

# DevOps

## Abstract

This page contains the technical task description for the DevOps Engineer.

## Requirements

We need to setup stable infrastructure that will allow the developers use CI/CD process and have several environments. To achieve that, we need:

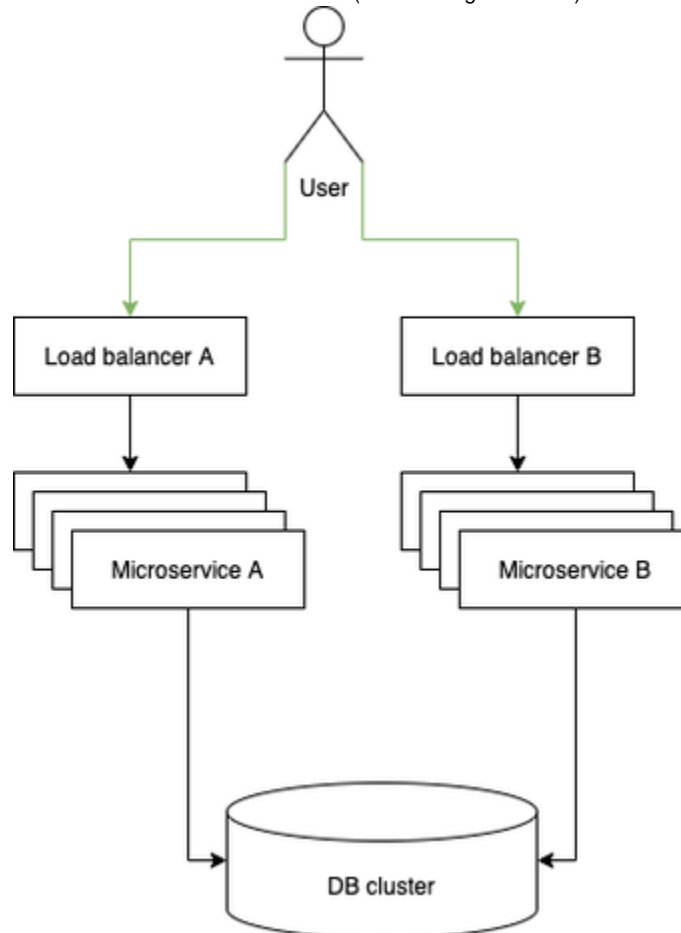
- Kubernetes cluster (that could be dynamically extended or shrunk)
- Deployment tool (Jenkins or ArgoCD)
- Relational database cluster (MySQL)
- Kafka instance
- Secrets vault storage
- NoSQL database cluster (MongoDB)
- Kibana & Grafana (should collect and visualize log data with the possibility to filter by environment, time, micro service type, etc)

By default, there should be three environments configured:

- Development environment. Should be redeployed on any change of repository **develop** branch. The redeployment should be invisible (e.g. startup new instances and after that shut down the old ones)
- Test environment. Should be redeployed by single click (in Jenkins or Argo). The source of the code is **test** branch of repository.
- Production environment. Should work in the same way as test, but the source of the code is **master** branch

We will have several micro services (let's assume N). We should have the possibility to redeploy any group of micro services independently.

Each group of micro services should be hidden behind the load balancer (see the diagram below).



The DB clusters shouldn't be accessible from the public network, but VPN/SSH tunnel should be configured for the developers that will allow to access the database from developers machines.

The vault should be configured for the dev, test and production environment. All the secrets should be there. The possibility to create and manage the secrets should be provided to development team.