



Two years in **production** with **Kubernetes**

an experience report

**lunar
way**[®]

Container Days 2019, Hamburg
Kasper Nissen (@phenex)



Two years in **production** with **Kubernetes**

an experience report

**lunar
way**[®]

Container Days 2019, Hamburg
Kasper Nissen (@phenex)

\$ whoami

Kasper Nissen (@phennex)

- Site Reliability Engineer / Cloud Architect @lunarway
- Cloud Native Computing Foundation Ambassador
- #CloudNativeAarhus #CloudNativeNordics
- kubecloud.io




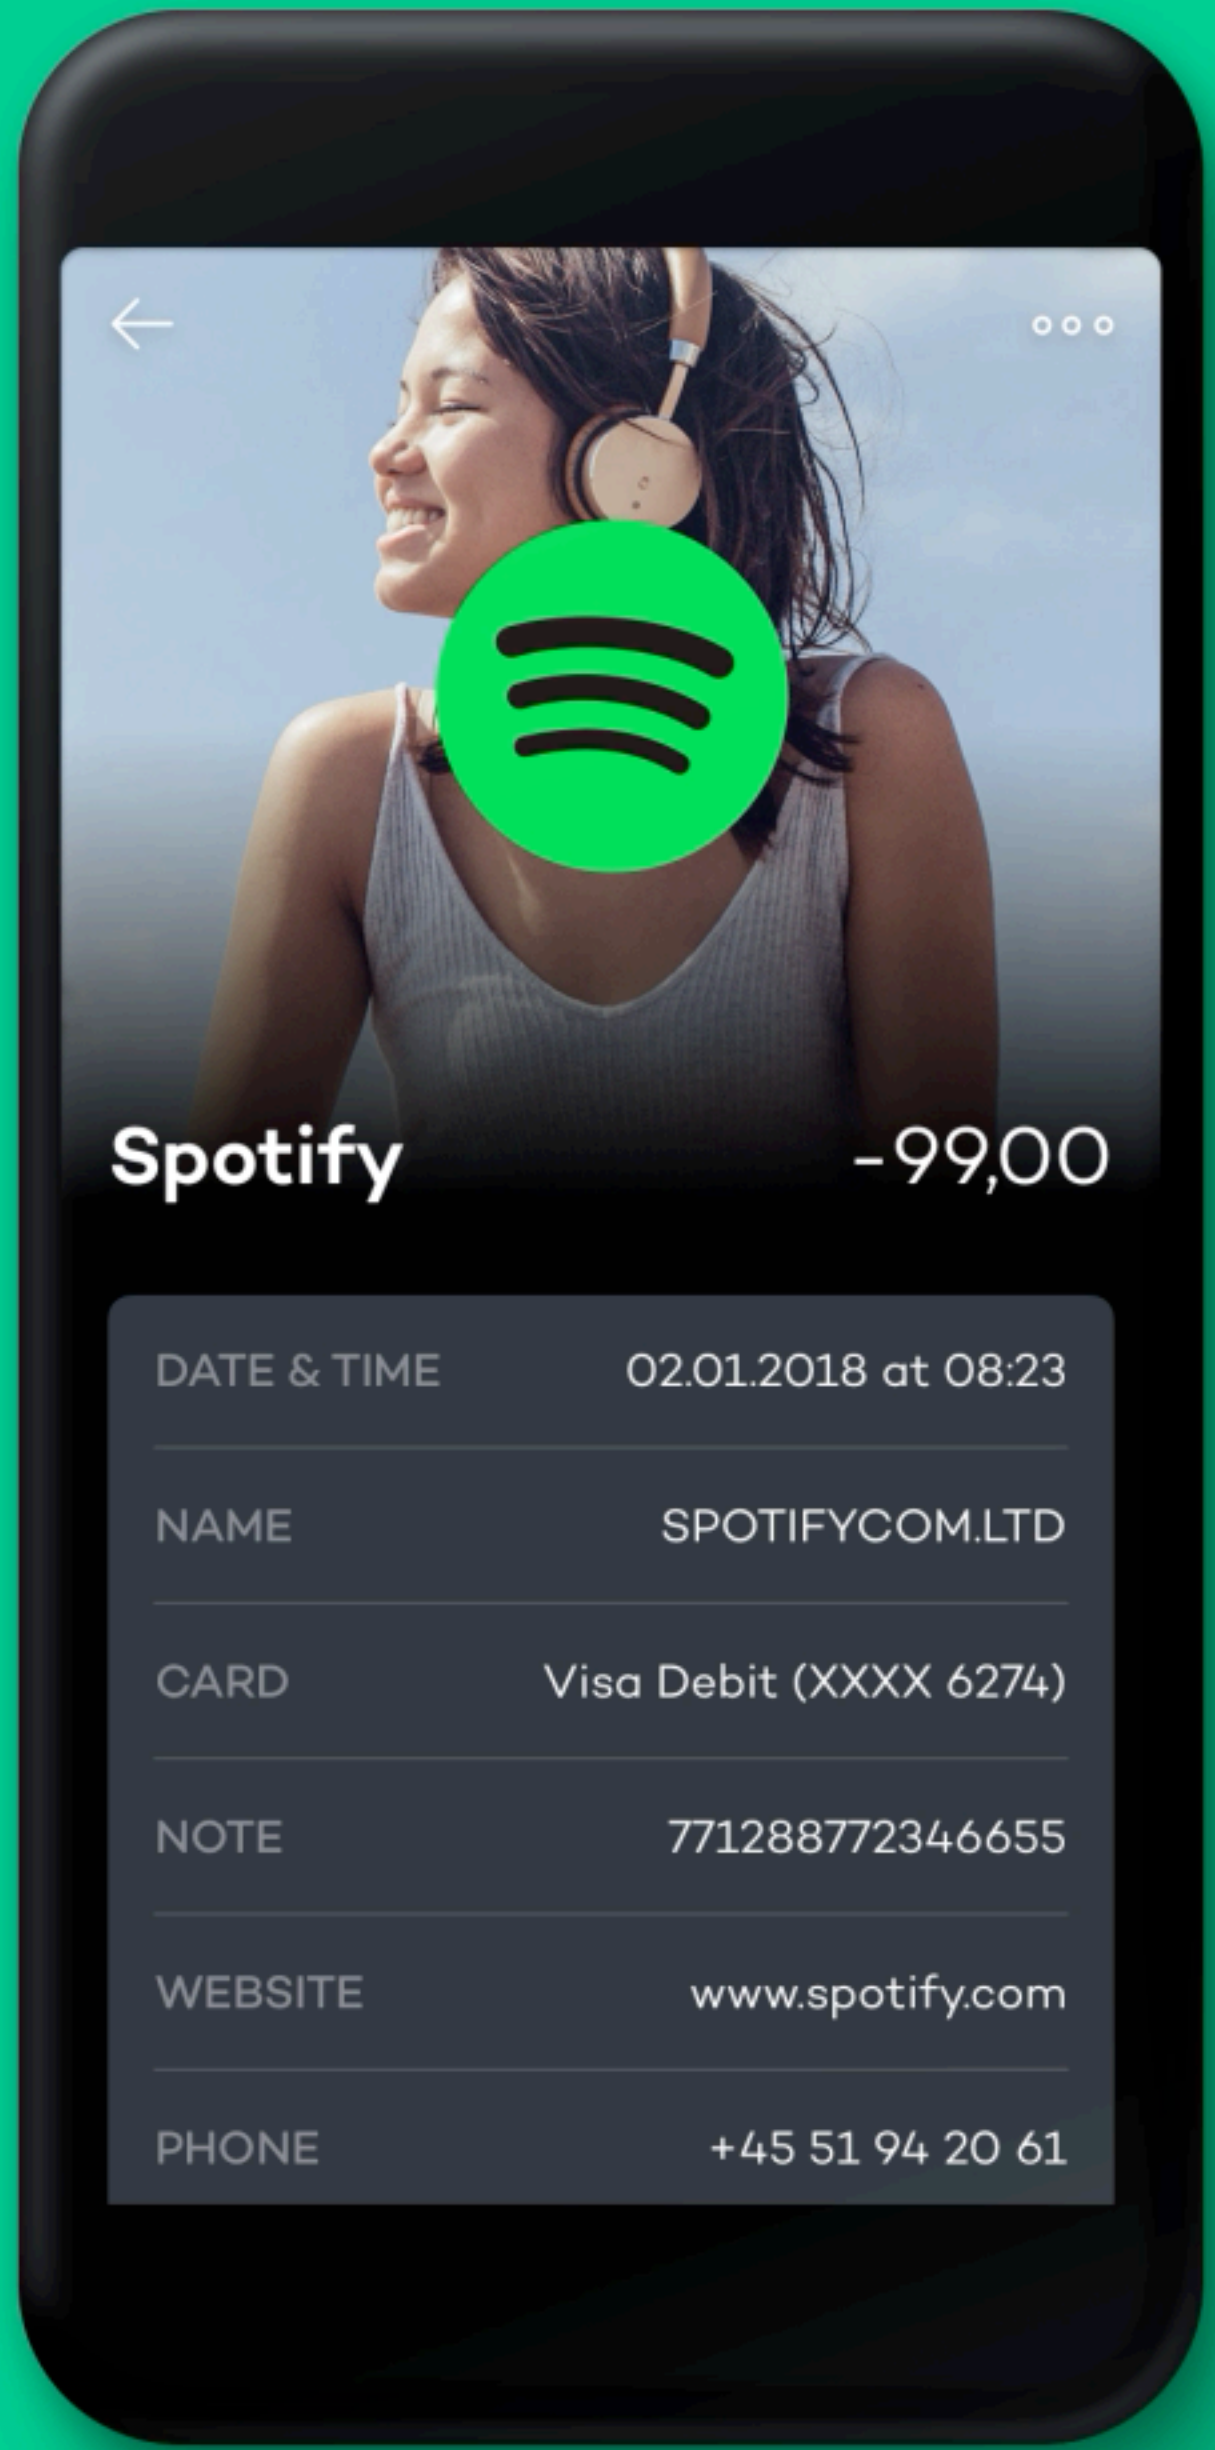
kubecloud.io

Foto credits: Lars Kruse, Aarhus University
Credits: [Pervasive Systems group](#), Section of Electrical and Computer Engineering, Department of Engineering, Aarhus University





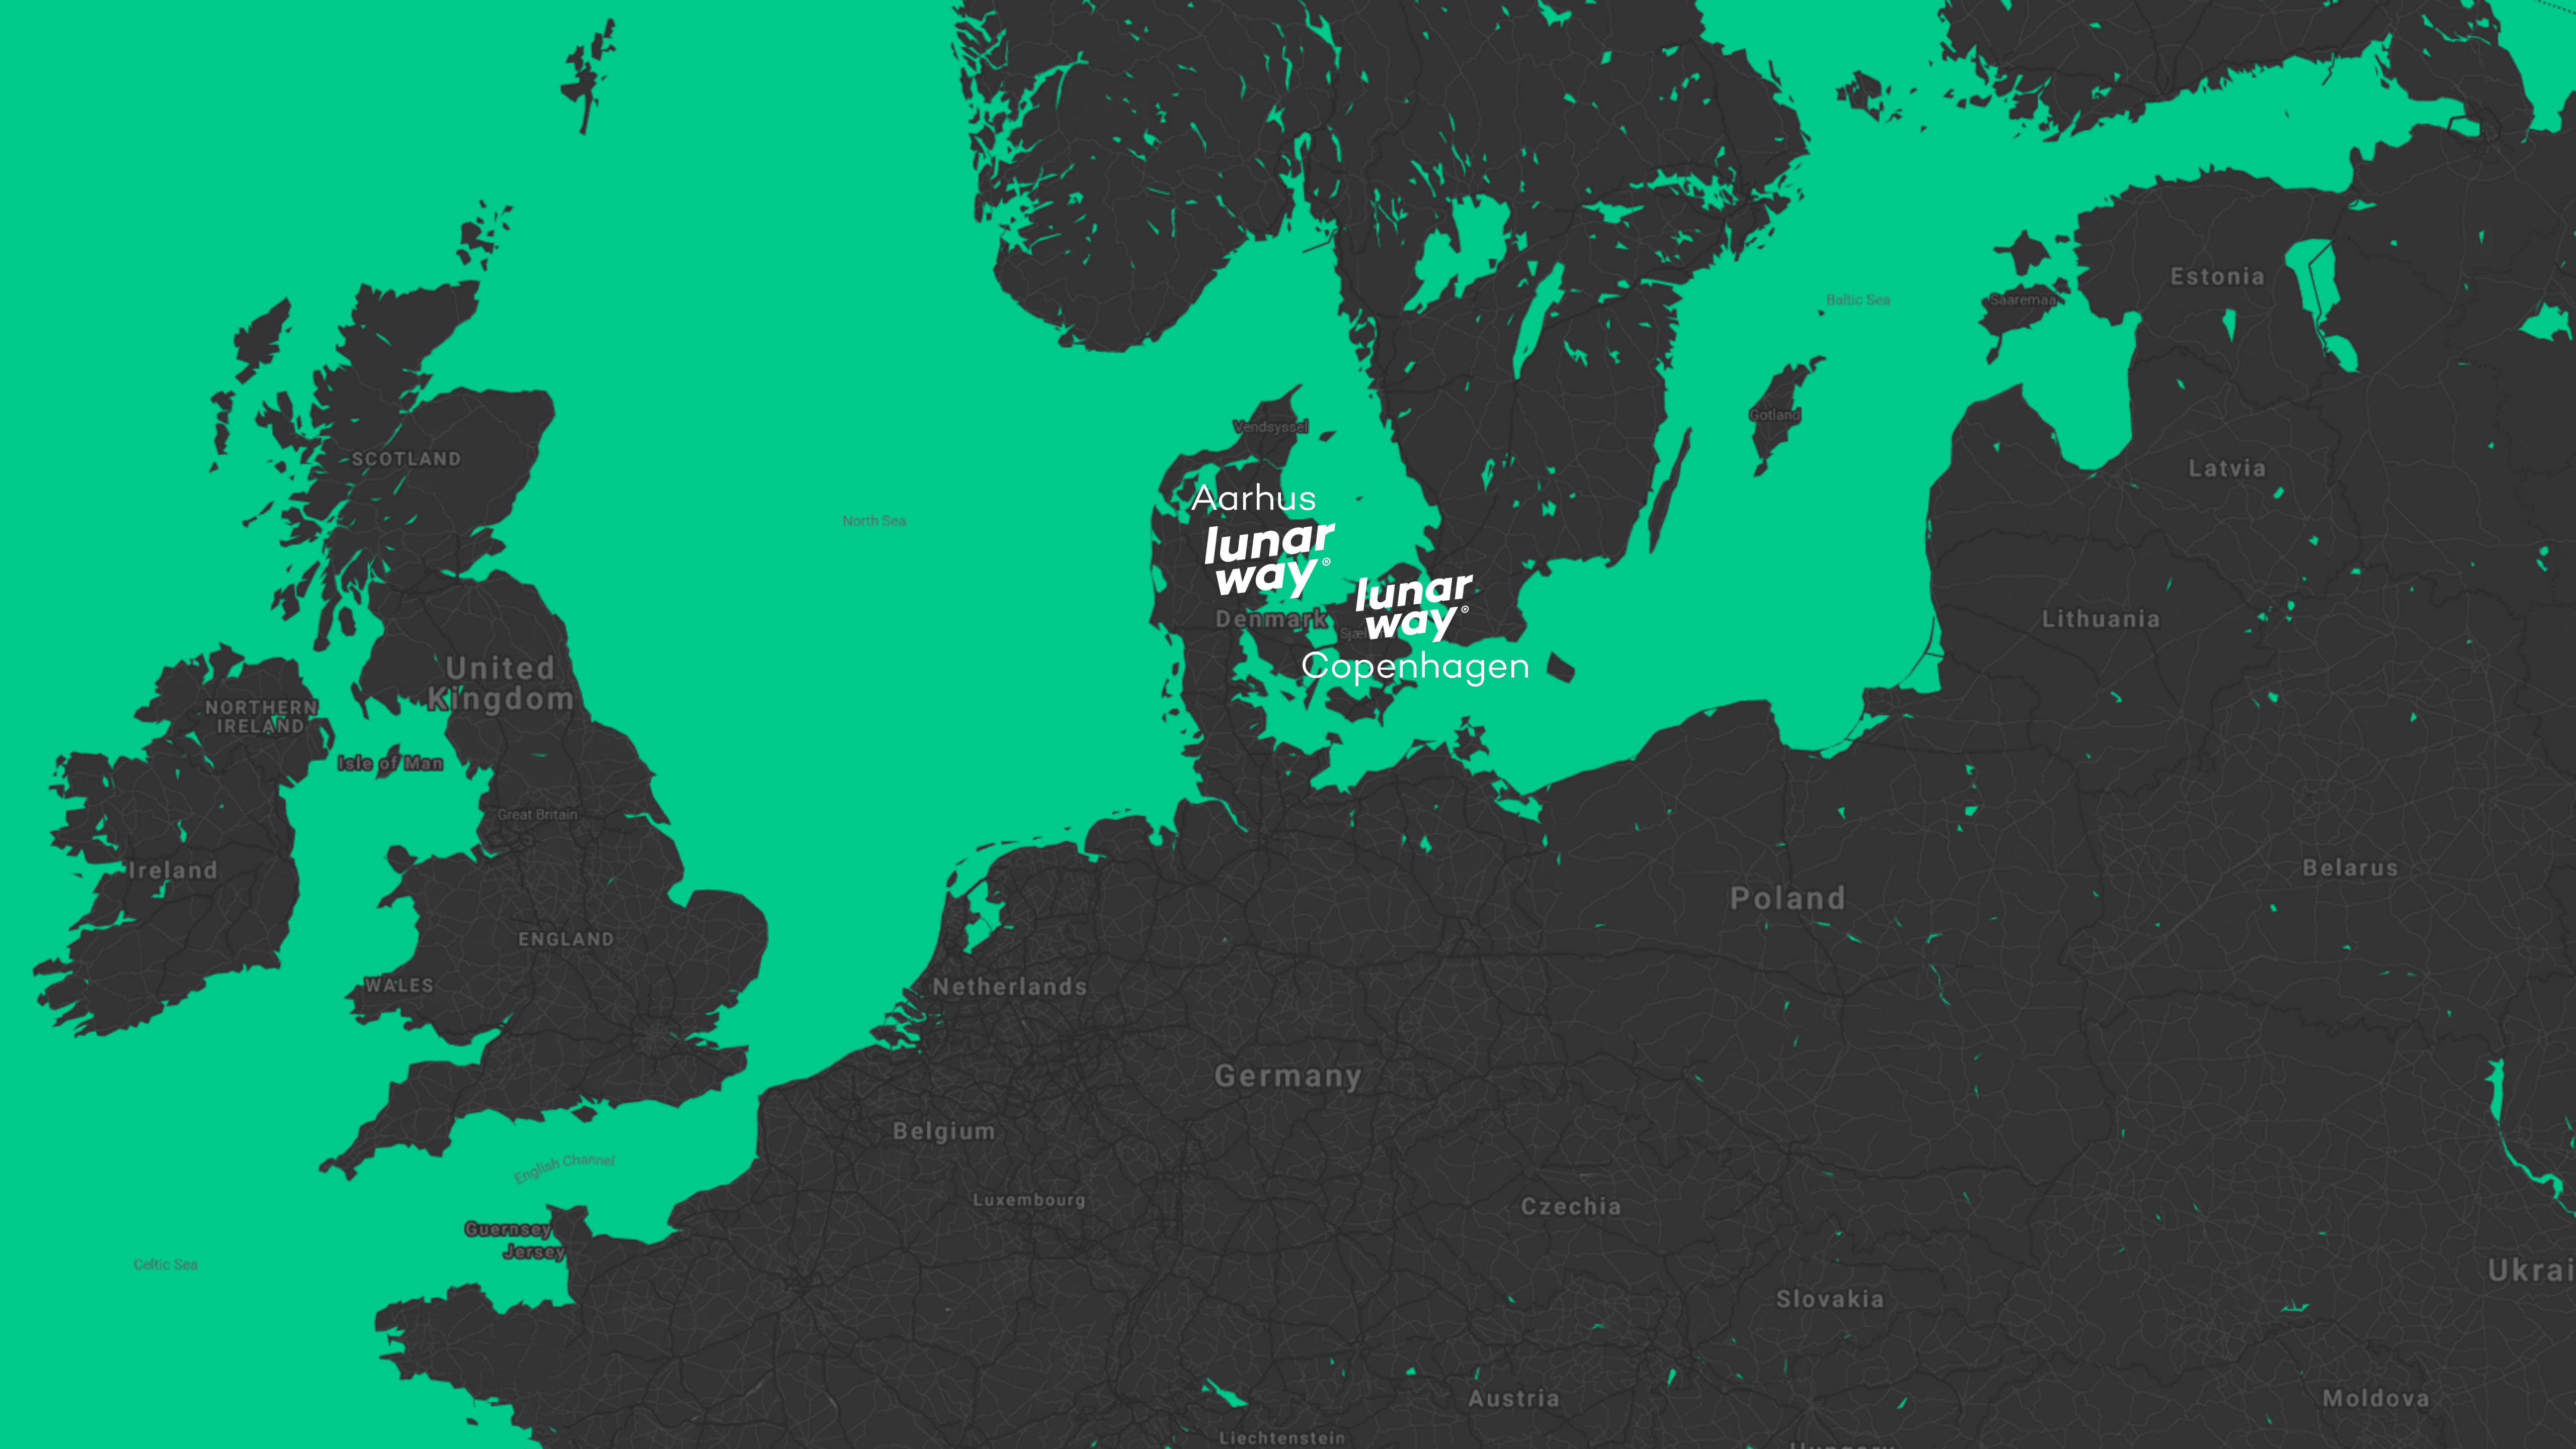
What is *lunar*
way[®] ?



