

Case study

Building a cloud-native application with Kubernetes

Neptune is a cloud-native data science lab that enables individuals and teams of data scientists to run multiple experiments simultaneously, thus shortening the time required to achieve their results. It also gives them the ability to collaborate, share results and manage the model training process.

Originally built by CodiLime, Neptune is currently being developed by CodiLime spin-off, Neptune Labs, Inc.

Challenge

Data scientists are currently in high demand as more and more companies use machine learning to support their businesses. Most of them are more scientists or mathematicians than engineers, and handling infrastructure is usually outside of their comfort zone. The business challenge here was to make their work as painless as possible.

This goal needed to be addressed, taking into account also another challenge. The need for computing power varies greatly between different machine learning projects, and even within a single project, the requirements for the number of GPUs can rapidly change.

Results & benefits

We decided to use Kubernetes as a layer of abstraction that separates data scientists from the low-level infrastructure problems.

Moreover, Neptune delivers a tailor-made autoscaling solution that is faster than those available off the shelf. It uses Kubernetes to smoothly handle the fluctuating amount of resources.

As a result, users are presented with a cost-effective platform that fits their needs.

The benefits of the solution include:

- It accommodates computing-heavy data experiments
- It offers flexible infrastructure management: on-demand scaling of the Kubernetes cluster
- GPUs are provisioned on demand only
- It has a cloud-native design

Usage statistics

100K
Experiments done
in the cloud

3000+
Users

15k+
Hours of GPU virtual
machines used

99.99%+
Availability

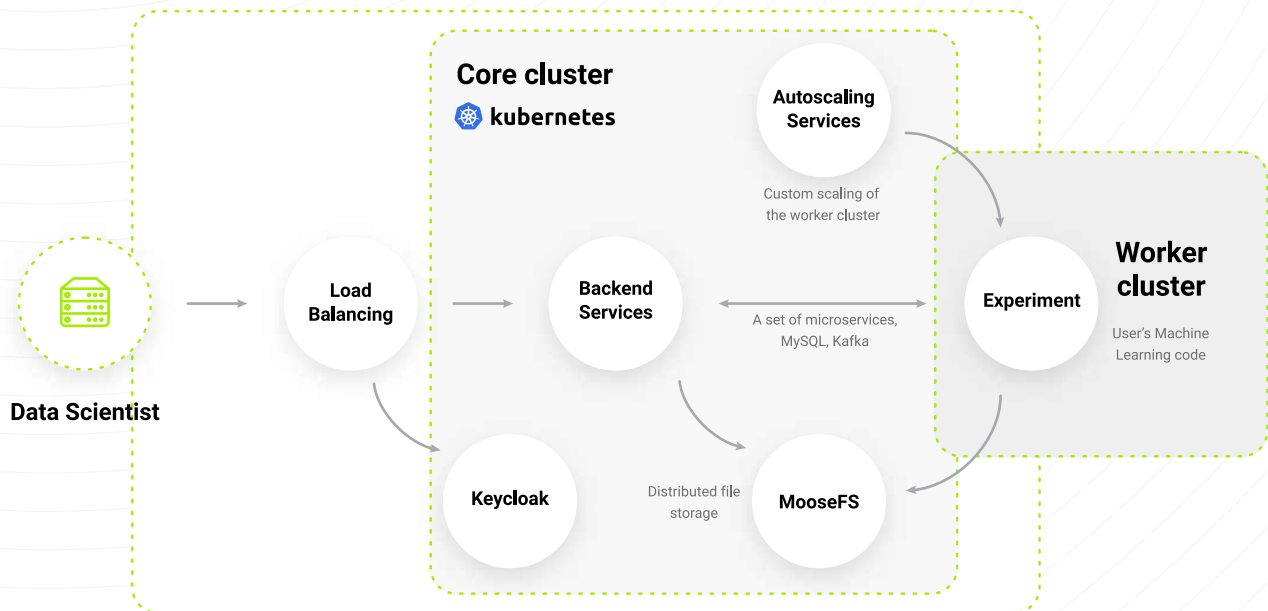
Solution

Neptune runs on Kubernetes and uses Helm templates to reduce the time needed to run new machines and start an experiment. The underlying Kubernetes cluster smoothes the process of establishing and closing experiment containers.

At the same time, by leveraging MooseFS (a distributed filesystem), Neptune ensures that all containers share access to the training dataset, making additional storage for every machine unnecessary.

Finally, Kubernetes makes Neptune infrastructure-agnostic, so it can be established in either a private or public cloud. Neptune can be run on a laptop, using cloud resources, or on a bare-metal infrastructure.

A graphic overview of Neptune



Technologies we used



About CodiLime

CodiLime helps companies build and transform their network applications by providing expert knowledge and technical teams experienced in building, developing and integrating SND & NFV solutions in a single-, multi- or hybrid-cloud environment.

Trusted by leaders since 2011



Should you need any further information regarding our services,
feel free to contact us at contact@codilime.com

For more info about CodiLime, go to codilime.com

CodiLime, Inc.

2100 Geng Road, Suite 210
Palo Alto, CA 94303, United States of America
Phone: **+1 650 285 2458**

CodiLime Sp. z o.o.

Krancowa 5
02-493 Warsaw, Poland
Phone: **+48 22 389 51 00**

