer the possibility of immediate integration of major cloud computing providers. This model also offers a novel way to set up lightweight compute sites, without the need of setting up a Grid stack.

During the last year we have been running several queues at clusters setup by institutes associated to ATLAS (ASGC, CERN, University of Chicago, University of Victoria) and on cloud providers (Amazon and Google), and have focused on increasing the stability and efficiency.

This contribution will discuss the advantages and challenges we have faced during our experience and also briefly introduce ongoing work to integrate less trivial (non pleasantly parallel) workloads.

Authors: BARREIRO MEGINO, Fernando Harald (University of Texas at Arlington); LIN, Fa-Hui (University of Texas at Arlington (US))

Presenter: BARREIRO MEGINO, Fernando Harald (University of Texas at Arlington)

Session Classification: block 1 - presentations

Contribution ID: 4

Type: **not specified**

Running a multi-tenant Kubernetes with GitOps

Wednesday 2 December 2020 09:20 (20 minutes)

Starting in October 2020, the PATH project is making a concerted effort to transition the centrally-run OSG services (such as websites, software repositories, information services) from ad-hoc deployment models to Kubernetes.

To do so, we needed a Kubernetes “home” and an operational model! In this talk, we’ll overview the work going on in the Tiger cluster at Morgridge, our current GitOps-based workflow with Flux, and how we see things fitting into the larger ecosystem of distributed services and federated Kubernetes.

Author: BOCKELMAN, Brian Paul (University of Wisconsin Madison (US))

Presenter: BOCKELMAN, Brian Paul (University of Wisconsin Madison (US))

Session Classification: block 3 - presentations

Contribution ID: 5

Type: **not specified**

Multi Cluster / Cloud Kubernetes for GPU Evaluation

Wednesday 2 December 2020 09:00 (20 minutes)

GPUs are scarce resources in many of our centers, including CERN.

This talk will quickly describe a multi cloud deployment with the goal of evaluating the performance of different workloads in all GPUs offered by GCP, Azure and AWS.

It will include some details about setting up clusters and GPUs in each of these clouds, and some preliminary results.

Author: BRITO DA ROCHA, Ricardo (CERN)

Presenter: BRITO DA ROCHA, Ricardo (CERN)

Session Classification: block 3 - presentations

Contribution ID: 6

Type: **not specified**

Debugging Kubernetes pod throughput with Calico CNI

Tuesday 1 December 2020 09:30 (20 minutes)

Exploring how the kubelet, with Calico as the CNI plugin, depends on the performance of the Kubernetes API server to be able to start pods quickly.

Author: Mr HARTLAND, Thomas George (CERN)

Presenter: Mr HARTLAND, Thomas George (CERN)

Session Classification: block 1 - presentations

Contribution ID: 7

Type: **not specified**

Experience with K8s at Coffea-Casa AF@UNL

Wednesday 2 December 2020 10:00 (20 minutes)

In this contribution we would like to share our experience designing an Analysis Facility for the columnar analysis utilizing the analysis package COFFEA at University of Nebraska-Lincoln and to describe our adventure on deploying different workloads and services at UNL Kubernetes cluster (Jupyterhub with Traefik integration, HTCondor, ServiceX and other infrastructure deployments).

Author: LUNDSTEDT, Carl (University of Nebraska Lincoln (US))

Co-author: SHADURA, Oksana (University of Nebraska Lincoln (US))

Presenter: LUNDSTEDT, Carl (University of Nebraska Lincoln (US))

Session Classification: block 3 - presentations

Contribution ID: 8

Type: **not specified**

Overview of CMSWEB Cluster in Kubernetes

Wednesday 2 December 2020 09:40 (20 minutes)

The CMS experiment heavily relies on the CMSWEB cluster to host critical services for its operational needs. The cluster is deployed on virtual machines (VMs) from the CERN OpenStack cloud and is manually maintained by operators and developers. The release cycle is composed of several steps, from building RPMs, their deployment, validation, and integration tests. To enhance the sustainability of the CMSWEB cluster, CMS decided to migrate its cluster to a containerized solution such as Docker, orchestrated with Kubernetes (k8s). This allows us to significantly reduce the release upgrade cycle, follow the end-to-end deployment procedure, and reduce operational cost.