Spotify

-99,00

DATE & TIME	02.01.2018 at 08:23
NAME	SPOTIFYCOM.LTD
CARD	Visa Debit (XXXX 6274)
NOTE	771288772346655
WEBSITE	www.spotify.com
PHONE	+45 51 94 20 61



Aarhus
lunar way

lunar way

Copenhagen

SCOTLAND

North Sea

Vendsyssel

Denmark

Sjælland

Baltic Sea

Gotland

Saaremaa

Estonia

Latvia

Lithuania

United Kingdom

NORTHERN IRELAND

Isle of Man

Great Britain

Ireland

ENGLAND

WALES

Netherlands

Poland

Belarus

Germany

Belgium

Luxembourg

Czechia

Slovakia

Ukraine

Celtic Sea

English Channel

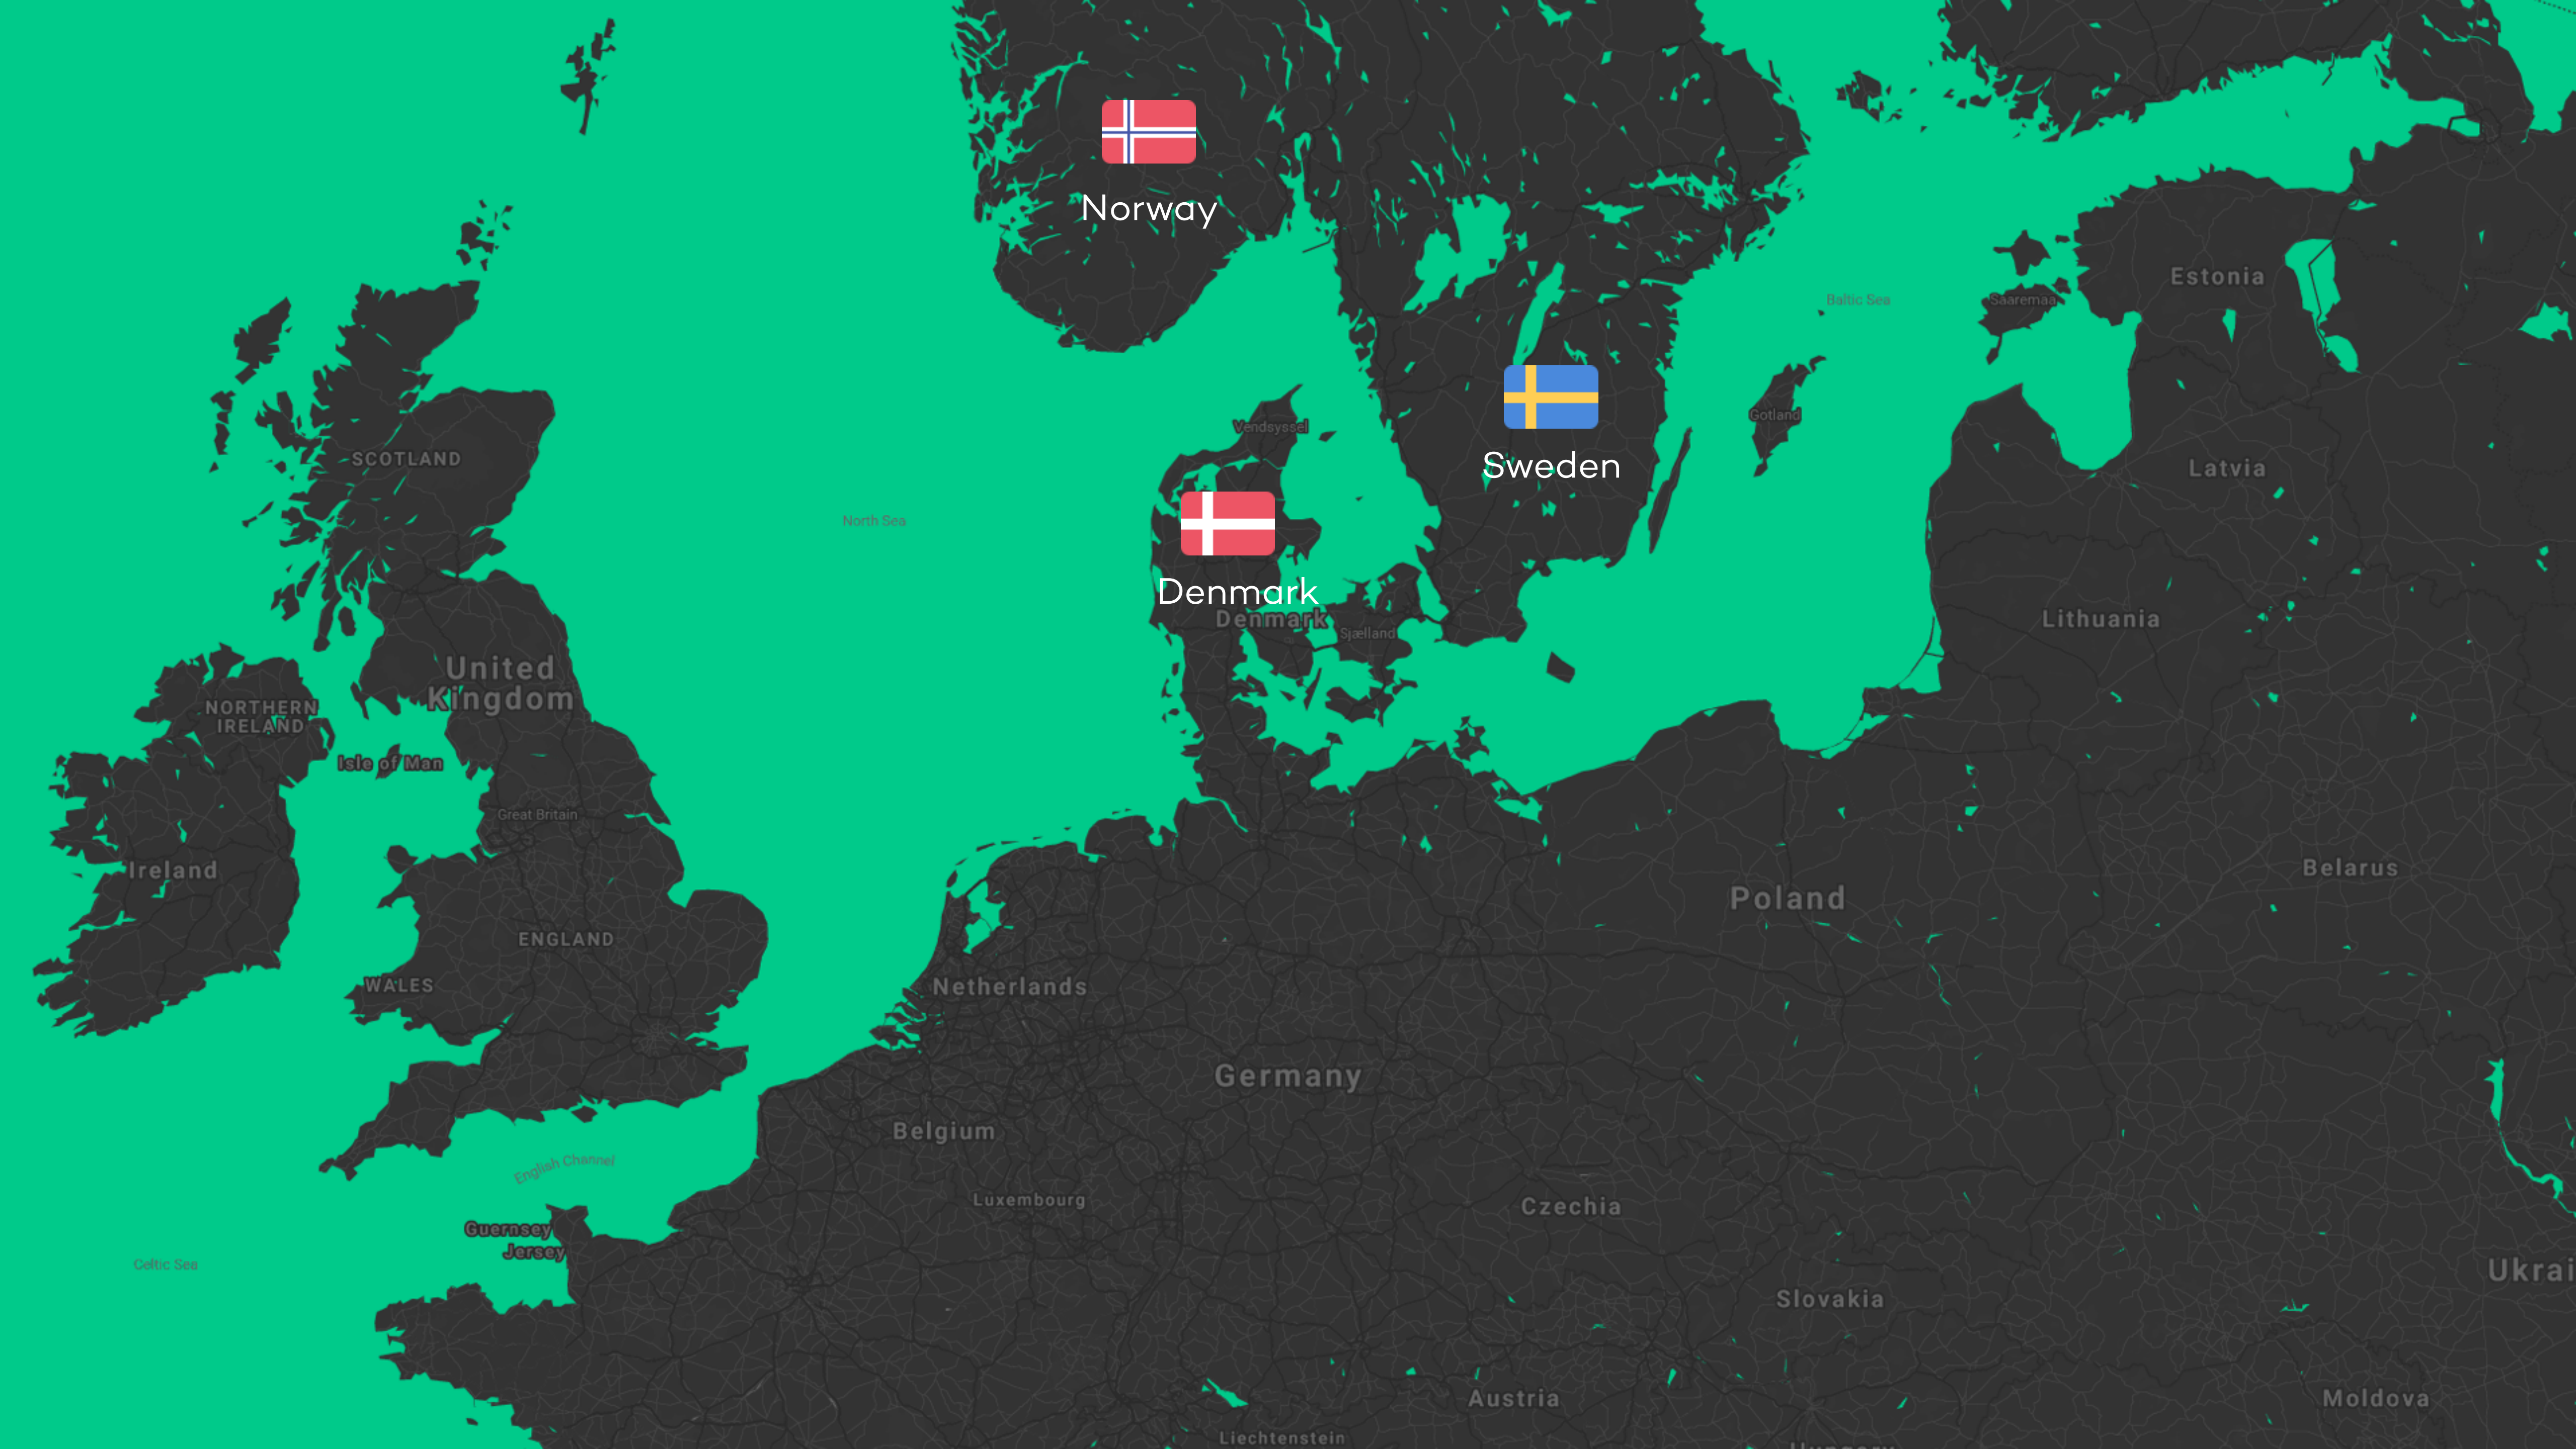
Guernsey
Jersey

Liechtenstein

Austria

Hungary

Moldova



Norway



Sweden



Denmark

SCOTLAND

NORTHERN IRELAND

United Kingdom

Ireland

Isle of Man

Great Britain

ENGLAND

WALES

Guernsey
Jersey

North Sea

Vendsyssel

Denmark

Sjælland

Baltic Sea

Saaremaa

Gotland

Estonia

Latvia

Lithuania

Poland

Belarus

Netherlands

Germany

Belgium

Luxembourg

Czechia

Slovakia

Austria

Liechtenstein

Hungary

Moldova

Ukraine



in numbers

85K+

Users

1M+

Transactions
(per month)

80+

Microservices

90+

Employees

100M+

€ Volume
(per month)

