

How to make kubernetes easy to use for thousands users

Authors: Alberto Pimpo, Alexandre Lossent, Aleksandra Wardzinska, Michal Kolodziejski (CERN)

Presenter: Ismael Posada Trobo (CERN)

Introduction

Kubernetes is the state-of-the-art platform for deploying applications, but it often introduces operational overhead and has a steep learning curve. CERN has thousands of active developers, and our goal is to enable them to deploy their code without needing in-depth knowledge of DevOps technologies. The solution must be both user-friendly and flexible as needed.

To achieve this, we implemented an infrastructure based on OKD, a Kubernetes distribution optimized for multitenancy. Our solution consists of multiple specialized OKD clusters, all managed through a unified web portal for streamlined application deployment. This allows users to completely avoid dealing with the Kubernetes layer, while still giving them the option to interact with it if needed.

We identified 4 kind of use cases and built 4 different clusters around them:

- One-click provisioning of popular apps (Grafana, Discourse, Nexus)
- Creation of websites using CMS (Wordpress, Drupal)
- Simple web hosting (Static content, simple PHP or Python scripts)
- Generic application hosting (monolithic or microservices based applications)



Creating specific purpose clusters using Operator SDK

One of the key technologies that enables us to run specific purpose clusters is Operator SDK.

Operators enable the codification of operational knowledge, allowing clusters to handle routine tasks like configuration updates, scaling, and failure recovery autonomously. This reduces the need for manual intervention and simplifies the deployment of specialized workloads, ensuring that clusters remain optimized for their intended use cases while maintaining Kubernetes-native flexibility and scalability.

Operators are used to provide one-click provisioning of popular apps (Grafana, Discourse, Nexus), but also to provide CMS managed websites (Wordpress, Drupal) and simple web hosting. The user can go to the unified portal, provision any of these applications and use the service with a complete abstraction from the Kubernetes level.



One-click code deployment with full control of the kubernetes resources

With OKD, we enable one-click code deployment while giving developers full control over Kubernetes resources. OKD provides a workflow called S2I (source to image). Users can provide a git repo and the cluster will automatically detect a strategy to deploy it.

S2I takes the source code, combines it with a predefined builder image (which includes the runtime and dependencies), and creates a ready-to-deploy container image. This process eliminates the need for Dockerfiles, making it easier and faster for developers to deploy their application. If a Dockerfile is present in the repo, this will be used to build the container image that will be deployed.

Furthermore, S2I deployments can be easily configured to be automatically redeployed at each new commit.

If the users have particular requirements on how the image has to be built, or want to use an already existing image, they can directly deploy an image from an external registry with one click.

If they want to have full control over their deployment, they can deploy their application using helm charts.

Git

Git Repo URL *

Validated

Show advanced Git options

Builder Image detected.

A Builder Image is recommended.



Red Hat OpenJDK 17 (UBI 8)

Edit Import Strategy

BUILDER JAVA OPENJDK

Build and run Java applications using Maven and OpenJDK 17 upon UBI8.

Sample repository: <https://github.com/jboss-openshift/openshift-quickstarts>

One portal to rule them all

To simplify access and streamline the user experience across multiple Kubernetes clusters, we developed an in-house portal using Python and React. This portal serves as a unified interface where users can manage all their applications, regardless of the cluster they reside on. It offers a single entry point for creating and deploying applications, effectively presenting a seamless and cohesive system.

The portal not only allows for quick and easy deployment of new applications but also provides visibility for the features offered by the other clusters.

This unified approach gives users the perception of interacting with a single, cohesive system, masking the complexity of the underlying infrastructure.

