



Contribution ID: 444 Contribution code: WED 02

Type: Poster

Exploiting Kubernetes to Simplify the Deployment and Management of the Multi-purpose CMS Pilot Job Factory

Wednesday 23 October 2024 15:33 (15 minutes)

GlideinWMS, a widely utilized workload management system in high-energy physics (HEP) research, serves as the backbone for efficient job provisioning across distributed computing resources. It is utilized by various experiments and organizations, including CMS, OSG, Dune, and FIFE, to create HTCondor pools as large as 600k cores. In particular, a shared factory service historically deployed at UCSD has been configured to interface with more than 500 routes to compute clusters.

As part of our team's initiative to modernize infrastructure and enhance scalability, we undertook the migration of the glideinWMS factory service into the Kubernetes environment. Leveraging the flexibility and orchestration capabilities of Kubernetes, we successfully deployed the factory service within the OSG Tiger Kubernetes cluster. The major benefits Kubernetes gives us is it streamlines the management and monitoring of the factory infrastructure, and improves fault tolerance through its resilient deployment strategies.

Through this case study, we aim to share insights, challenges, and best practices encountered during the migration process. Our experience underscores the benefits of embracing containerization and Kubernetes orchestration for HEP computing infrastructure, paving the way for scalability and resilience in distributed computing environments.

Primary authors: DOST, Jeffrey Michael (Univ. of California San Diego (US)); MASCHERONI, Marco (Univ. of California San Diego (US))

Co-authors: BOCKELMAN, Brian Paul (University of Wisconsin Madison (US)); Mr WALSWORTH, Colby (University of California San Diego); KIZINEVIC, Edita (CERN); WURTHWEIN, Frank (UCSD); THILTGES, John (University of Nebraska Lincoln (US)); ALBERT, Merina (Fermi National Accelerator Lab. (US)); ZOKAITE, Vaiva (Vilnius University (LT))

Presenter: MASCHERONI, Marco (Univ. of California San Diego (US))

Session Classification: Poster session

Track Classification: Track 4 - Distributed Computing