3

k8s clusters

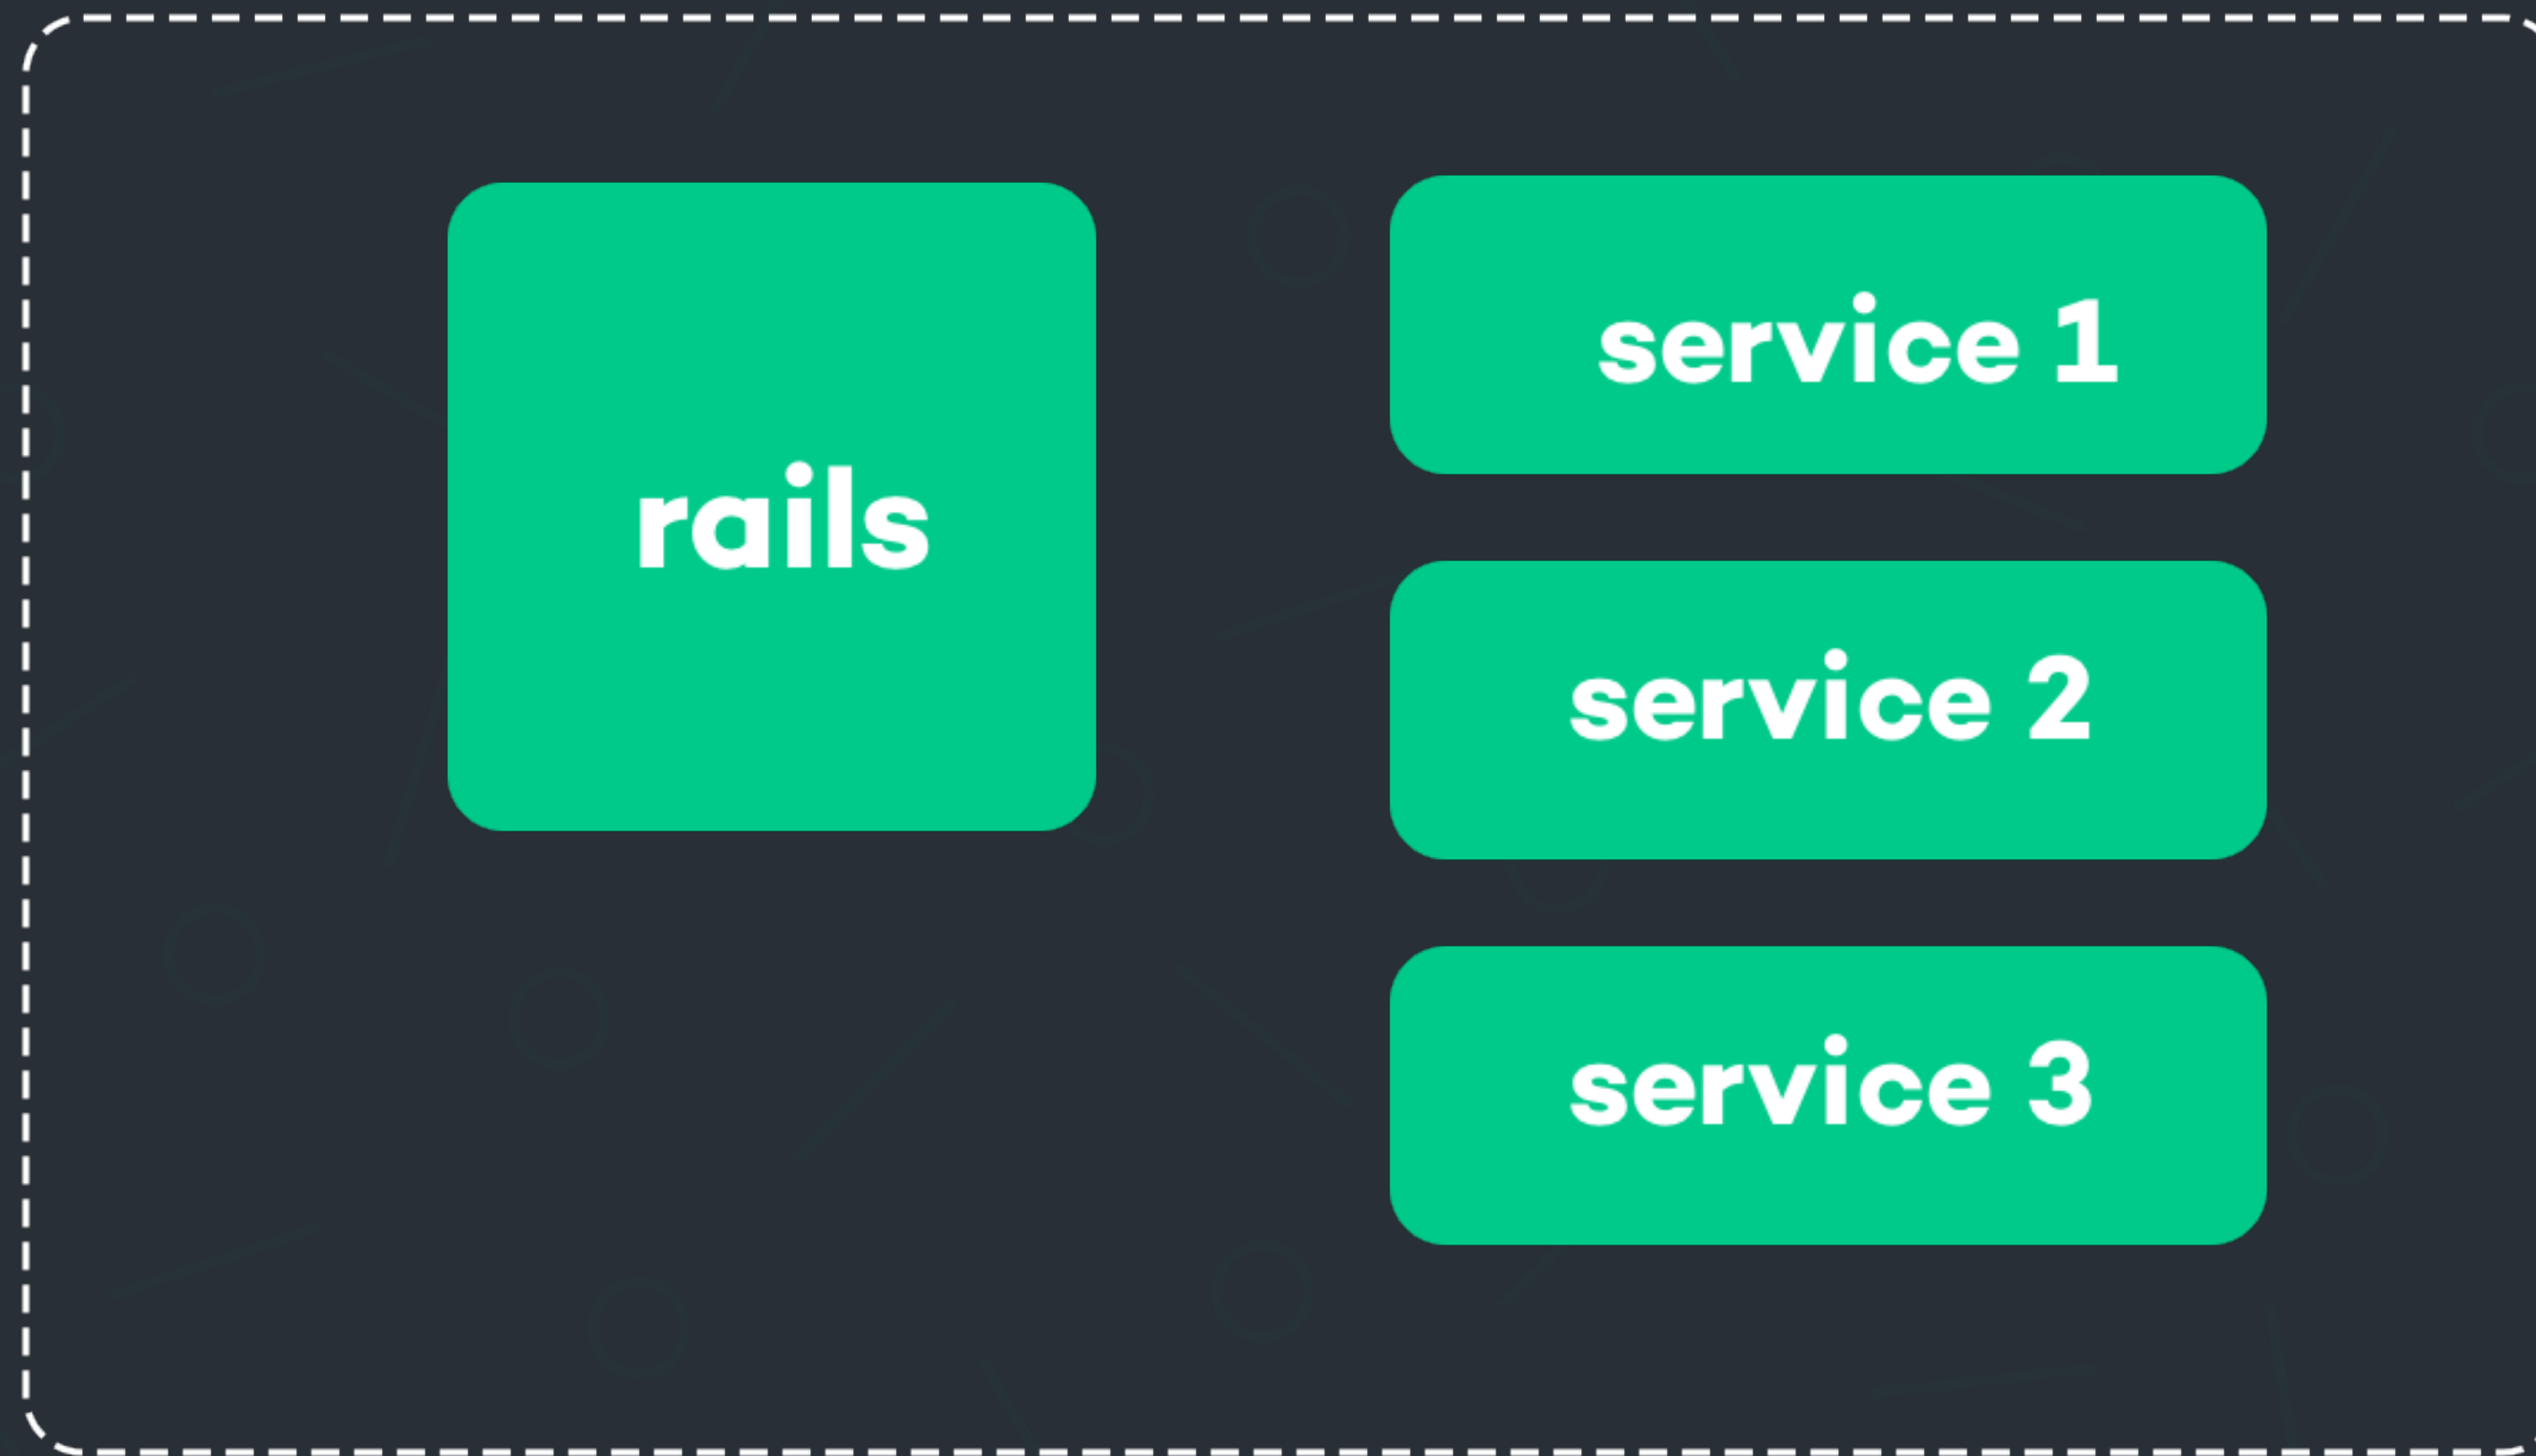
**So, where did we start our
cloud native journey?**



DU

Deployable Unit (DU)





DU

Deployable Unit (DU)

