

# Kubernetes Database-as-a-Service

Degree programme : BSc in Computer Science | Specialisation : Data Engineering  
Thesis advisor : Prof. Dr. Mascha Kurpicz-Briki  
Expert : Dr. Joachim Wolfgang Kaltz  
Industrial partner : Bedag Informatik AG, Bern

This Bachelor's Thesis brings forth an open-source Kubernetes Operator able to provide a self-service mechanism for database provisioning in external DB management solutions. End-users such as software developers can request database instances by writing a simple YAML configuration file. Database secrets are stored using Kubernetes Secrets, giving applications transparent, secure access to user credentials.

## Introduction

It is no secret that we are recently witnessing a massive increase in the development of cloud-native applications all around the industry. Companies are embracing the transparency and collaboration enabled by open-source initiatives. In this modern day and age, the awareness around data ownership and the administration of mission critical infrastructure has also increased due to growing concerns in organizations bound to strict compliance requirements. This project aims at creating a Database-as-a-Service (DBaaS) using an innovative, modern approach.

## Use case

There are examples where an organization cannot or does not want to host their critical data in cloud environments and is looking for a way to bridge the gap between their Kubernetes clusters and on-premise DBMS solutions. Database administrators can retain full control on the life cycle of database instances by creating stored procedures or equivalent mechanisms for each DBMS. Organizations can configure an opaque provisioning system for databases where data and business logic is kept as close as possible to their location.

## Kubernetes operators

An operator is a specific pattern used to handle the life cycle of Kubernetes resources. Its goal is to capture the human natural way of performing tasks in order to automate processes that would otherwise be carried out manually. Due to the infinite number of possibilities when deploying and administering an application, Kubernetes can be extended using the operator pattern with the intention of encapsulating complex business logic, such as interacting with external services and performing a series of tasks. The Operator decouples the configuration from the implementation and ensures a well-defined boundary between the Kubernetes and database worlds.

## Enterprise software

In an effort to bring the Operator to an acceptable degree of quality and be deployed in the infrastructure of Bedag Informatik AG, the Operator includes features such as level-based logging and metrics, PostgreSQL, MariaDB and Microsoft SQL Server support, unit and integration testing, rate-limiting, credential rotation and more. Finally, the Operator can be easily deployed using the enclosed Helm Chart, ensuring a robust deployment method.



Cristiano Colangelo  
cristiano.colangelo@bluewin.ch

