Updated 07/27/2023

Sign up

Kanister training: Application data management on Kubernetes

3 days (21 hours)

Presentation

Kanister training will teach you how to capture application-specific data management tasks. The technology supports all the details of running on Kubernetes. In this course, you'll restrict data management tasks into plans that enable uniform operations at large scale. You'll provide a set of required data management primitives. Customize plans to meet the specific needs of your environment and workloads without starting from scratch. With this tool, you leverage a robust collection of plans for common persistent-state applications. With Kanister, create CustomResourcesDefinitions in which you can fully manage the object lifecycle via the kubectl command line. The framework uses Kubernetes' code generation tools to create client libraries for its CustomResources. As with all our training courses, this one will introduce you to the latest stable version of Kanister (at the time of writing, Kanister v0.89).

Objectives

- Understanding Kanister's architecture
- How to set up a personalized Kanister plan
- Manage Kubernetes execution and monitoring details
- Master key Kanister concepts

Target audience

- Developers
- Engineers
- DevOps

Prerequisites

Basic knowledge of Kubernetes.

Kanister training program

Tool introduction

- Kanister tool presentation
- Design objectives
  - Extensible
  - API-Driven
  - Application-Centric
- Installation
  - Kanister configuration
  - Control of custom resource definitions (CRDs)
  - Building and deploying from source

Architecture

- Workflows
  - Creating an ActionSet
  - Examining the Blueprint environment
- Customized resources
  - CustomResourceDefinitions (CRD) creation
  - Using Kubernetes code generation
- Plans
- Set of actions
- Profiles
- Joystick

Tasks

- Automate ActionSet creation with Argo Cron workflows
- Separate controller and data path logs
  - Deployment configuration
  - Newspaper segregation
- Change Kanister's logging level

Tools

- Kanctl
- Kando
- Docker Image

The functions

- Existing functions
  - KubeExec
  - KubeExecAll
  - KubeTask
  - ScaleWorkload
- The data
- Volume snapshot
- Instant RDS
- CubeOps
- WaitV2

Model parameters

- Rendering templates
- Objects
  - StatefulSet
  - Deployments
  - Name space
  - PVC
- Artifacts
  - Input artifacts
  - Output artifacts
- Configuration

Troubleshooting

- Webhook validation for Blueprints
- Firewall check

Companies concerned

This training course is aimed at both individuals and companies, large or small, wishing to train their teams in a new advanced computer technology, or to acquire specific business knowledge or modern methods.

Teaching methods

Practical course: 60% Practical, 40% Theory. Training material distributed in digital format to all participants.

Organization

The course alternates theoretical inputs from the trainer supported by examples and

brainstorming sessions and group work.

Validation

At the end of the session, a multiple-choice questionnaire verifies the correct acquisition of skills.

Sanction

A certificate will be issued to each trainee who completes the course.