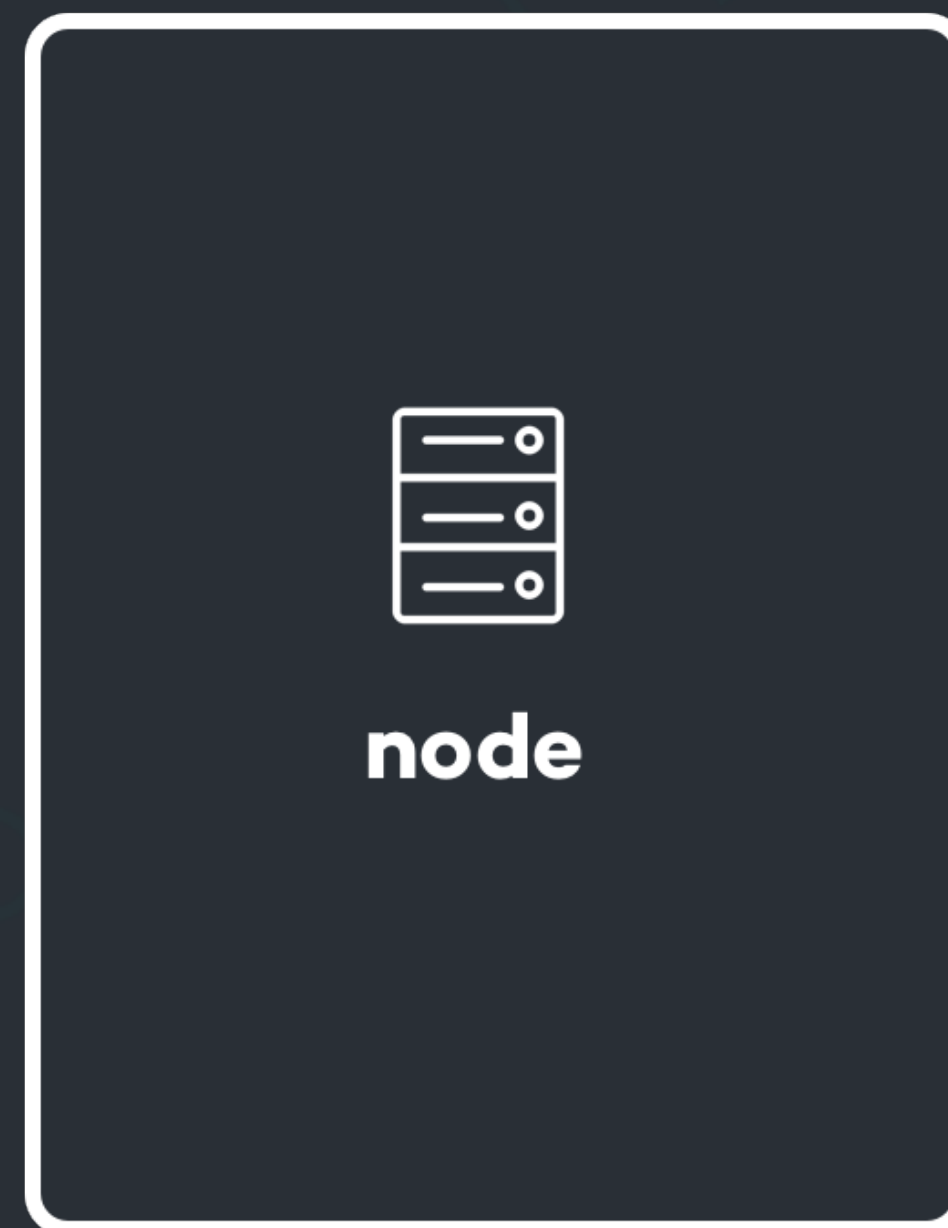

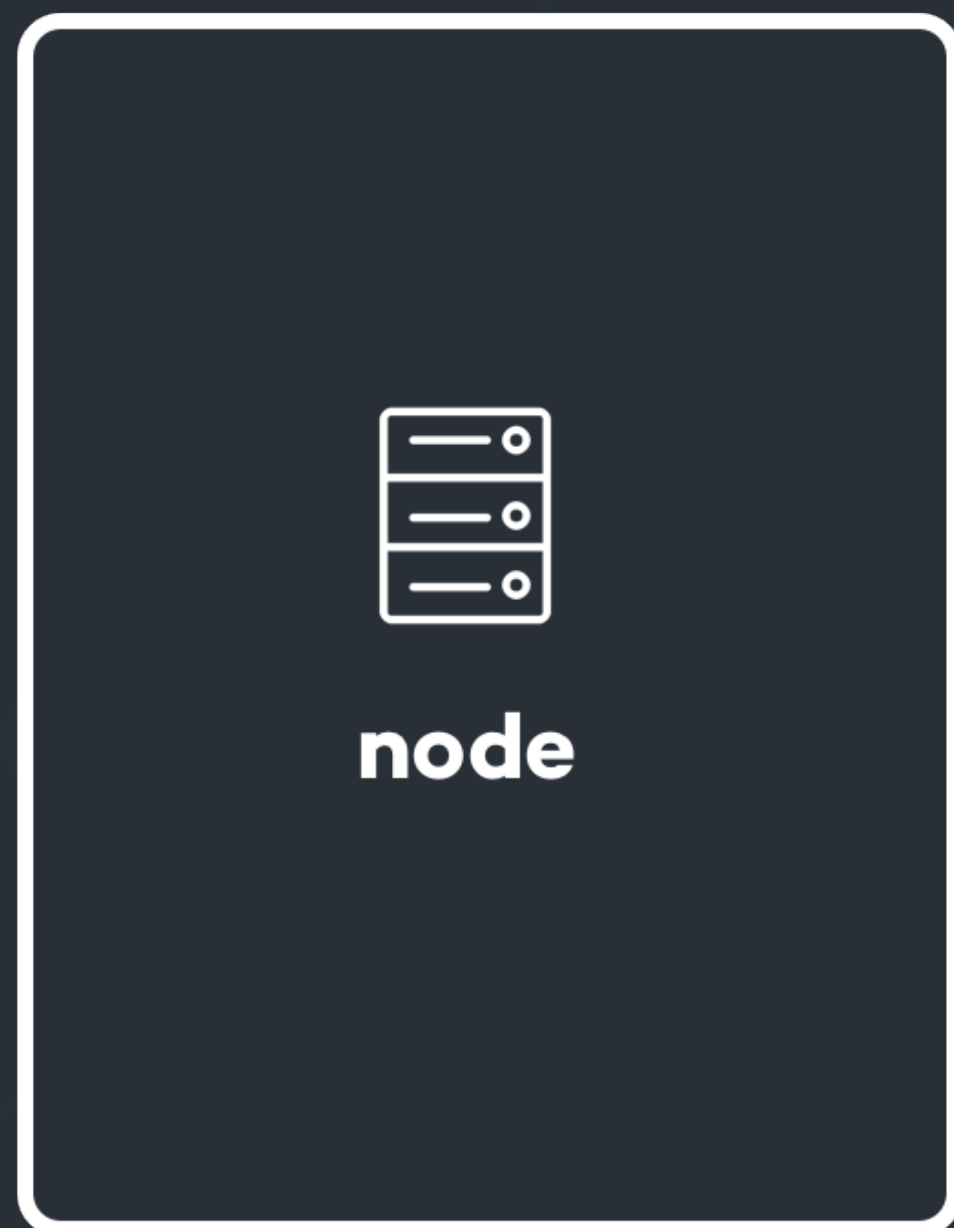



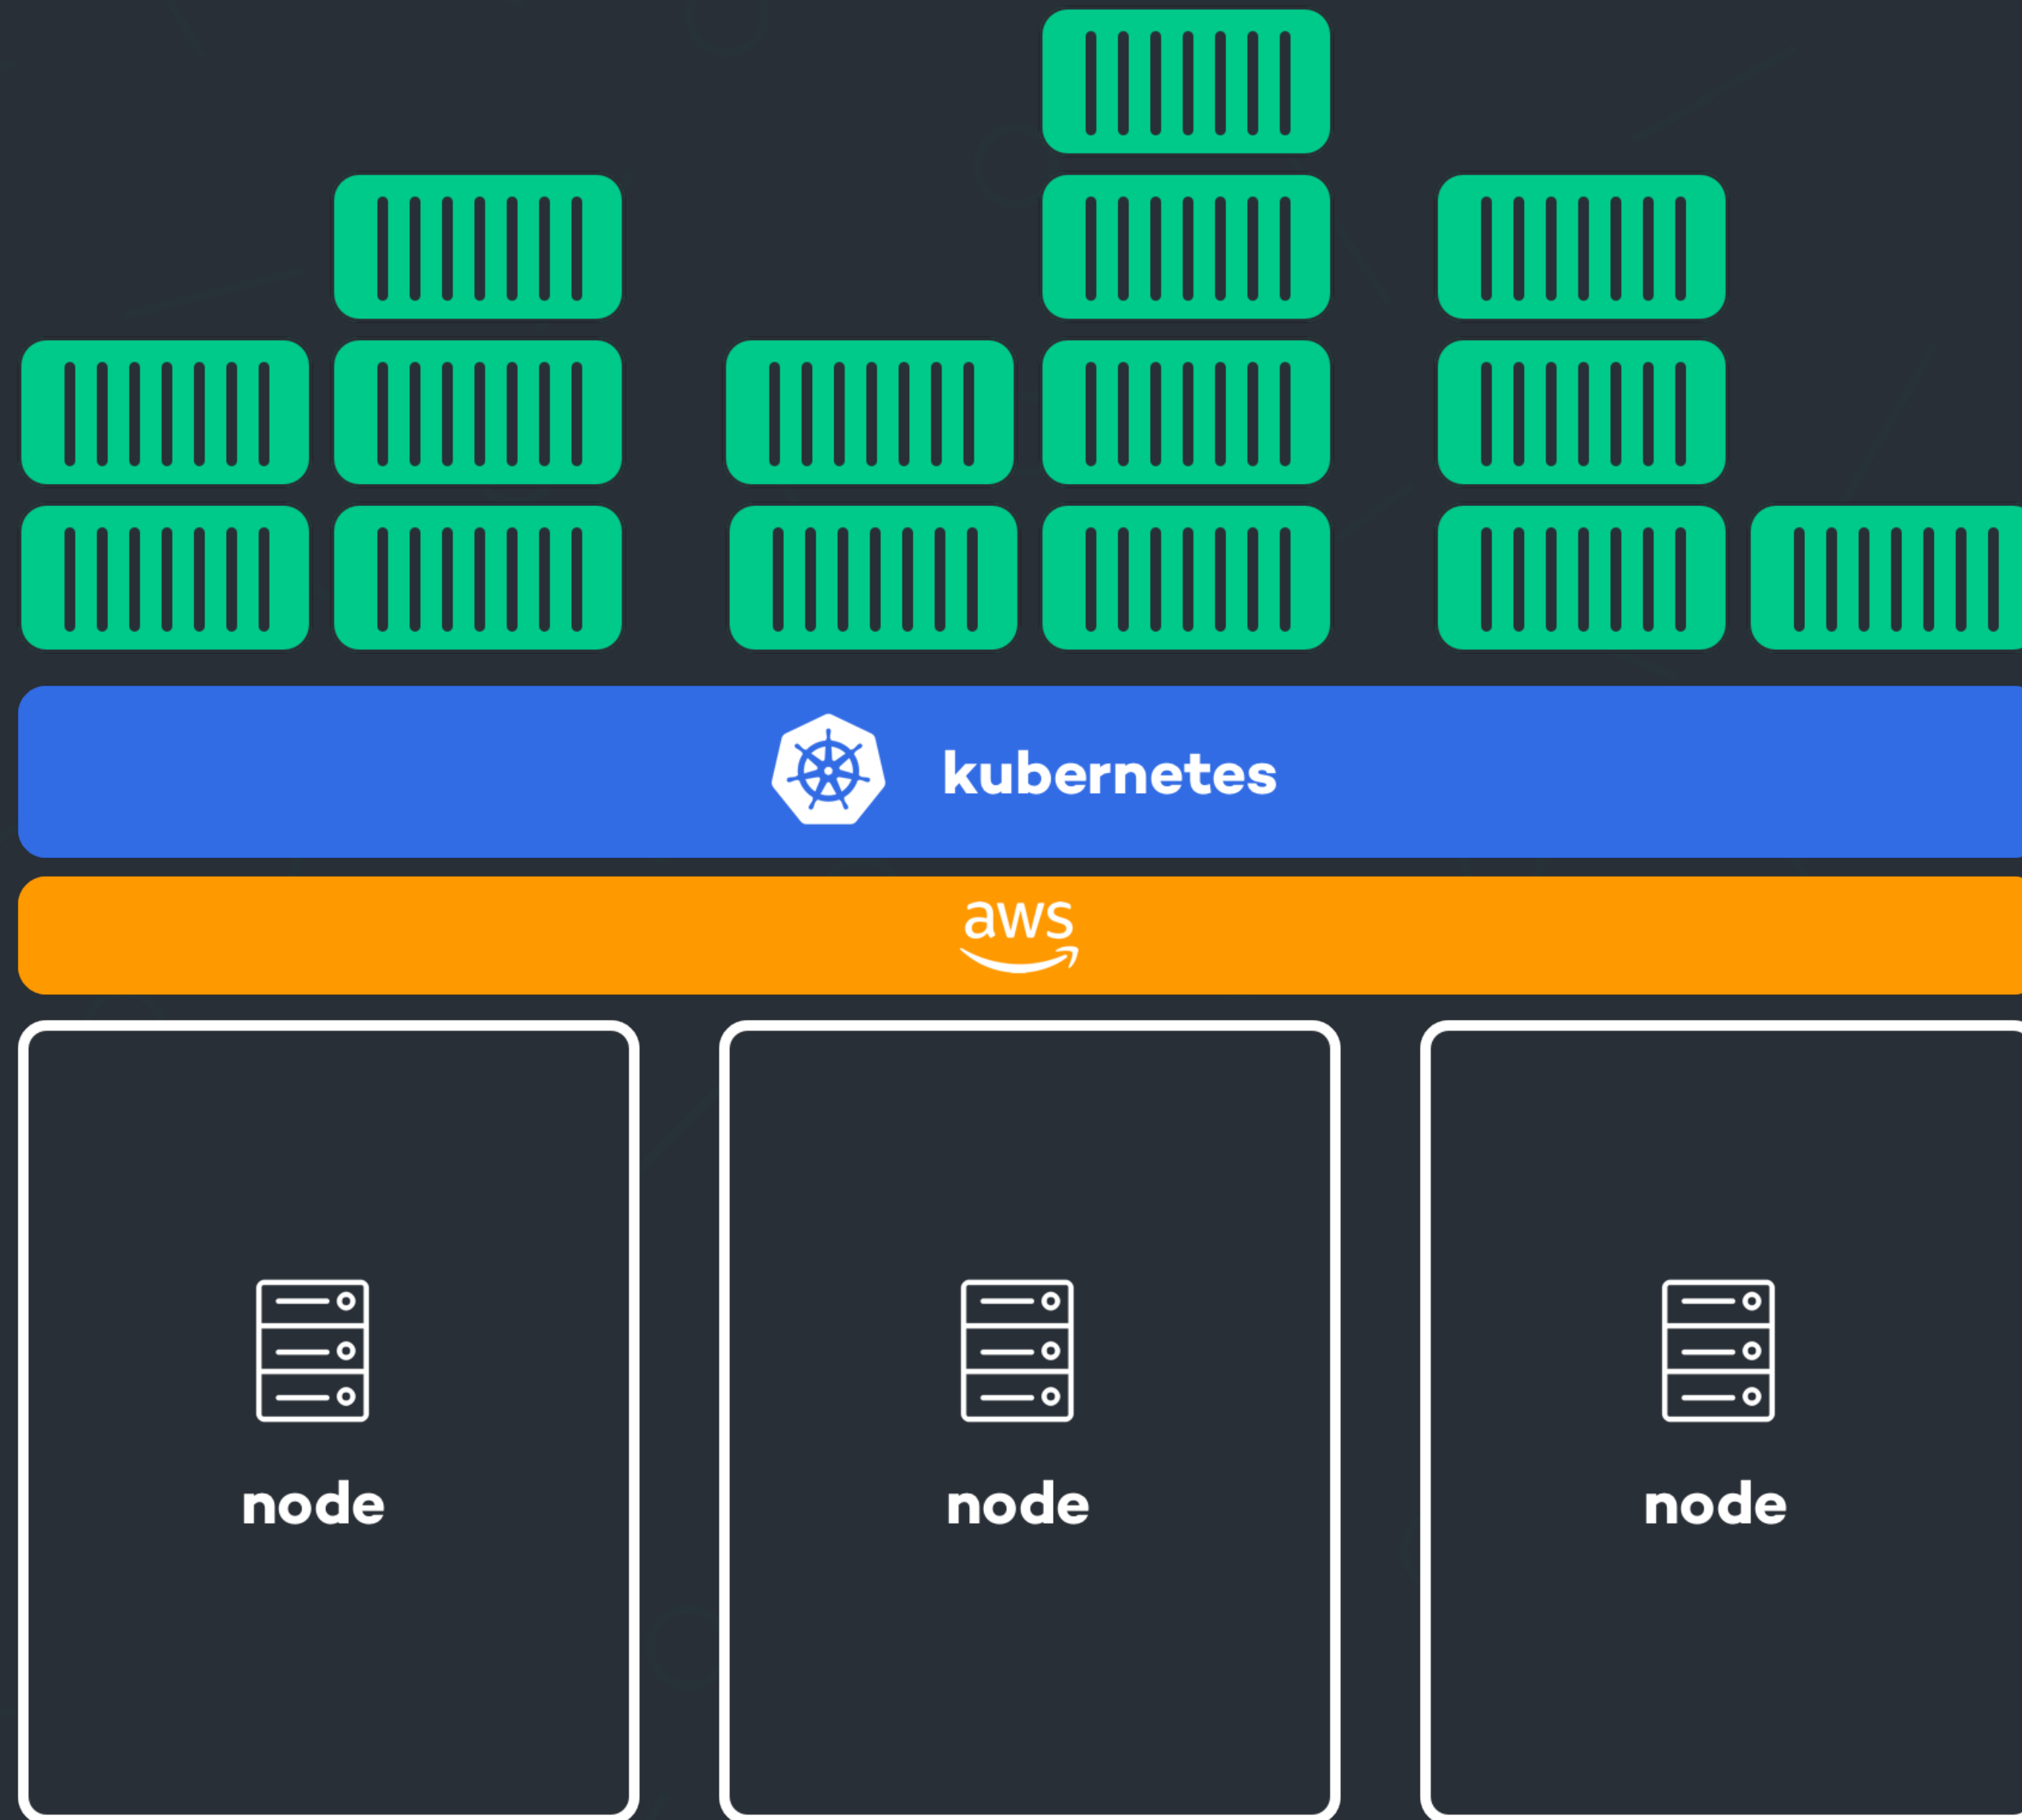
application-oriented





machine-oriented

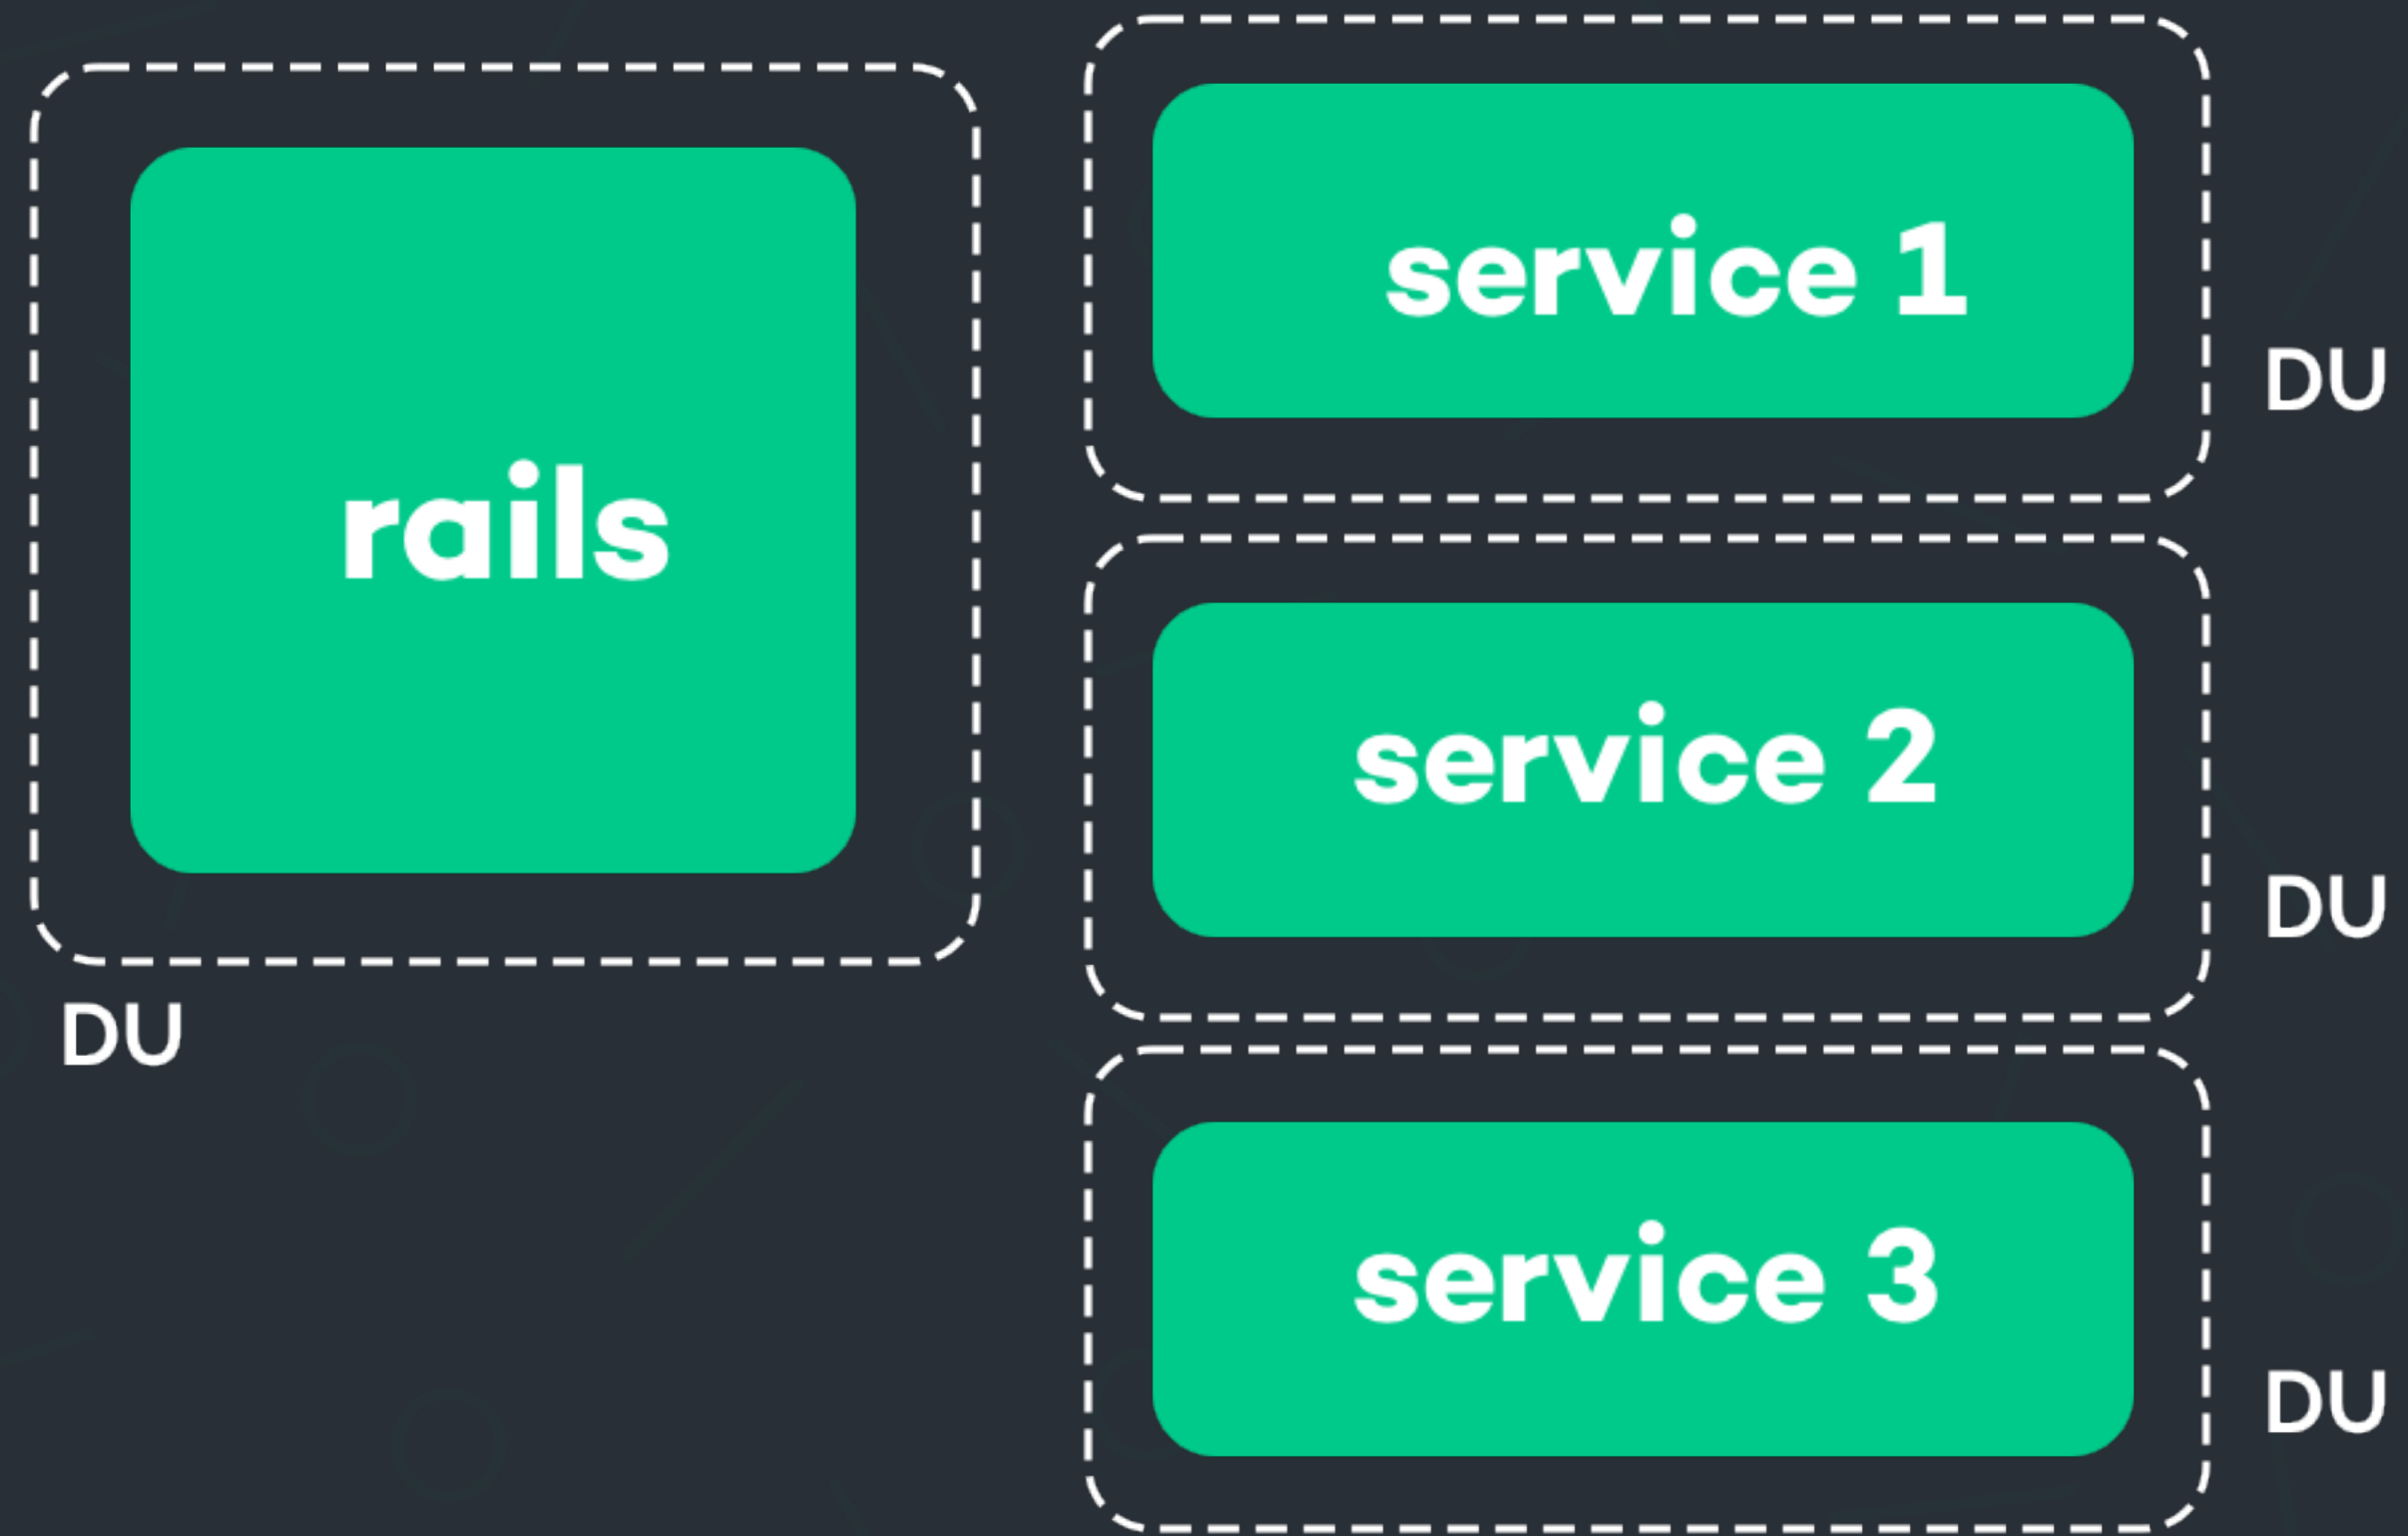
“Unlock velocity by shifting the focus of deployment from being machine-oriented to be application-oriented”

Inspired by “Borg, Omega, and Kubernetes” by Brendan Burns et. al.



