Neptune.ai

case study

codilime ✕ .∿ neptune.ai

Building a cloud-native application for a data science lab

Neptune.ai started out as a cloud-native data science lab. Individuals and teams of data scientists were able to run multiple experiments simultaneously, thus shortening the time required to achieve their results. Neptune took on the complexity of infrastructure management in the experimentation process, while facilitating collaboration, sharing results, and managing the process of building an ML model.

Currently, Neptune's mission is to give production ML teams the same level of control, while building models, as software developers have during applications shipping.

→

Challenge

When starting this project, both the CodiLime and Neptune.ai teams were aware that building machine learning models requires massive computing power. To achieve that the correct selection of technologies and devices was crucial at this point.

At the beginning, there was a spike in the demand for computing power. We were able to meet these requirements, and by the end they had dropped sharply.

The CodiLime team helped with building a platform that is easy to manage and cost-effective for users.

What we delivered

Initial implementation of the service, including a custom Kubernetes autoscaler

1

Infrastructure design and implementation
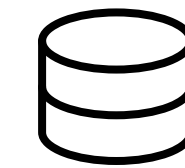
2

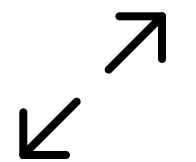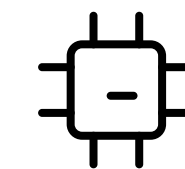**IT support**

3

IT advice

4

Results & benefits

Successful accommodation of computing-heavy data experiments

Flexible infrastructure management: on-demand scaling on the Kubernetes cluster

GPUs provisioned on demand only

A cloud-native design

Client's testimonial

neptune.ai

CodiLime's expertise in software engineering was indispensable at the beginning of Neptune's journey. CodiLime helped us build and release the first version of Neptune and convince investors that our product is a real game changer for data scientists.

Without any hesitation, I can recommend CodiLime as a reliable technology partner for every tech startup.

Piotr Niedźwiedź
CEO of Neptune.ai

→

The scope of the project

As Neptune runs on Kubernetes, the whole process of establishing and closing VMs becomes seamless and smooth.