Recently, we have performed the migration of some CMSWEB services from the VM cluster to the Kubernetes. This talk gives an overview of the current CMSWEB cluster. We describe the new architecture of the CMSWEB cluster in Kubernetes and its implementation strategy. We'll discuss how we create docker images of the services and deploy them in this cluster using the service deployment cycle. Furthermore, we'll discuss about monitoring of these services. In the end, we'll discuss our future plan related to this cluster.

Authors: IMRAN, Muhammad (National Centre for Physics (PK)); KUZNETSOV, Valentin Y (Cornell University (US)); Dr PFEIFFER, Andreas (CERN); PAPARRIGOPOULOS, Panos (CERN)

Presenter: IMRAN, Muhammad (National Centre for Physics (PK))

Session Classification: block 3 - presentations

Contribution ID: 9

Type: **not specified**

Fermilab Experience with OKD (OpenShift)

Tuesday 1 December 2020 09:10 (20 minutes)

Fermilab has made the strategic decision to deploy OKD, the open source version of Red Hat OpenShift, for Kubernetes container management. We will discuss our experience so far with OKD and describe some of the challenges we faced deploying a variety of applications.

Author: TIRADANI, Anthony (Fermilab)

Presenters: TIRADANI, Anthony (Fermilab); TIRADANI, Anthony Richard (Fermi National Accelerator Lab. (US))

Session Classification: block 1 - presentations

Contribution ID: 10

Type: **not specified**

Lazy Image Pulling with Stargz

Tuesday 1 December 2020 10:30 (20 minutes)

Container images allowed having reproducible environments and container orchestration lets users parallelize and create elaborate workflows with tools like Argo or just Kubernetes jobs. It is easy to create very large images and when parallelizing jobs, the time and cost of pulling container images can increase significantly. Golang developers proposed the seekable tar gunzip format (stargz) to address the issue for their CI by downloading files from the container registry only when they are needed. This presentation describes the current state of lazy loading images with containerd and stargz and we will present the results of our benchmarks.

Author: TRIGAZIS, Spyridon (CERN)

Presenter: TRIGAZIS, Spyridon (CERN)

Session Classification: block 1 - presentations

Contribution ID: 11

Type: **not specified**

What's new with SLATE?

Wednesday 2 December 2020 10:40 (20 minutes)

Will review progress over the past year with SLATE - including new containerized apps, storage provisioner, security policies for federated operations

Authors: GARDNER JR, Robert William (University of Chicago (US)); BRYANT, lincoln

Presenter: BRYANT, lincoln

Session Classification: block 3 - presentations

Contribution ID: 12

Type: **not specified**

Test REANA Deployment at BNL

Wednesday 2 December 2020 10:20 (20 minutes)

In this presentation we'll discuss our experiences deploying a test REANA instance on a k8s cluster at BNL.

Author: HOLLOWELL, Christopher Henry (Brookhaven National Laboratory (US))

Presenter: HOLLOWELL, Christopher Henry (Brookhaven National Laboratory (US))

Session Classification: block 3 - presentations

Contribution ID: 13

Type: **not specified**

Kubernetes at UVic

Wednesday 2 December 2020 11:00 (20 minutes)

I will describe Kubernetes cluster deployment at UVic, including batch computing and APEL accounting for ATLAS.

Author: TAYLOR, Ryan (University of Victoria (CA))

Presenter: TAYLOR, Ryan (University of Victoria (CA))

Session Classification: block 3 - presentations

Contribution ID: 14

Type: **not specified**

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

Contribution ID: 15

Type: **not specified**

Packaging and using services in Kubernetes

Wednesday 2 December 2020 11:20 (20 minutes)

OSG lessons learned distributing service container images and experiences contributing to and deploying services with SLATE

Author: LIN, Brian Hua (University of Wisconsin - Madison)

Presenter: LIN, Brian Hua (University of Wisconsin - Madison)

Session Classification: block 3 - presentations

Contribution ID: **16**

Type: **not specified**

Discussion

Tuesday 1 December 2020 11:40 (1h 20m)

Session Classification: Overflow & Open Discussion

Contribution ID: 17

Type: **not specified**

Discussion

Wednesday 2 December 2020 12:05 (55 minutes)

Session Classification: Overflow & Open Discussion