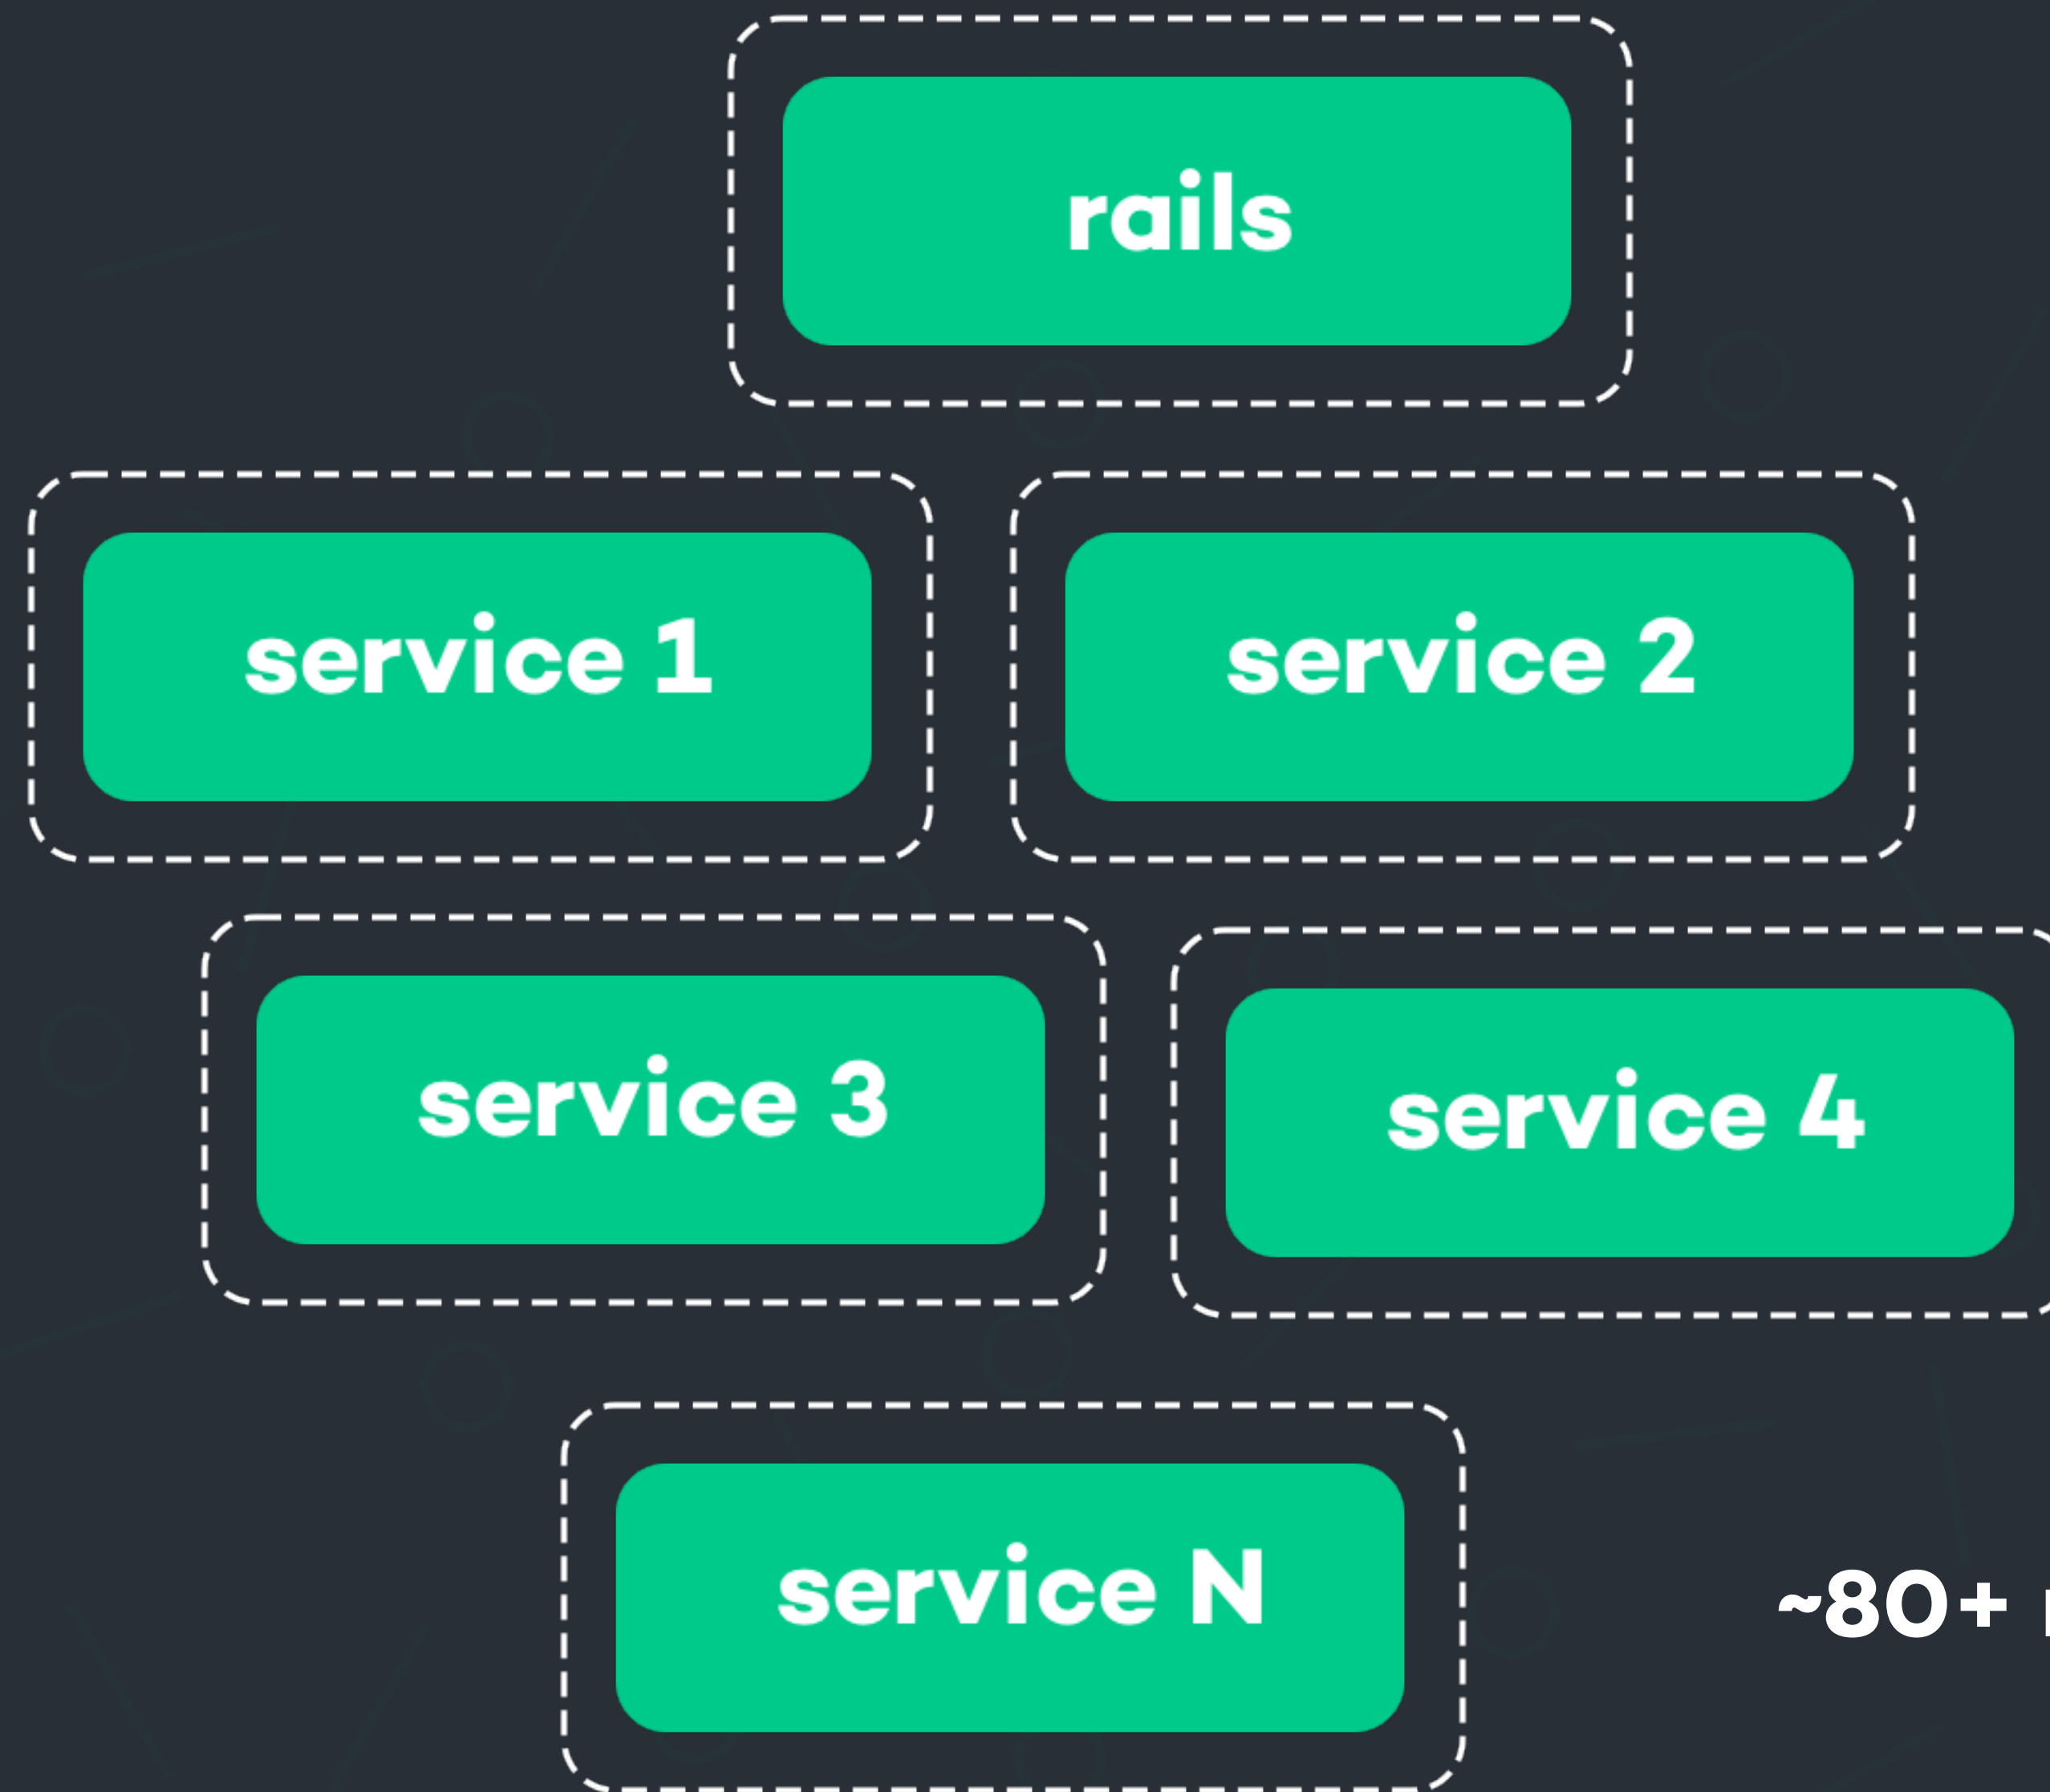






Deployable Unit (DU)





~80+ microservices

”Kubernetes is the **foundation
on which we **build**”**

**“Kubernetes is a platform
for creating platforms”**

Bryan Liles, CloudNativeCon+KubeCon Europe 2019

“Kubernetes is not a destination. It’s part of the journey”

Bryan Liles, CloudNativeCon+KubeCon Europe 2019

**lunar
way®**

**lunar
way®**

Today's topics

1. Finding the right abstraction
2. Release Management
3. Operations/Maintenance

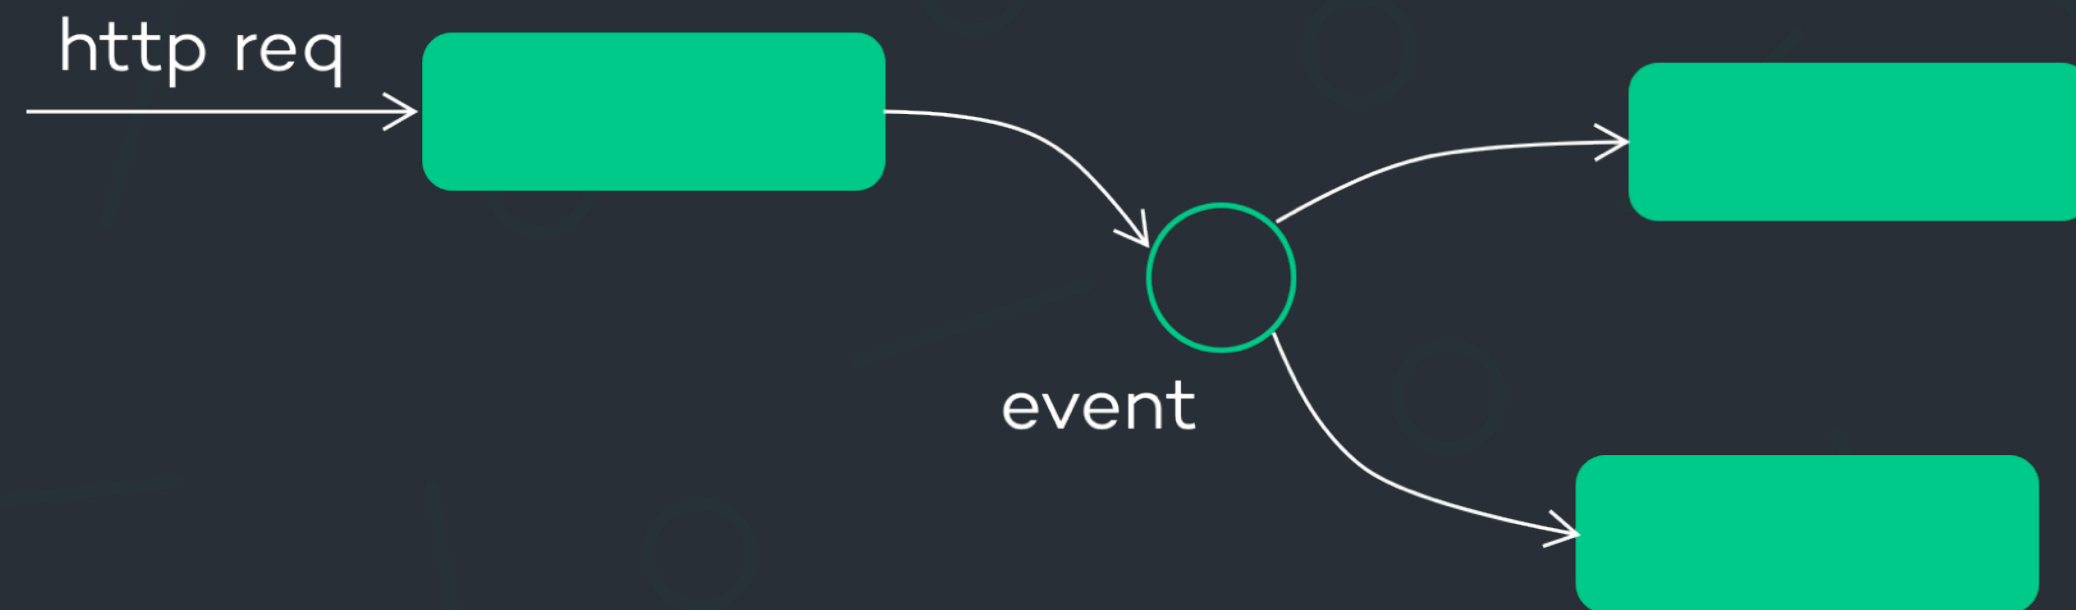
“Good decisions come from experience, and experience comes from bad decisions.”

Anonymous quote from the Internet


GOTO Copenhagen 2017



Shared Core Dependencies



Asynchronous first



Self-contained



Self-contained

Scripts

Kubernetes resources

Dockerfiles

Jenkins pipeline



Ah man, I have 8 services, and I have this deadline on a new feature....

You need to upgrade your **Dockerfile** to **node:8.16.0-alpine** to patch security vulnerabilities



**Feature
Squads**



**Platform
team**

**We needed to find the right
abstraction...**


“Developers shouldn’t really care about Dockerfiles and Kubernetes resources”



Introducing shuttle

A CLI for handling shared **build and deploy** tools between many projects **no matter what technologies** the project is using.