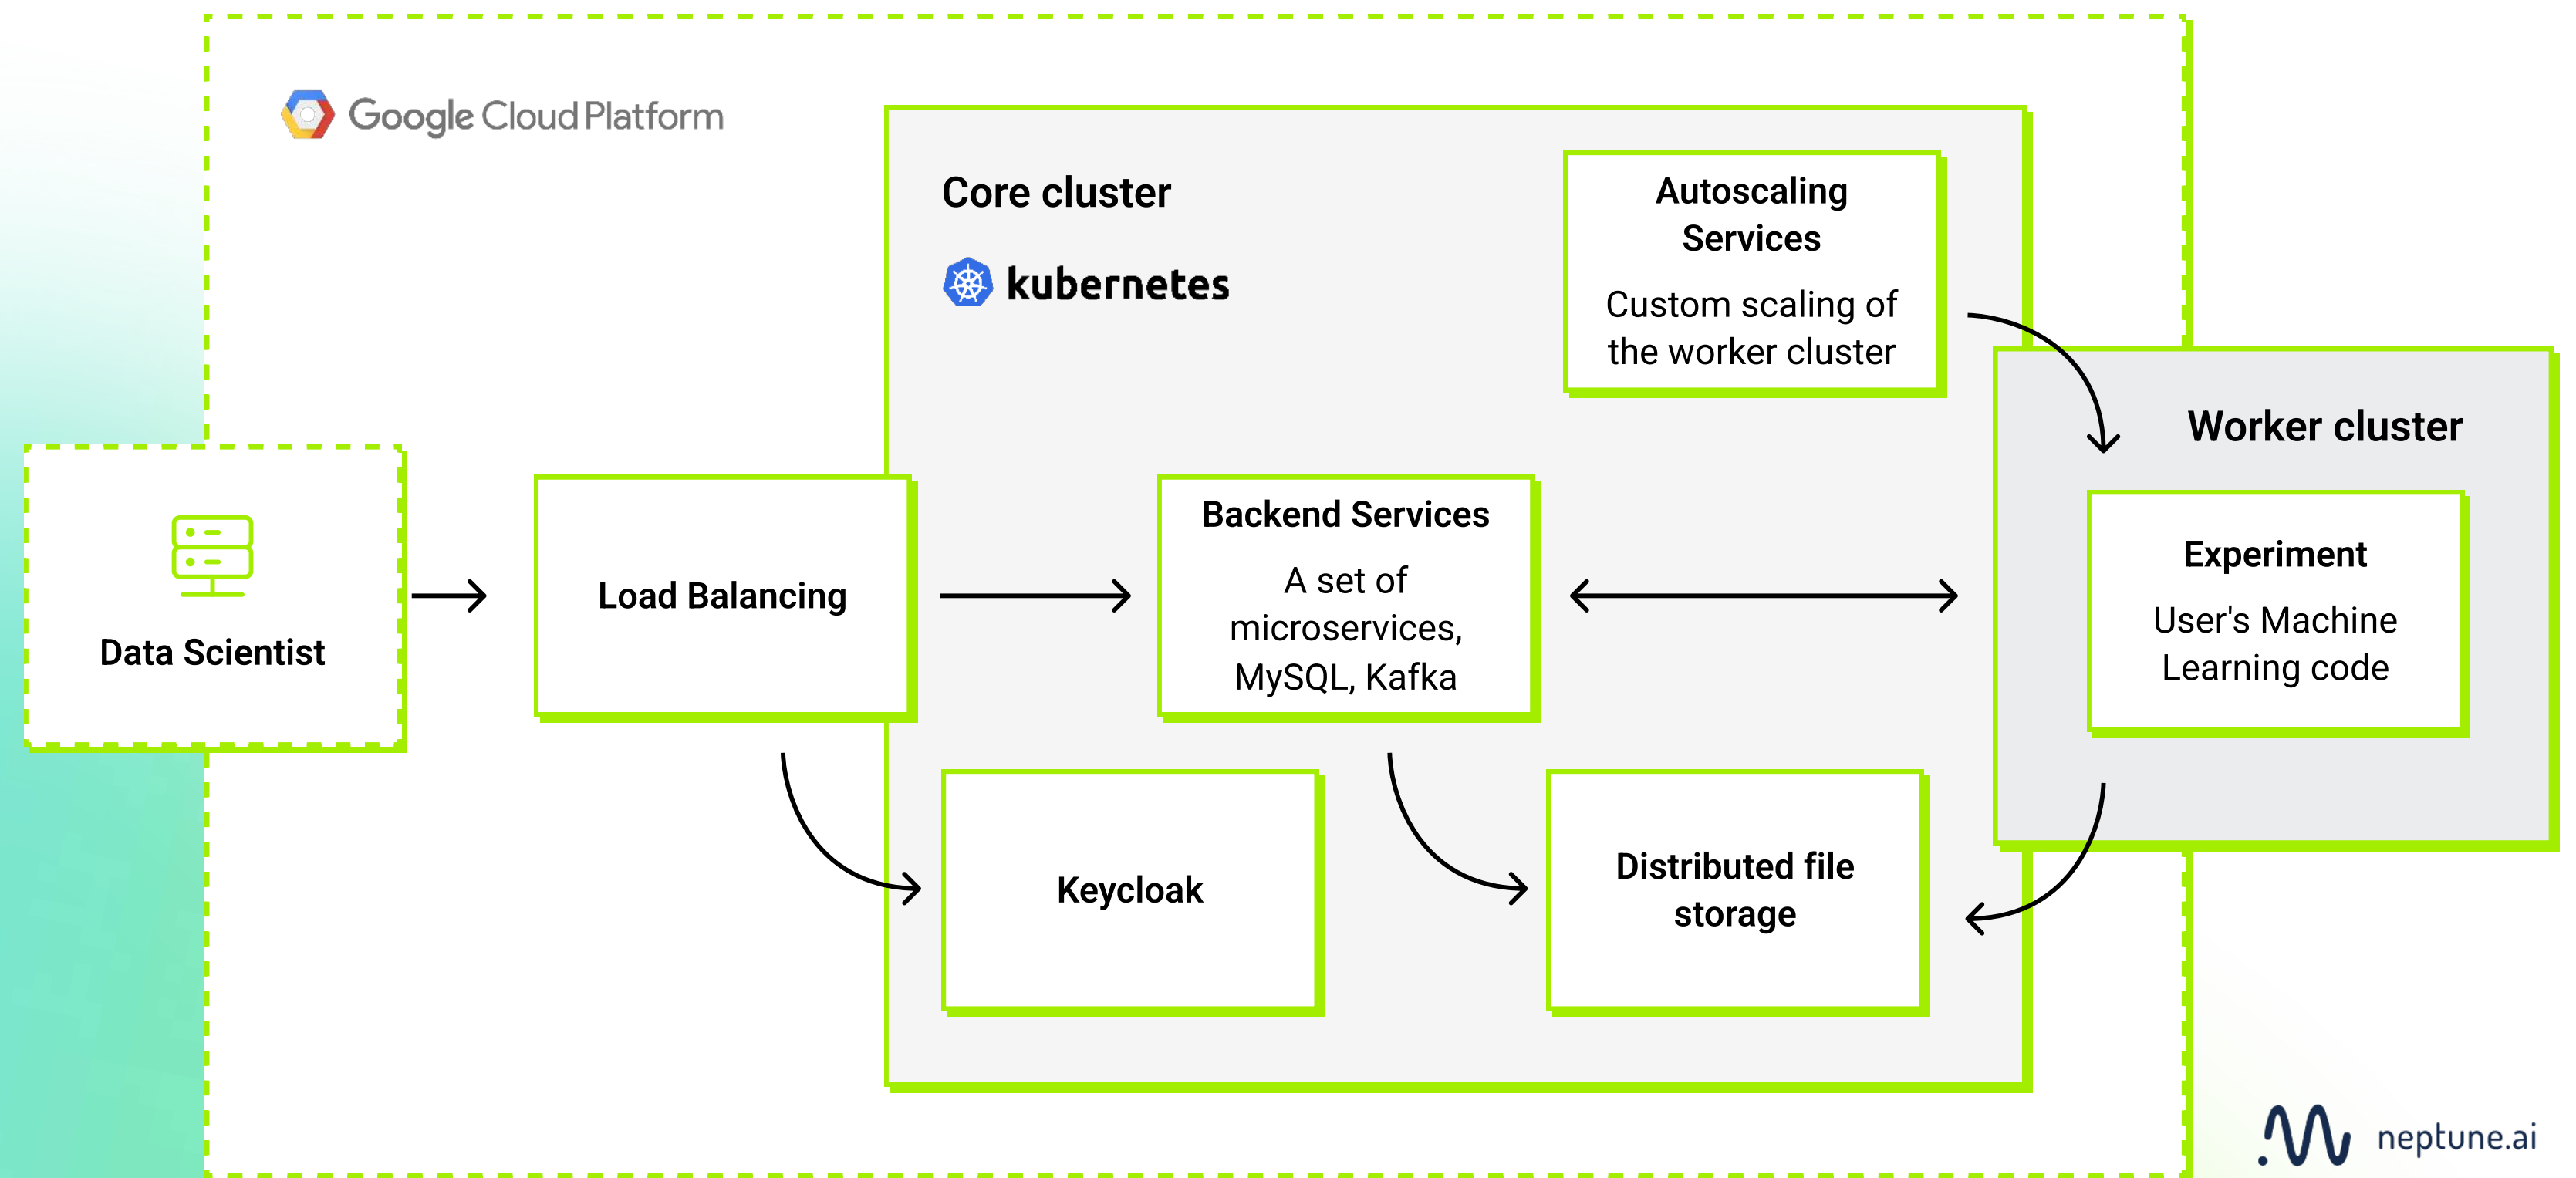Using Helm templates and a custom scaling service:

