

Building a Kubernetes infrastructure for CERN's Content Management Systems

Thursday, 20 May 2021 15:39 (13 minutes)

The infrastructure behind `home.cern` and 1000 other Drupal websites serves more than 15,000 unique visitors daily. To best serve the site owners, a small engineering team needs development speed to adapt to their evolving needs and operational velocity to troubleshoot emerging problems rapidly. We designed a new Web Frameworks platform by extending Kubernetes to replace the ageing physical infrastructure and reduce the dependency on homebrew components.

The new platform is modular, built around standard components and thus less complex to operate. Some requirements are covered solely by upstream open source projects, whereas others by components shared across CERN's web hosting platforms. We leverage the Operator framework and the Kubernetes API to get observability, policy enforcement, access control and auditing, and high availability for free. Thanks to containers and namespaces, websites are isolated. This isolation clarifies security boundaries and minimizes attack surface, while empowering site owners.

In this work we present the open-source design of the new system and contrast it with the one it replaces, demonstrating how we drastically reduced our technical debt.

Primary author: SAMARAS-TSAKIRIS, Konstantinos (CERN)

Co-authors: RAJULA, Vineet Reddy (CERN); BORGES AURINDO BARROS, Francisco (Universidade de Lisboa (PT)); ALVAREZ FERNANDEZ, Eduardo (CERN); WAGNER, Andreas (CERN)

Presenter: SAMARAS-TSAKIRIS, Konstantinos (CERN)

Session Classification: Virtualisation

Track Classification: Distributed Computing, Data Management and Facilities