<https://github.com/lunarway/shuttle>



Scripts Templates
pipeline scripts


```
shuttle template deployment.yaml
```

```
shuttle run build
```

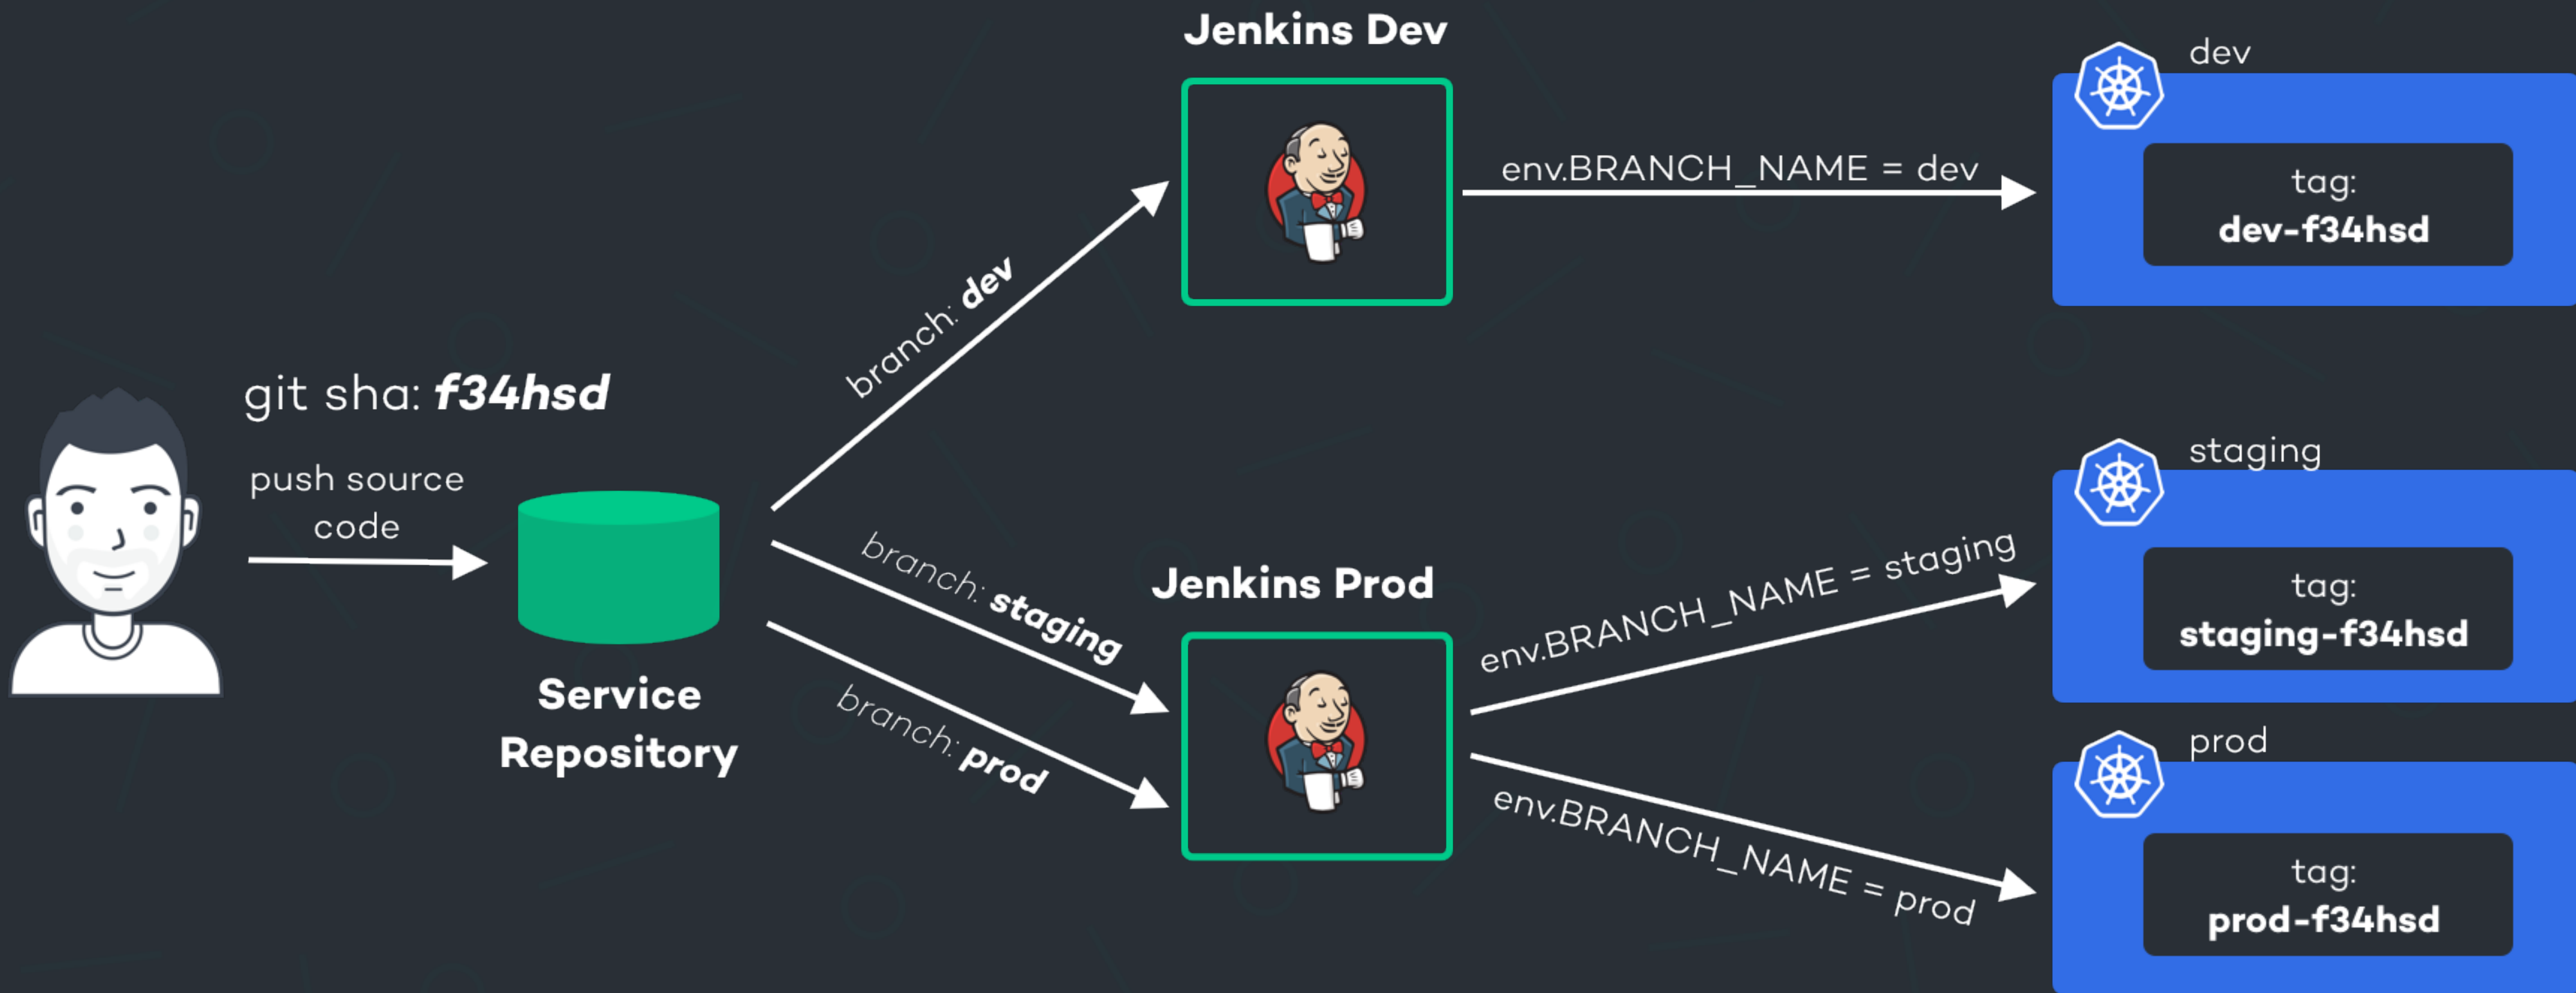
```
shuttle run generate_config
```

```
shuttle run secrets
```

```
shuttle run test
```

**”Centralization of Dockerfiles,
Pipelines, and Kubernetes
resources was the key to
unlock velocity for the
platform team”**

Next order of business. Release Management.



Jenkinsfile
Switch on `env.BRANCH_NAME`

Problems?

Dynamic generation of k8s resources, resulted in **no audit trail**

...

No single source of truth for a **disaster recovery situation**

...