- Neptune was able to reduce the time needed to run new machines,

- and finally launch the whole experiment.

With CodiLime's help, Neptune ensured that all VMs shared access to the training dataset being used, making experimentation faster and providing non-trivial cost optimizations. Using Kubernetes' granular access control, Neptune.ai was able to guarantee that only authorized workloads had access to their data.
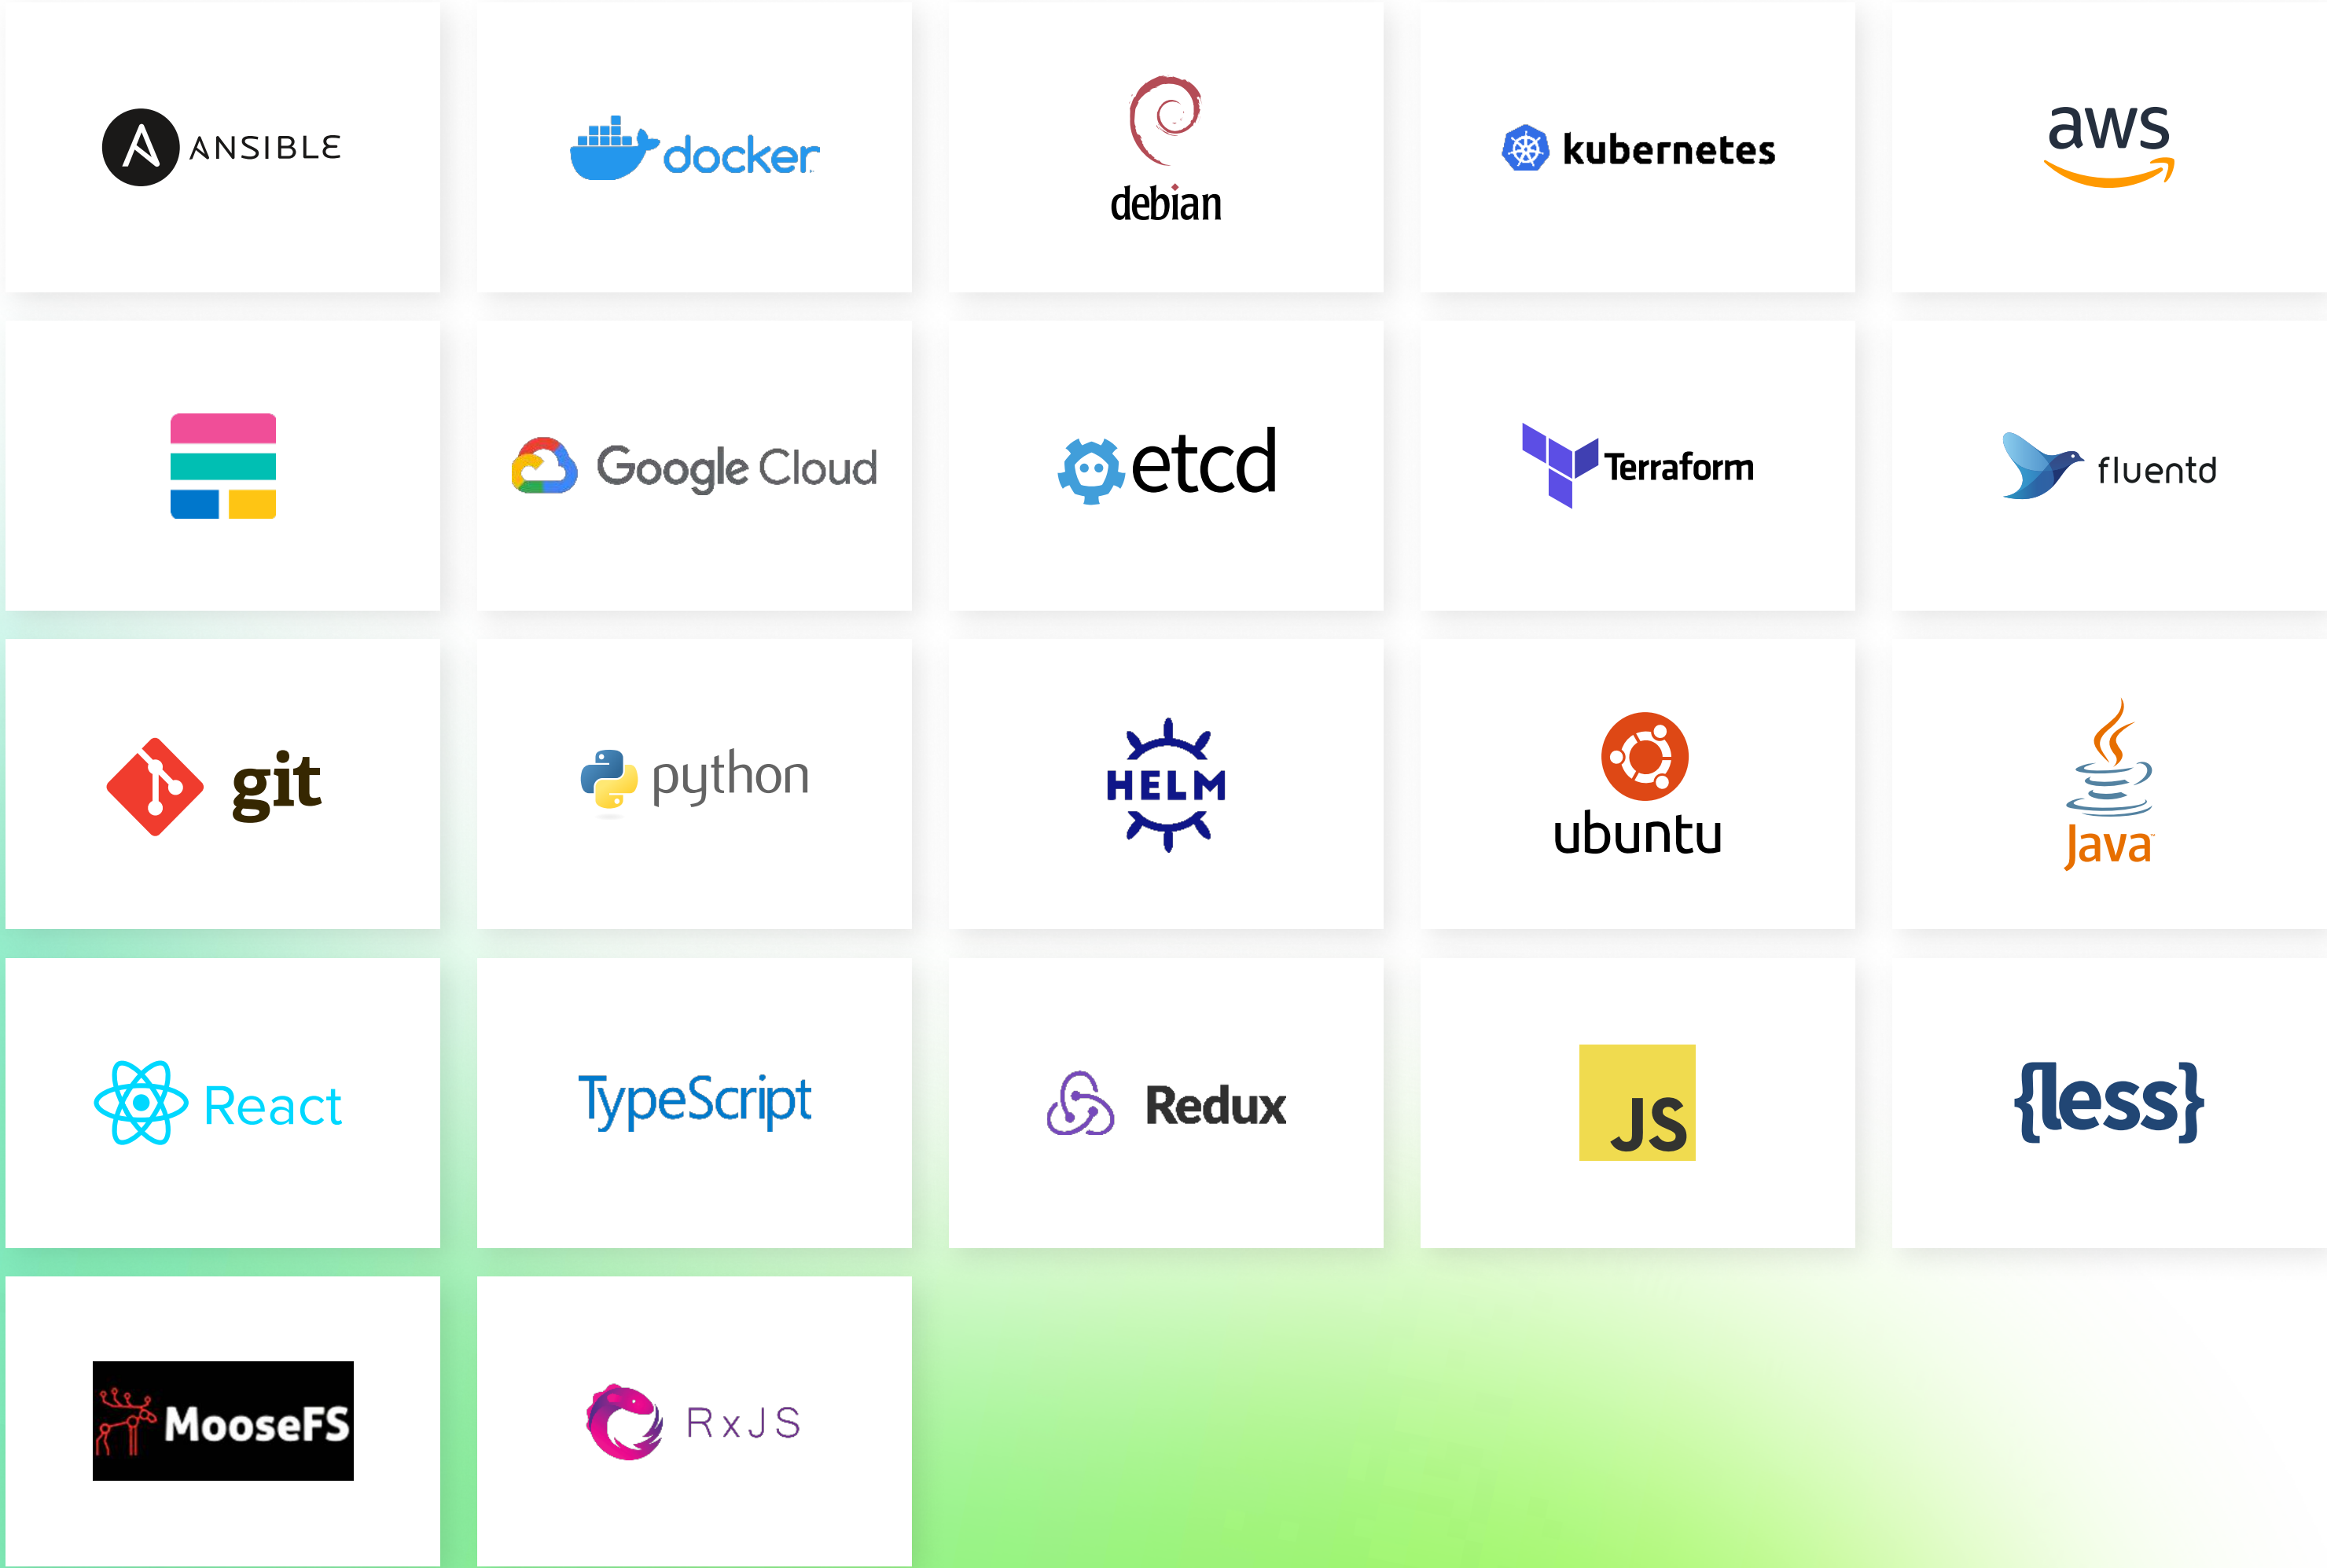
Last but not least, the usage of Kubernetes makes Neptune infrastructure-agnostic, so it can be established in both a private or public cloud.

→

The scope of the project

Google Cloud Platform

Core cluster

kubernetes

Data Scientist

Load Balancing

Autoscaling Services
Custom scaling of the worker cluster

Backend Services
A set of microservices, MySQL, Kafka

Worker cluster

Experiment
User's Machine Learning code

Keycloak

Distributed file storage

neptune.ai

Technologies and tools

Ansible

docker

debian

kubernetes

aws

Google Cloud

etcd

Terraform

fluentd

git

python

HELM

ubuntu

Java

React

TypeScript

Redux

JS

{less}

MooseFS

RxJS

About CodiLime

codilime

Since 2011, CodiLime has been the engineering partner of choice for semiconductor companies, networking services, telecom services, and software solution providers.

We have come a long way – from a startup to a company that hires more than 350 top-notch software developers, network engineers, DevOps experts, and solution architects. **We focus on five N.E.E.D.S. - Networks, Equipment, Environment, Data and Security.**

We aim to link engineering talent with business domain expertise. Everything to provide our clients with delivery excellence and custom-tailored solutions.

Check out what our partners have said about us and how they evaluated our cooperation.

Go to about us page