A lot of wasted time because of **duplicate builds**



This guy looks really innocent



**But, do you really trust him to have
the keys to your house?**

**“Continuous Integration (CI)
and
Continuous Deployment (CD)
are two different concerns.”**

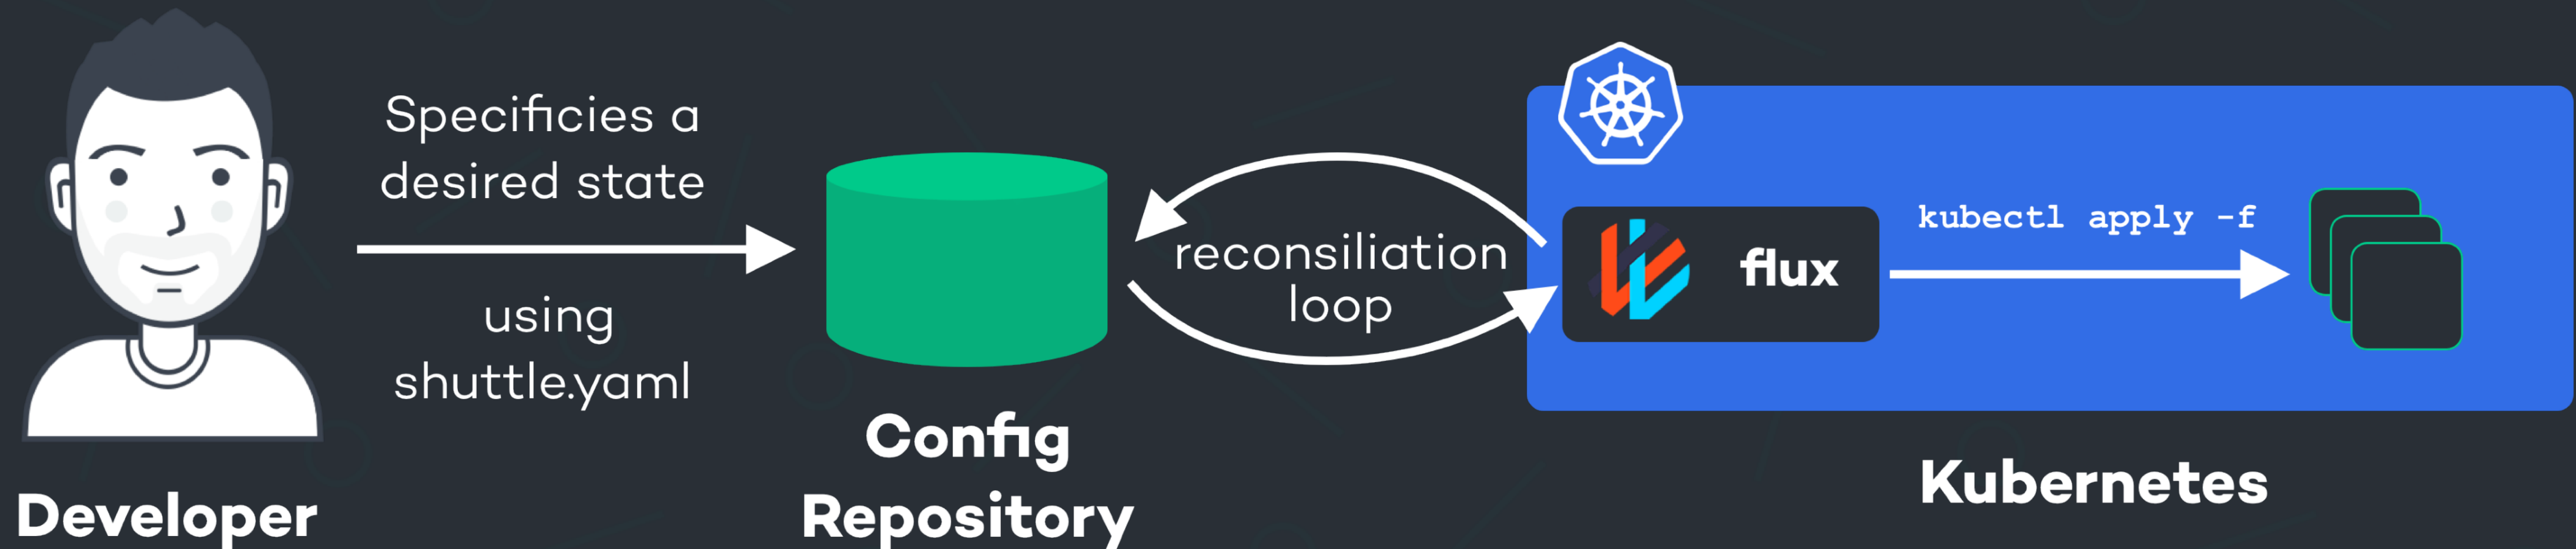


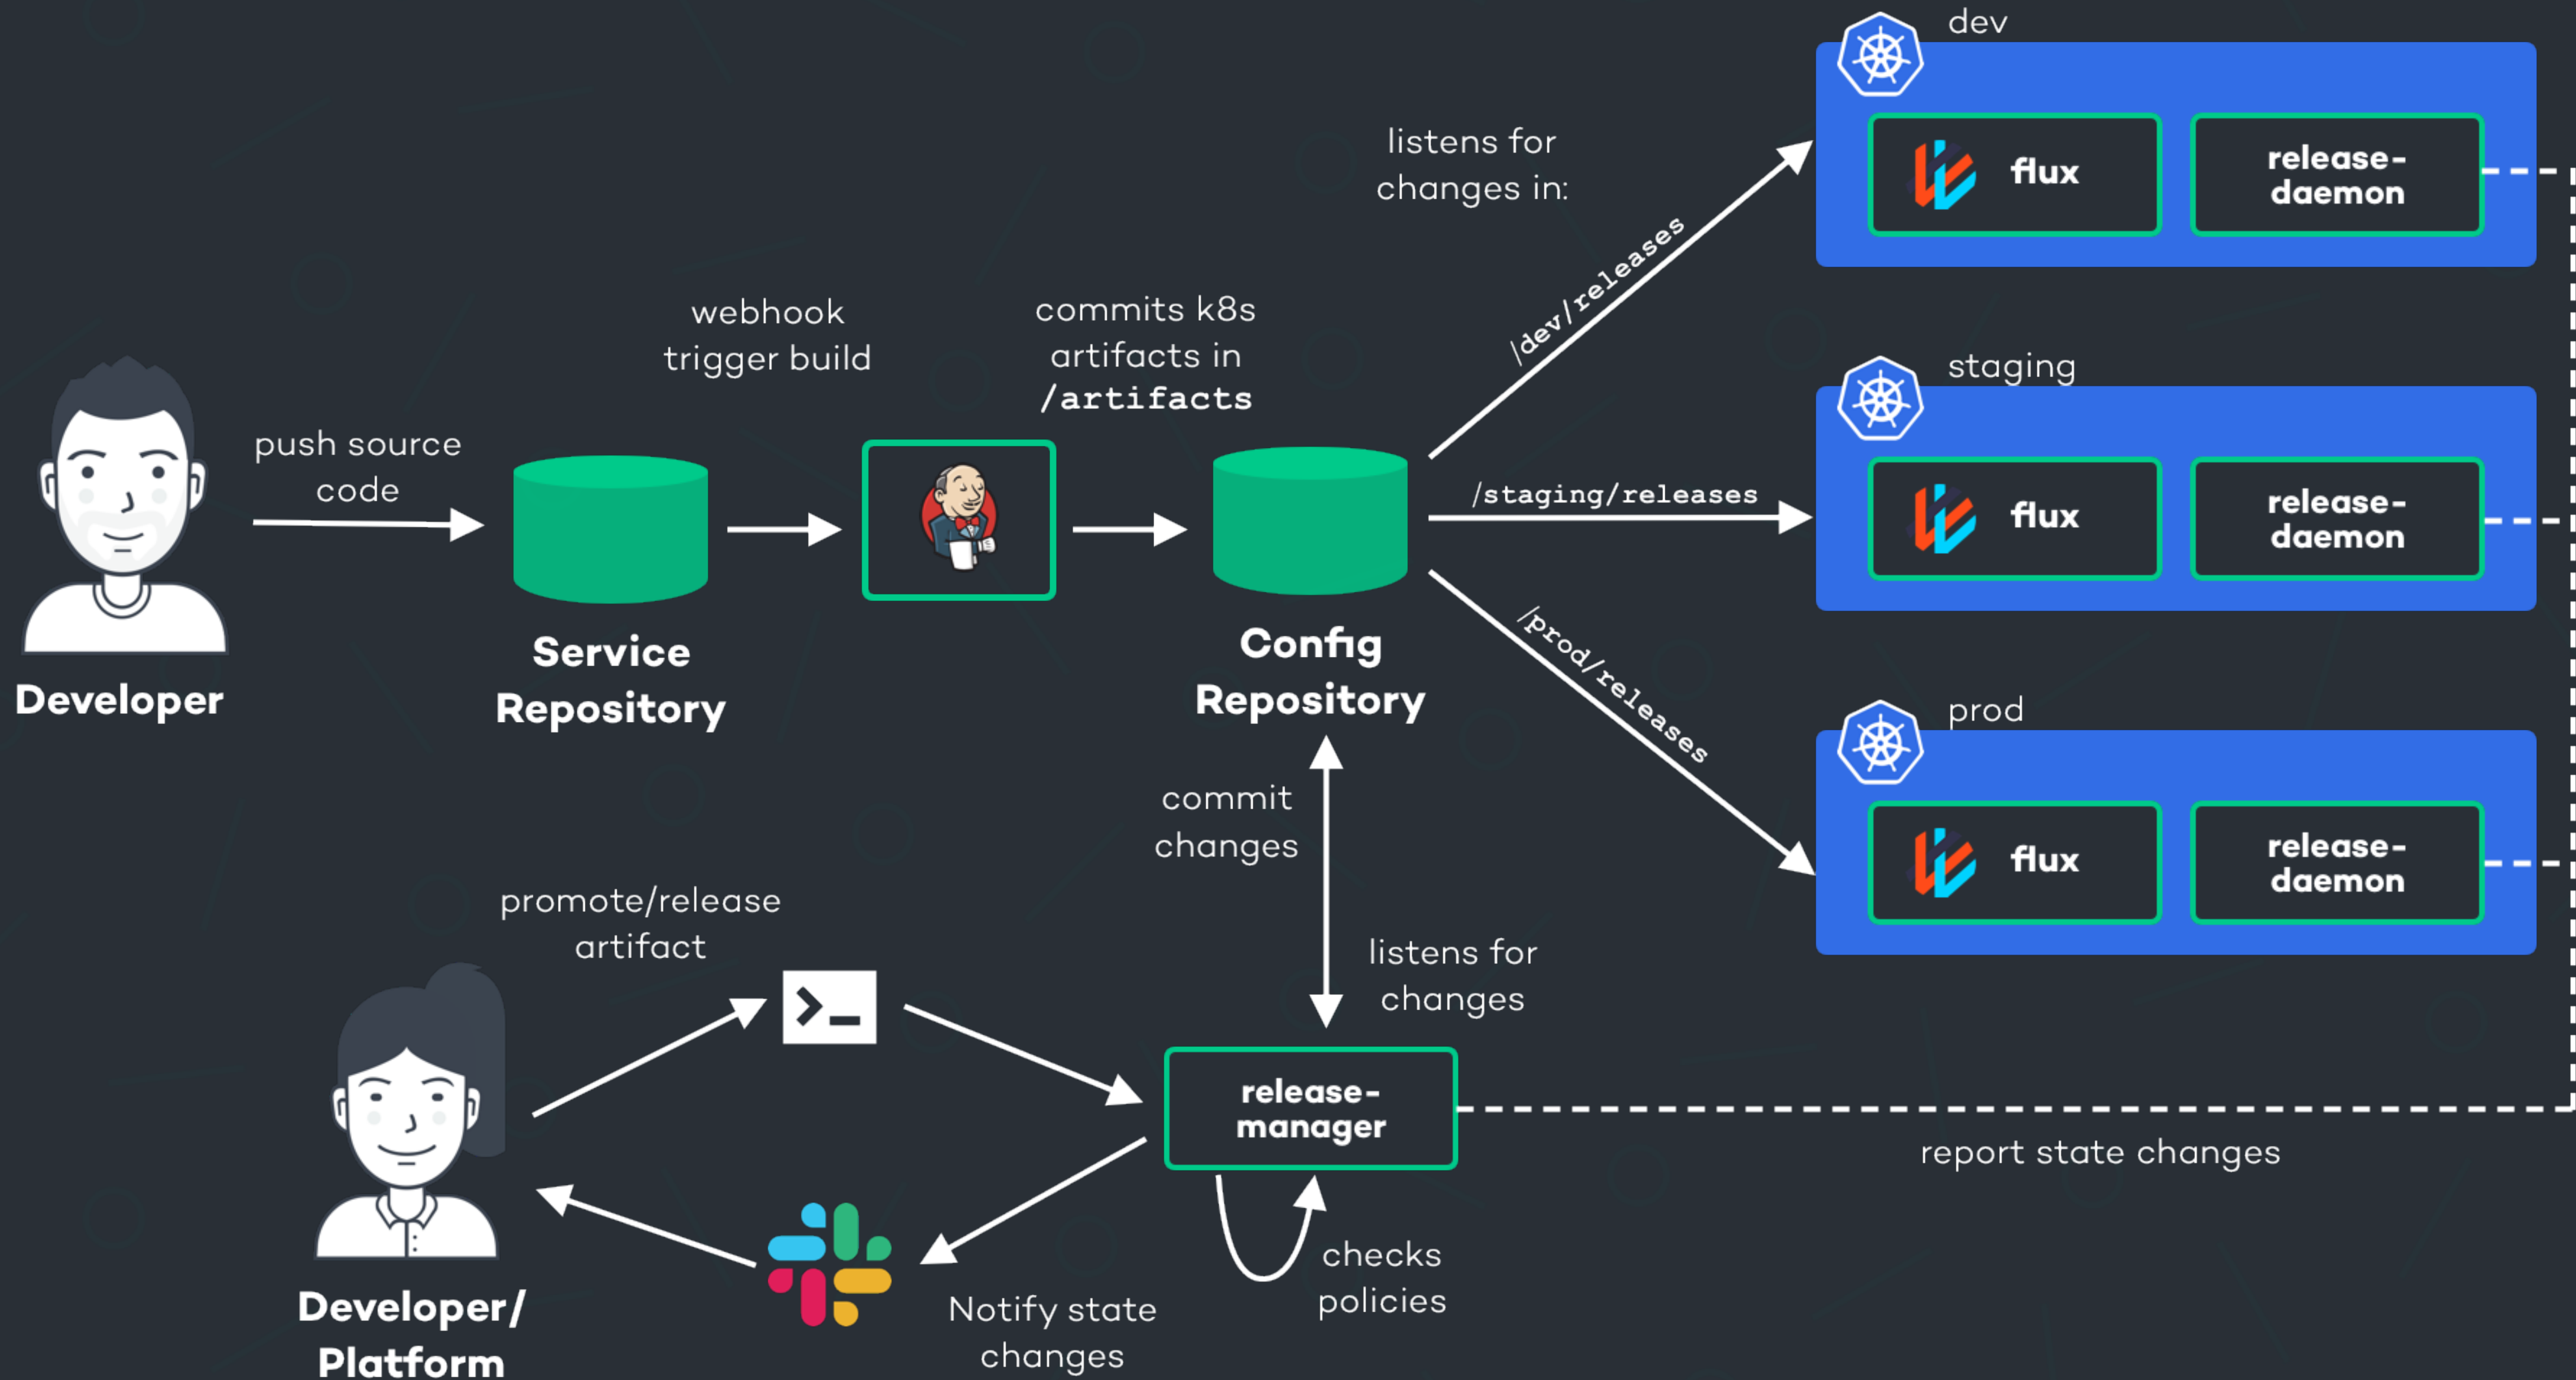
Release Management

Meet Ham



Using Git as the **source of truth**





hamctl

(ham cuddle or ham c-t-l or ham control)

```
hamctl promote --env dev
```

```
hamctl release --branch master --env dev
```

```
hamctl status
```

```
hamctl policy apply auto-release --branch master --env prod
```



Did it solve the problems?

Dynamic generation of k8s resources, resulted in **no audit trail** ✓

...





No single source of truth for a **disaster recovery situation** ✓

...

A lot of wasted time because of **duplicate builds** ✓

Operations & Maintenance

Setup, and how to keep everything updated

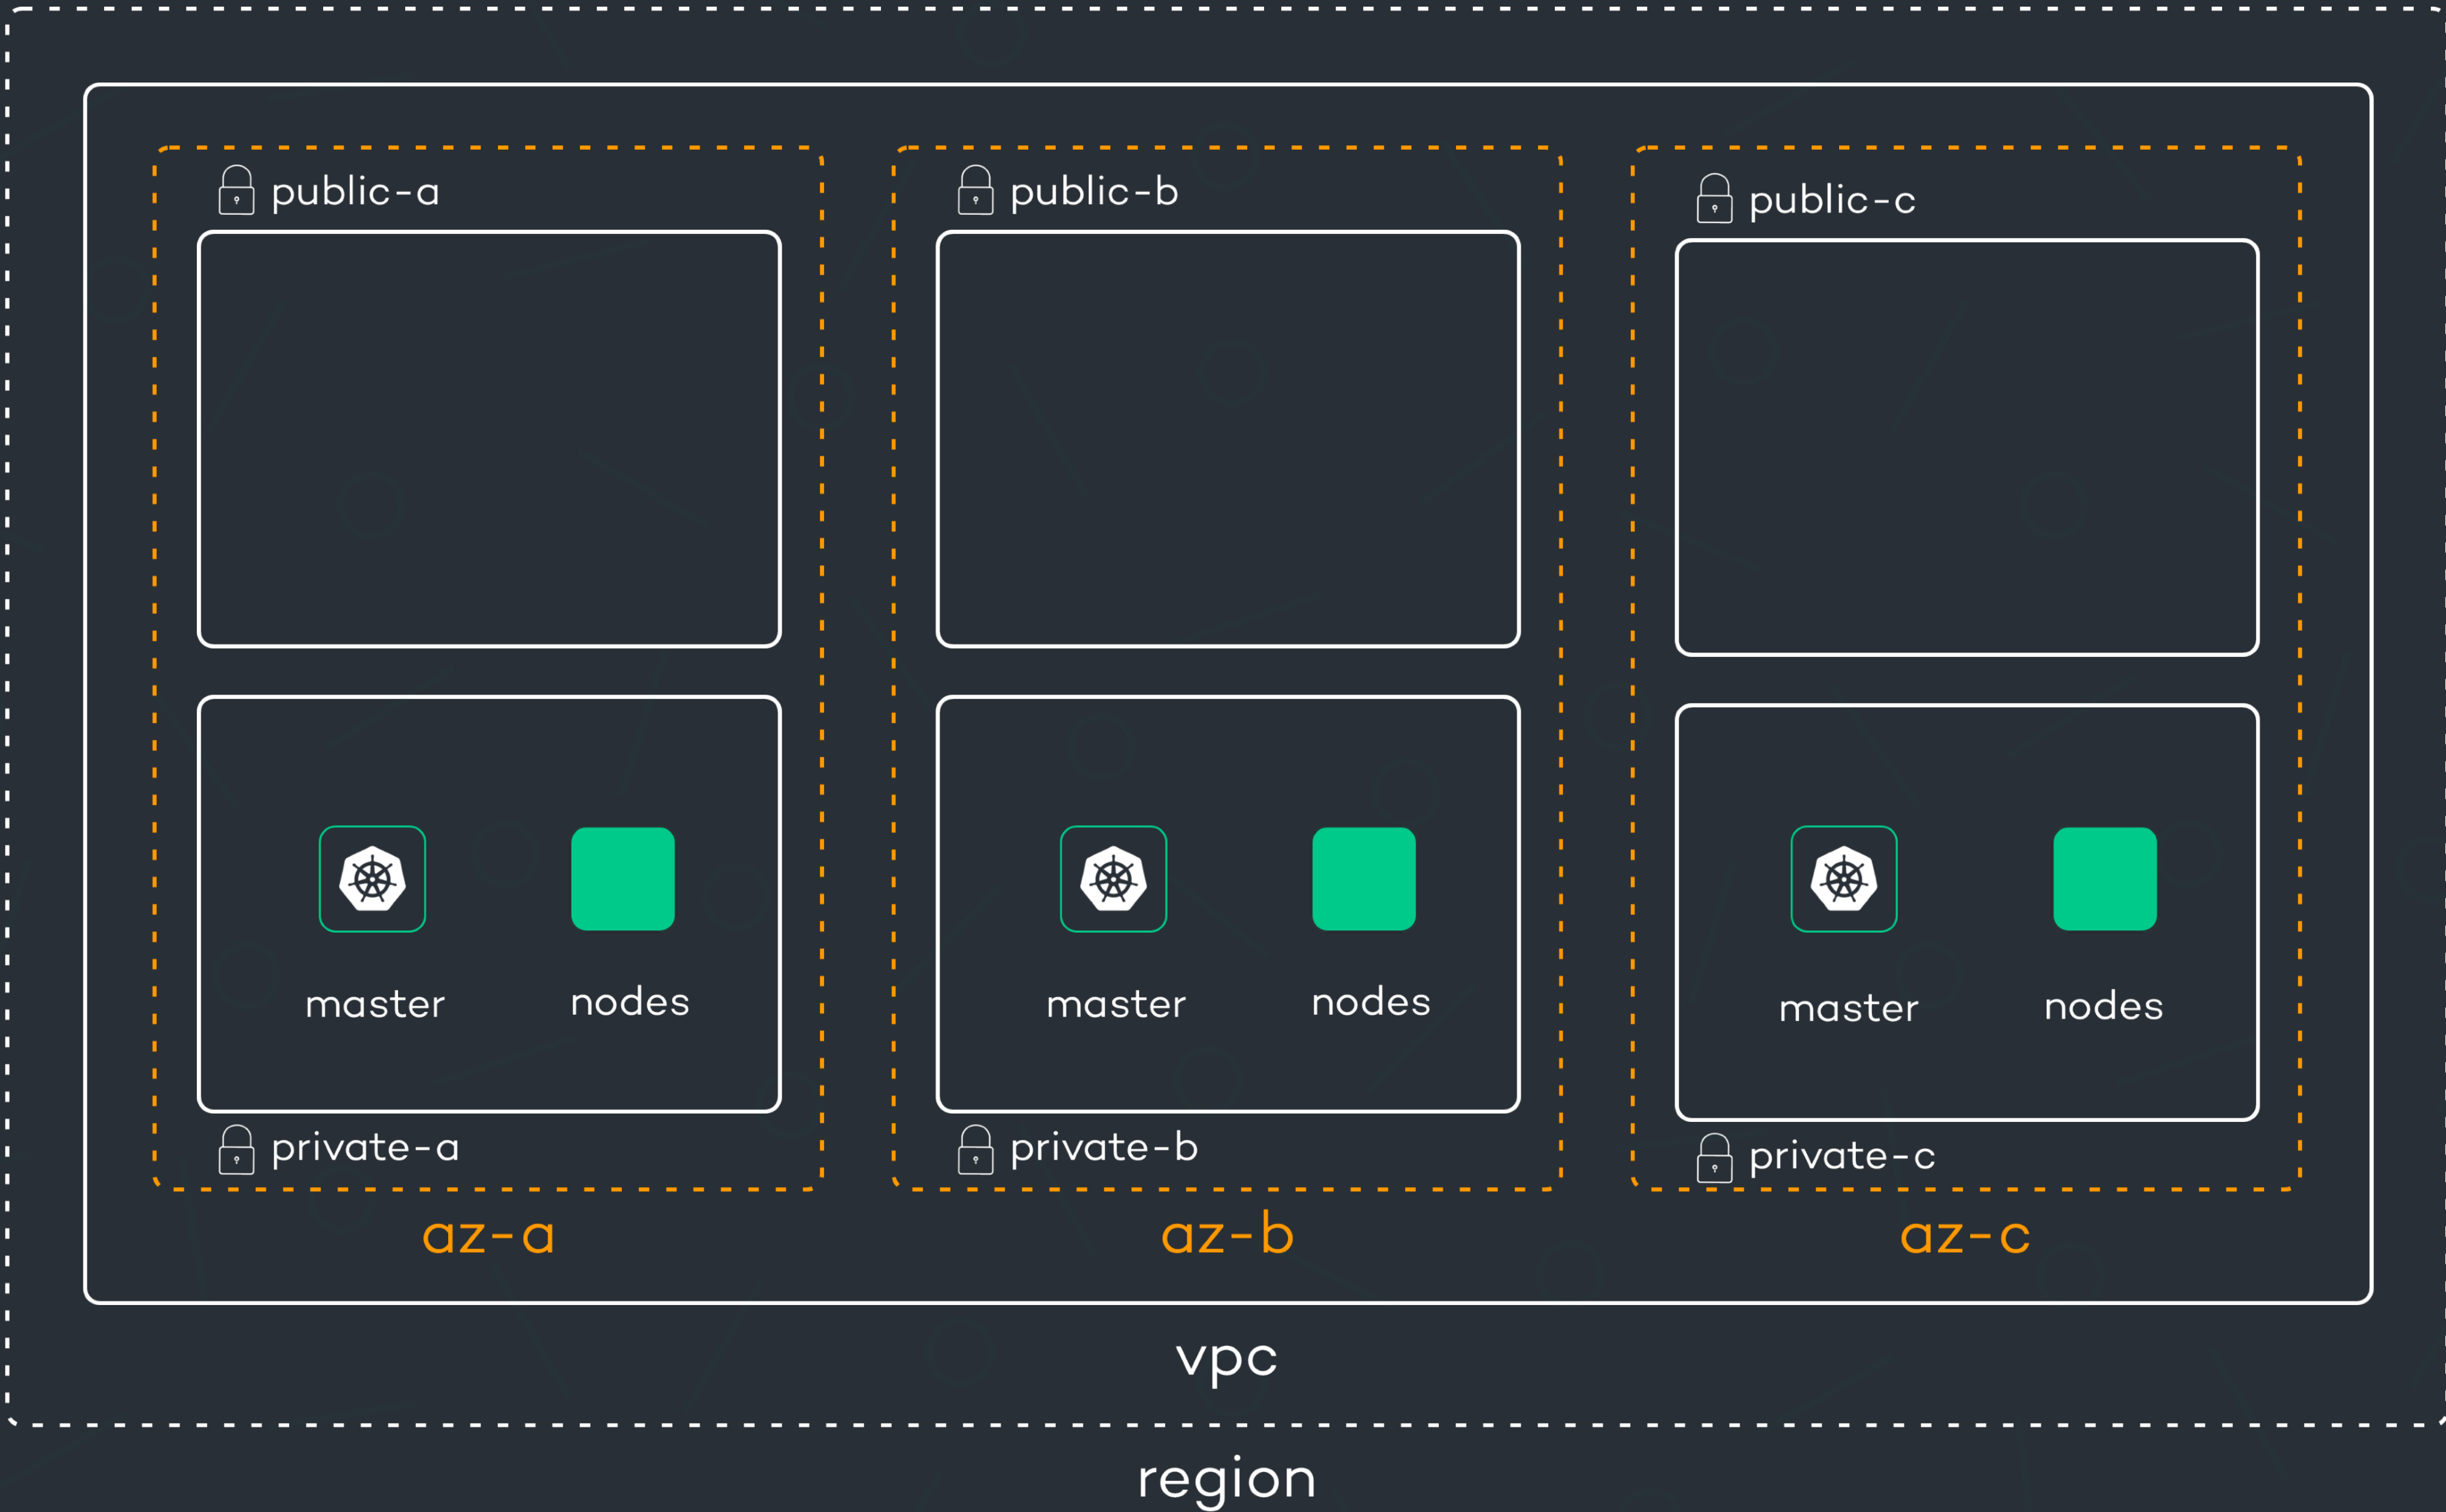


kops is basically, **kubectl** for
clusters

kops is a **semi-managed** solution primarily for **AWS**


```
$ export ZONES=${ZONES:-"eu-west-1a,eu-west-1b,eu-west-1c"}
$ export KOPS_STATE_STORE="s3://my-state-store"

$ kops create cluster k8s-clusters.example.com \
  --node-count 3 \
  --zones $ZONES \
  --master-zones $ZONES \
  --networking weave \
  --topology private \
  --bastion="true" \
  --yes
```



kops

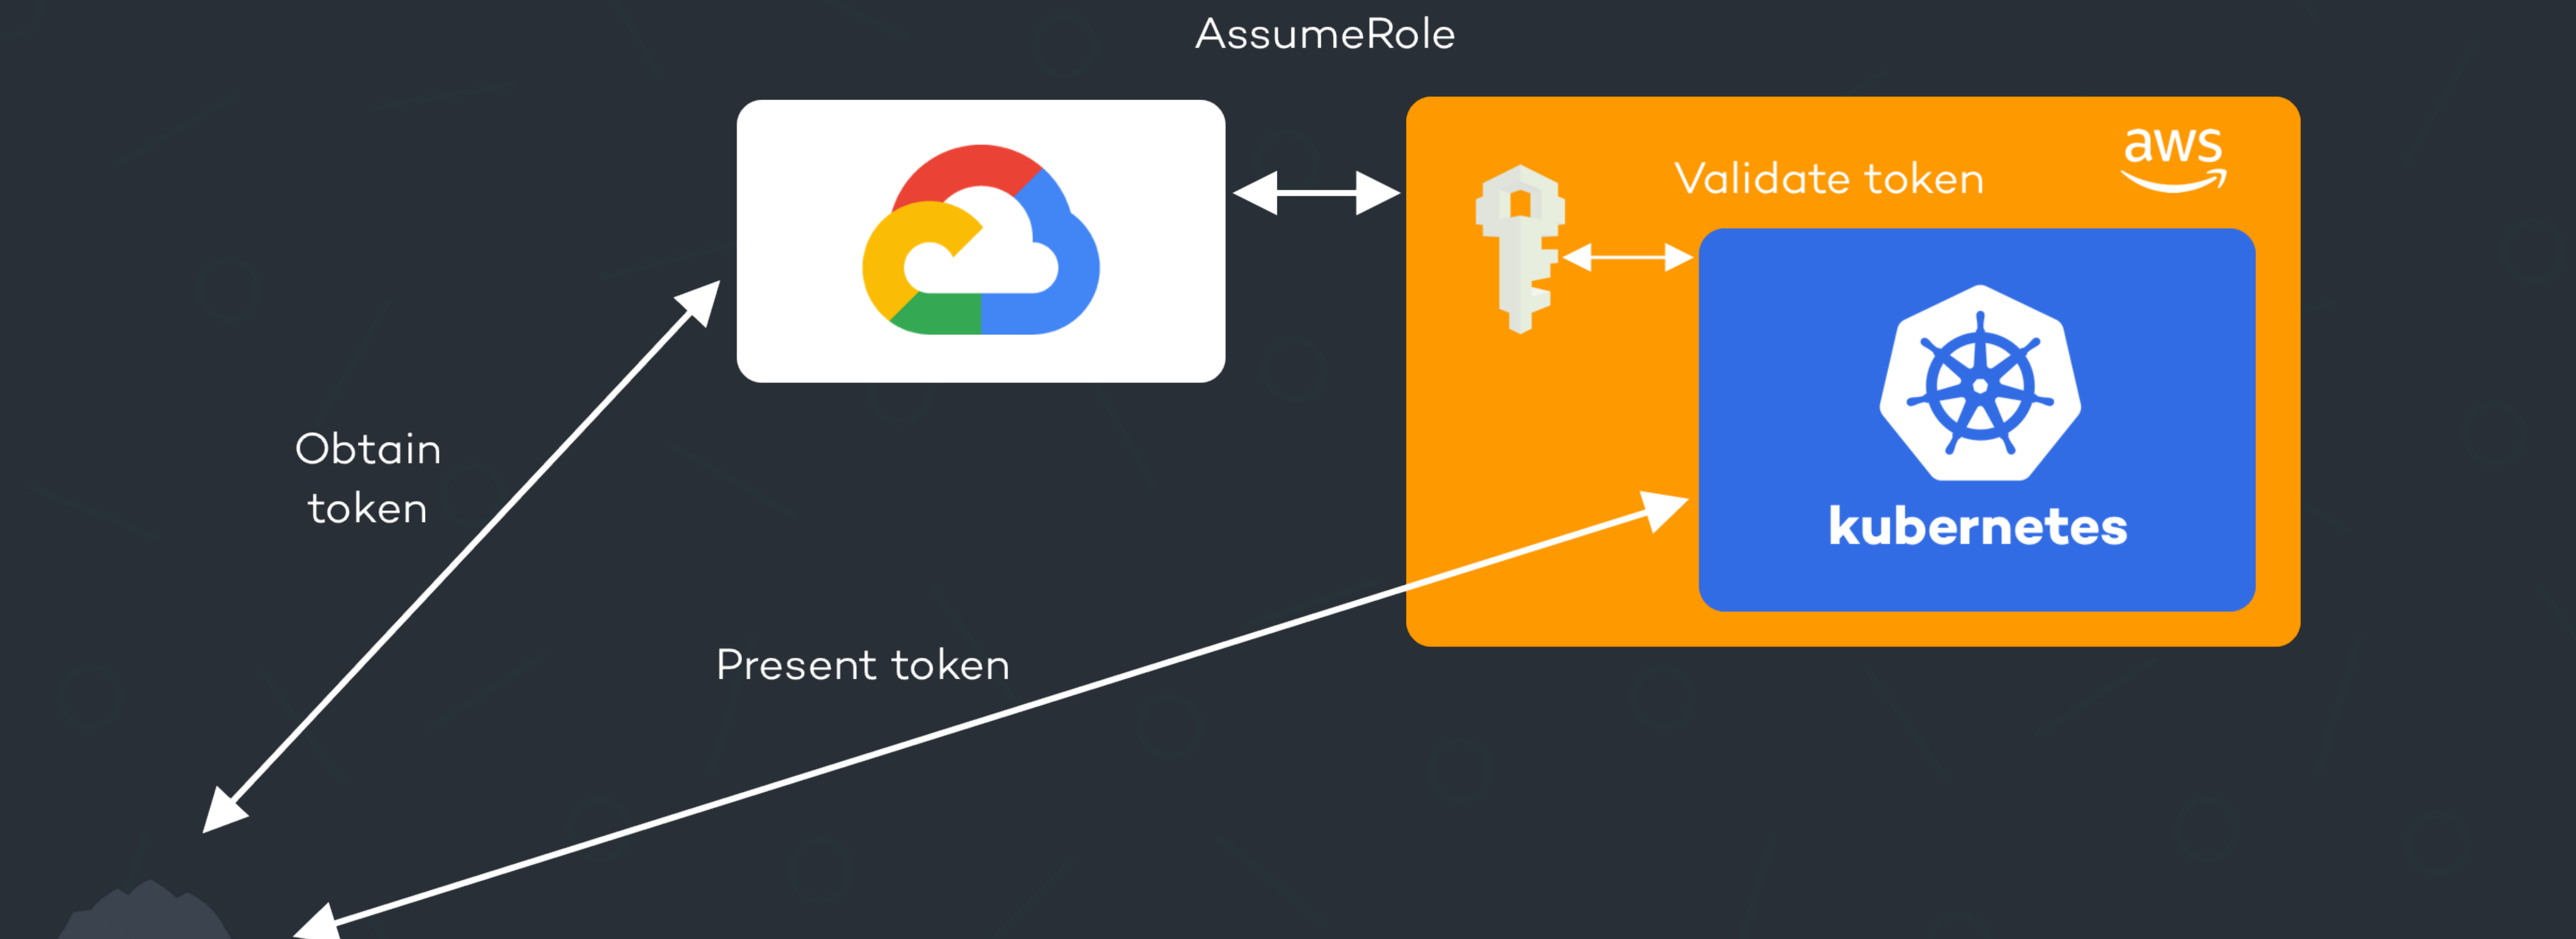
```
apiVersion: kops/v1alpha2
kind: InstanceGroup
metadata:
  creationTimestamp: 2019-06-13T10:17:11Z
  labels:
    kops.k8s.io/cluster: k8s.test.lunarway.com
  name: nodes
spec:
  image: kope.io/k8s-1.12-debian-stretch-amd64-hvm-ebs-2019-05-13
  machineType: m4.large
  maxSize: 3
  minSize: 3
  nodeLabels:
    kops.k8s.io/instancegroup: nodes
  role: Node
  subnets:
  - eu-west-1a
  - eu-west-1b
  - eu-west-1c
```

cluster-api

```
apiVersion: "cluster.k8s.io/v1alpha1"
kind: MachineDeployment
metadata:
  name: sample-machinedeployment
spec:
  replicas: 3
  selector:
    matchLabels:
      foo: bar
  template:
    metadata:
      labels:
        foo: bar
    spec:
      providerSpec:
        value:
          apiVersion: "gceproviderconfig/v1alpha1"
          kind: "GCEProviderConfig"
          roles:
            - Node
          project: "${GCPLOUD_PROJECT}"
          zone: "us-central1-c"
          machineType: "n1-standard-2"
          os: "ubuntu-1604-lts"
      versions:
        kubelet: 1.9.4
```



Accessing the clusters?




```

$ k8s dev
Switched to context "dev".
Google Password:
Open the Google App, and tap 'Yes' on the prompt to sign in ...
Assuming arn:aws:iam::1234567890:role/GoogleAdminLogin
Credentials Expiration: 2019-06-12 15:47:33+00:00


```

kops rolling-update cluster


1.



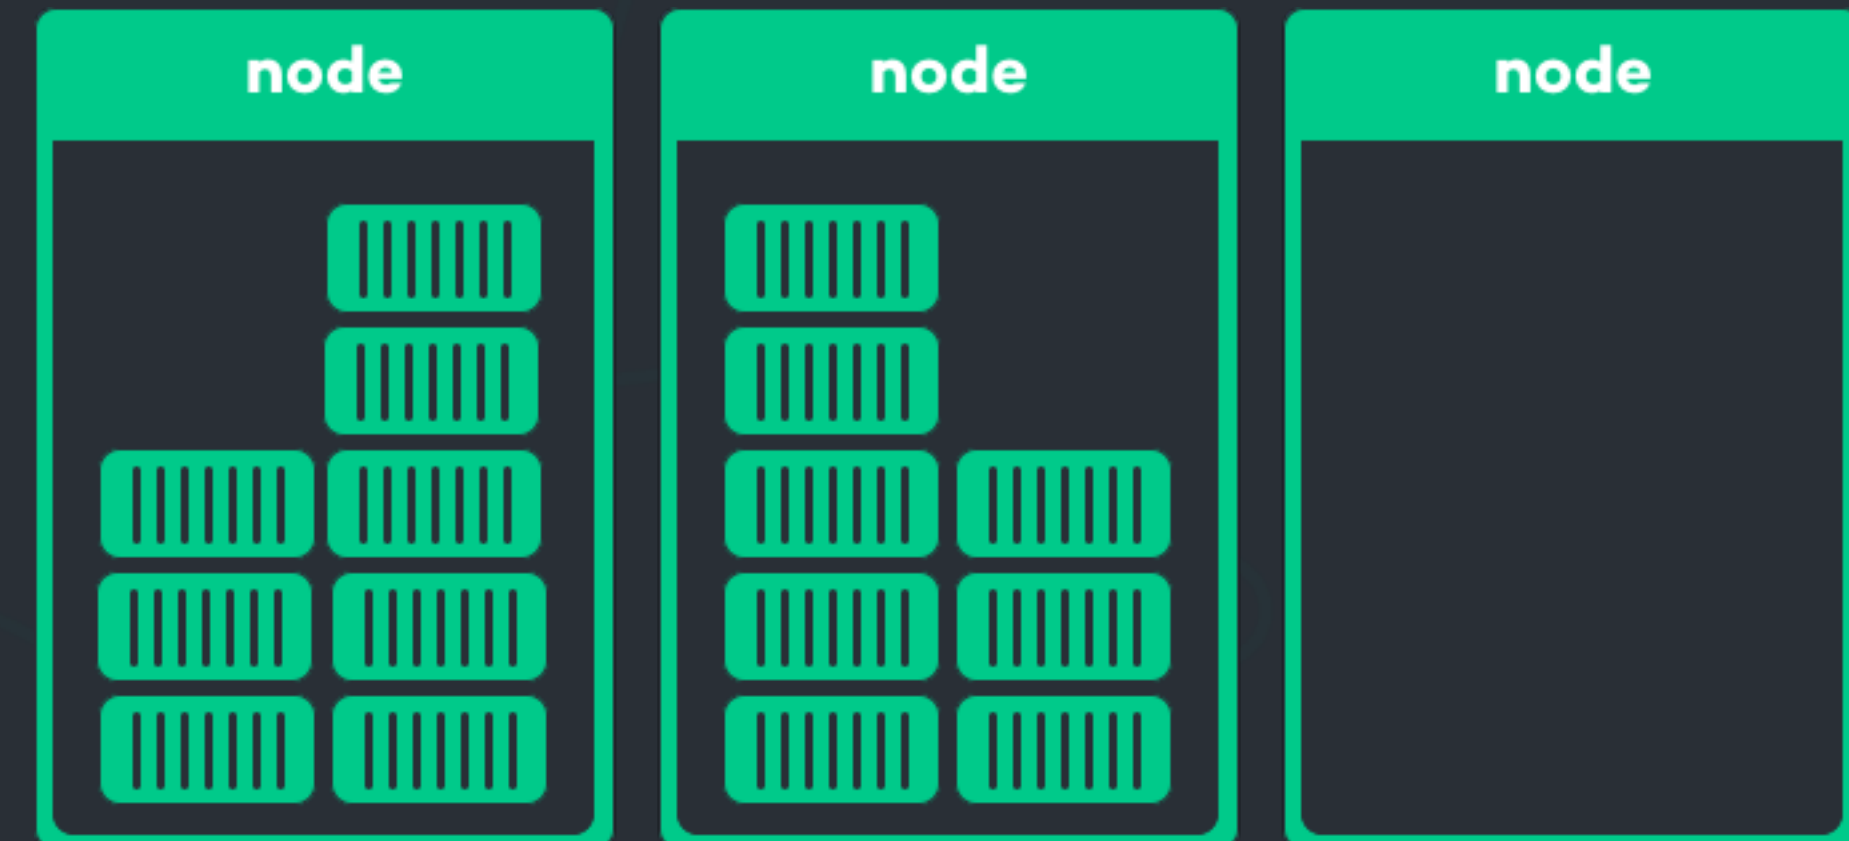
2.



3.



4.



Problems?

Congestion on nodes because of `docker pull`

...


Under `provisioned` cluster

...


`Rebalancing` needed after roll, to spread and fill empty node

dextre roll nodes

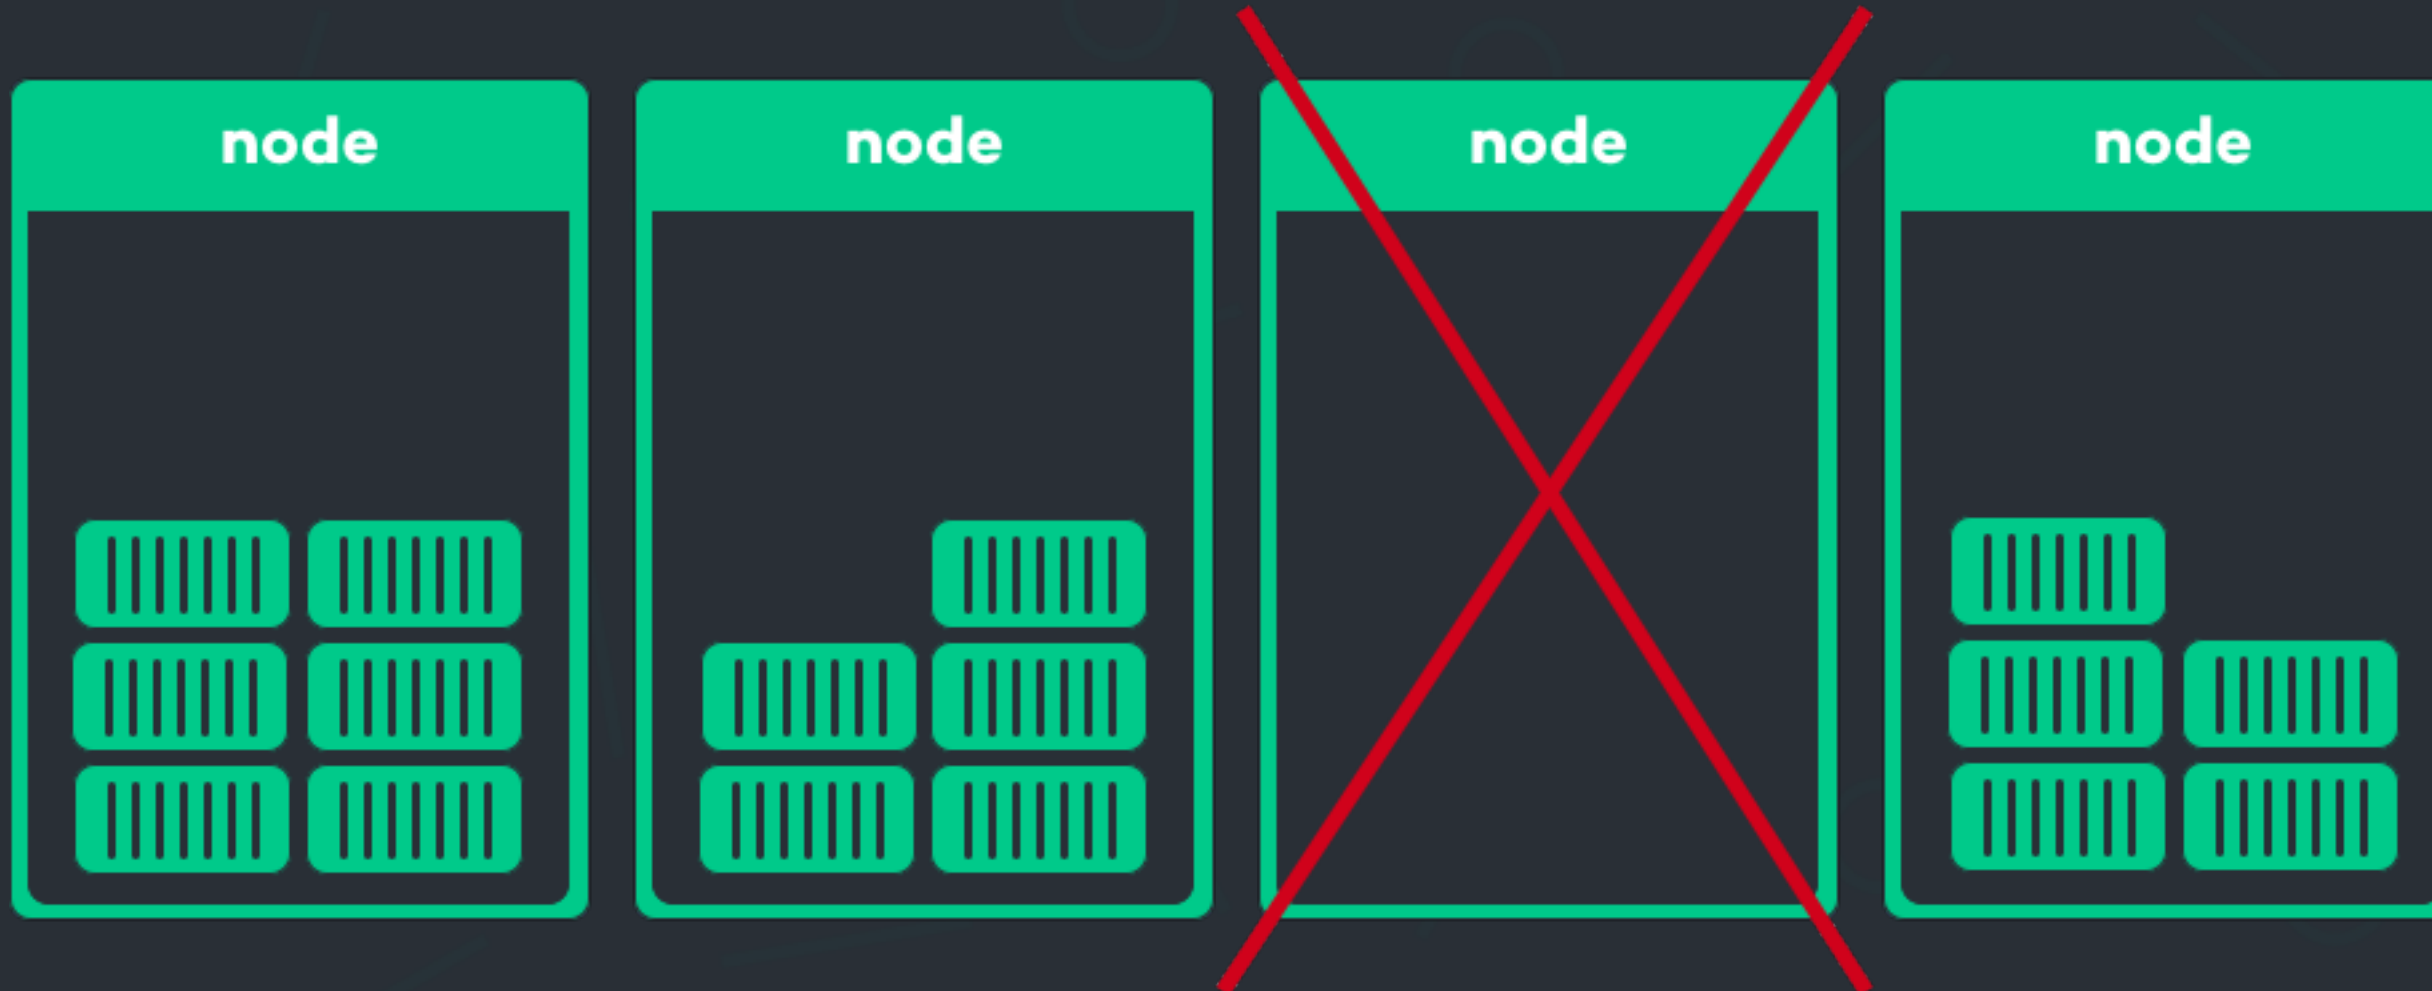
1.




2.



3.



4.



Did it solve the problems?

Congestion on nodes because of **docker pull** ✓

...

Under **provisioned** cluster ✓

...

Rebalancing needed after roll, to spread and fill empty node ✓

So, was Kubernetes the right decision, despite our size?

Yes, we wouldn't have been able to manage all those services without it.

”Adopting **Kubernetes is not just adopting another piece of software, it’s **introducing a new culture**”**

”Kubernetes is the **foundation
on which we build”**

“Kubernetes is not a destination. It’s part of the journey”

Bryan Liles, CloudNativeCon+KubeCon Europe 2019

**lunar
way®**

**lunar
way®**

Thank you!
Questions?

@phennex

kni@lunarway.com

lunar
way[®]

We are hiring!
jobs.lunarway.com