

***lunar
way***[®]

A tale of 1,5 years with Cloud Native technologies at Lunar Way

GOTO Nights CPH August 2018
Kasper Nissen - @phennex



Who?

Kasper Nissen (@phennex)

- Cloud Architect / SRE @lunarway
- Previous; LEGO Systems, IT Minds, Drivelogger
- Organiser & Co-Founder of Cloud Native Aarhus
- MSc. Computer Engineering
- Founder Cloud Native DK Slack Community
- Occasional speaker at meet ups and conferences



lunar
way[®]

AGENDA

- **Cloud Native Infrastructure**
- **Container Orchestration with Kubernetes**
- **Monitoring with Prometheus**
- **Log Management with fluentd and Humio**

Cloud Native Infrastructure

Cloud Native, the CNCF definition

Cloud native technologies empower organizations to build and run **scalable applications** in modern, **dynamic environments** such as public, private, and hybrid clouds. **Containers, service meshes, microservices, immutable infrastructure, and declarative APIs** exemplify this approach.

These techniques enable **loosely coupled** systems that are **resilient, manageable, and observable**. Combined with robust automation, they allow engineers to make **high-impact changes frequently** and **predictably** with **minimal toil**.

The Cloud Native Computing Foundation seeks to drive adoption of this **paradigm** by fostering and sustaining an ecosystem of **open source, vendor-neutral projects**. We democratize state-of-the-art patterns to make these innovations accessible for everyone.

Key characteristics

- Scalable systems
- Dynamic environments
- Containers
- Immutable Infrastructure
- Observability and manageability
- Open source

Why do this?



Speed!

App Definition and Development

Database and Data Warehouse

Streaming

Source Code Management

Application Definition and Image Build

Continuous Integration / Continuous Delivery (CI/CD)

Orchestration & Management

Scheduling & Orchestration

Coordination & Service Discovery

Service Management

Runtime

Cloud-Native Storage

Container Runtime

Cloud-Native Network

Provisioning

Host Management / Tooling

Infrastructure Automation

Container Registries


Secure Images

Key Management

Cloud

Public

Private



l.cncf.io

This landscape is intended as a map through the previously uncharted terrain of cloud native technologies. There are many routes to deploying a cloud native application, with CNCF Projects representing a particularly well-traveled path.

CLOUD NATIVE Landscape

CLOUD NATIVE COMPUTING FOUNDATION

Redpoint Amplify

Platforms

Certified Kubernetes - Distribution

Certified Kubernetes - Hosted

Certified Kubernetes - Installer

Non-Certified Kubernetes

PaaS/Container Service

Kubernetes Certified Service Provider

Kubernetes Training Partner

Observability & Analysis

Monitoring


Logging

Tracing

Serverless



Special



Tools

Security

Framework

Platform

Hosted

Installable

s.cncf.io

Serverless computing refers to a new model of cloud native computing, enabled by architectures that do not require server management to build and run applications. This landscape illustrates a finer-grained deployment model where applications, bundled as one or more functions, are uploaded to a platform and then executed, scaled, and billed in response to the exact demand needed at the moment.

Cloud Native Landscape

Today, we will setup...

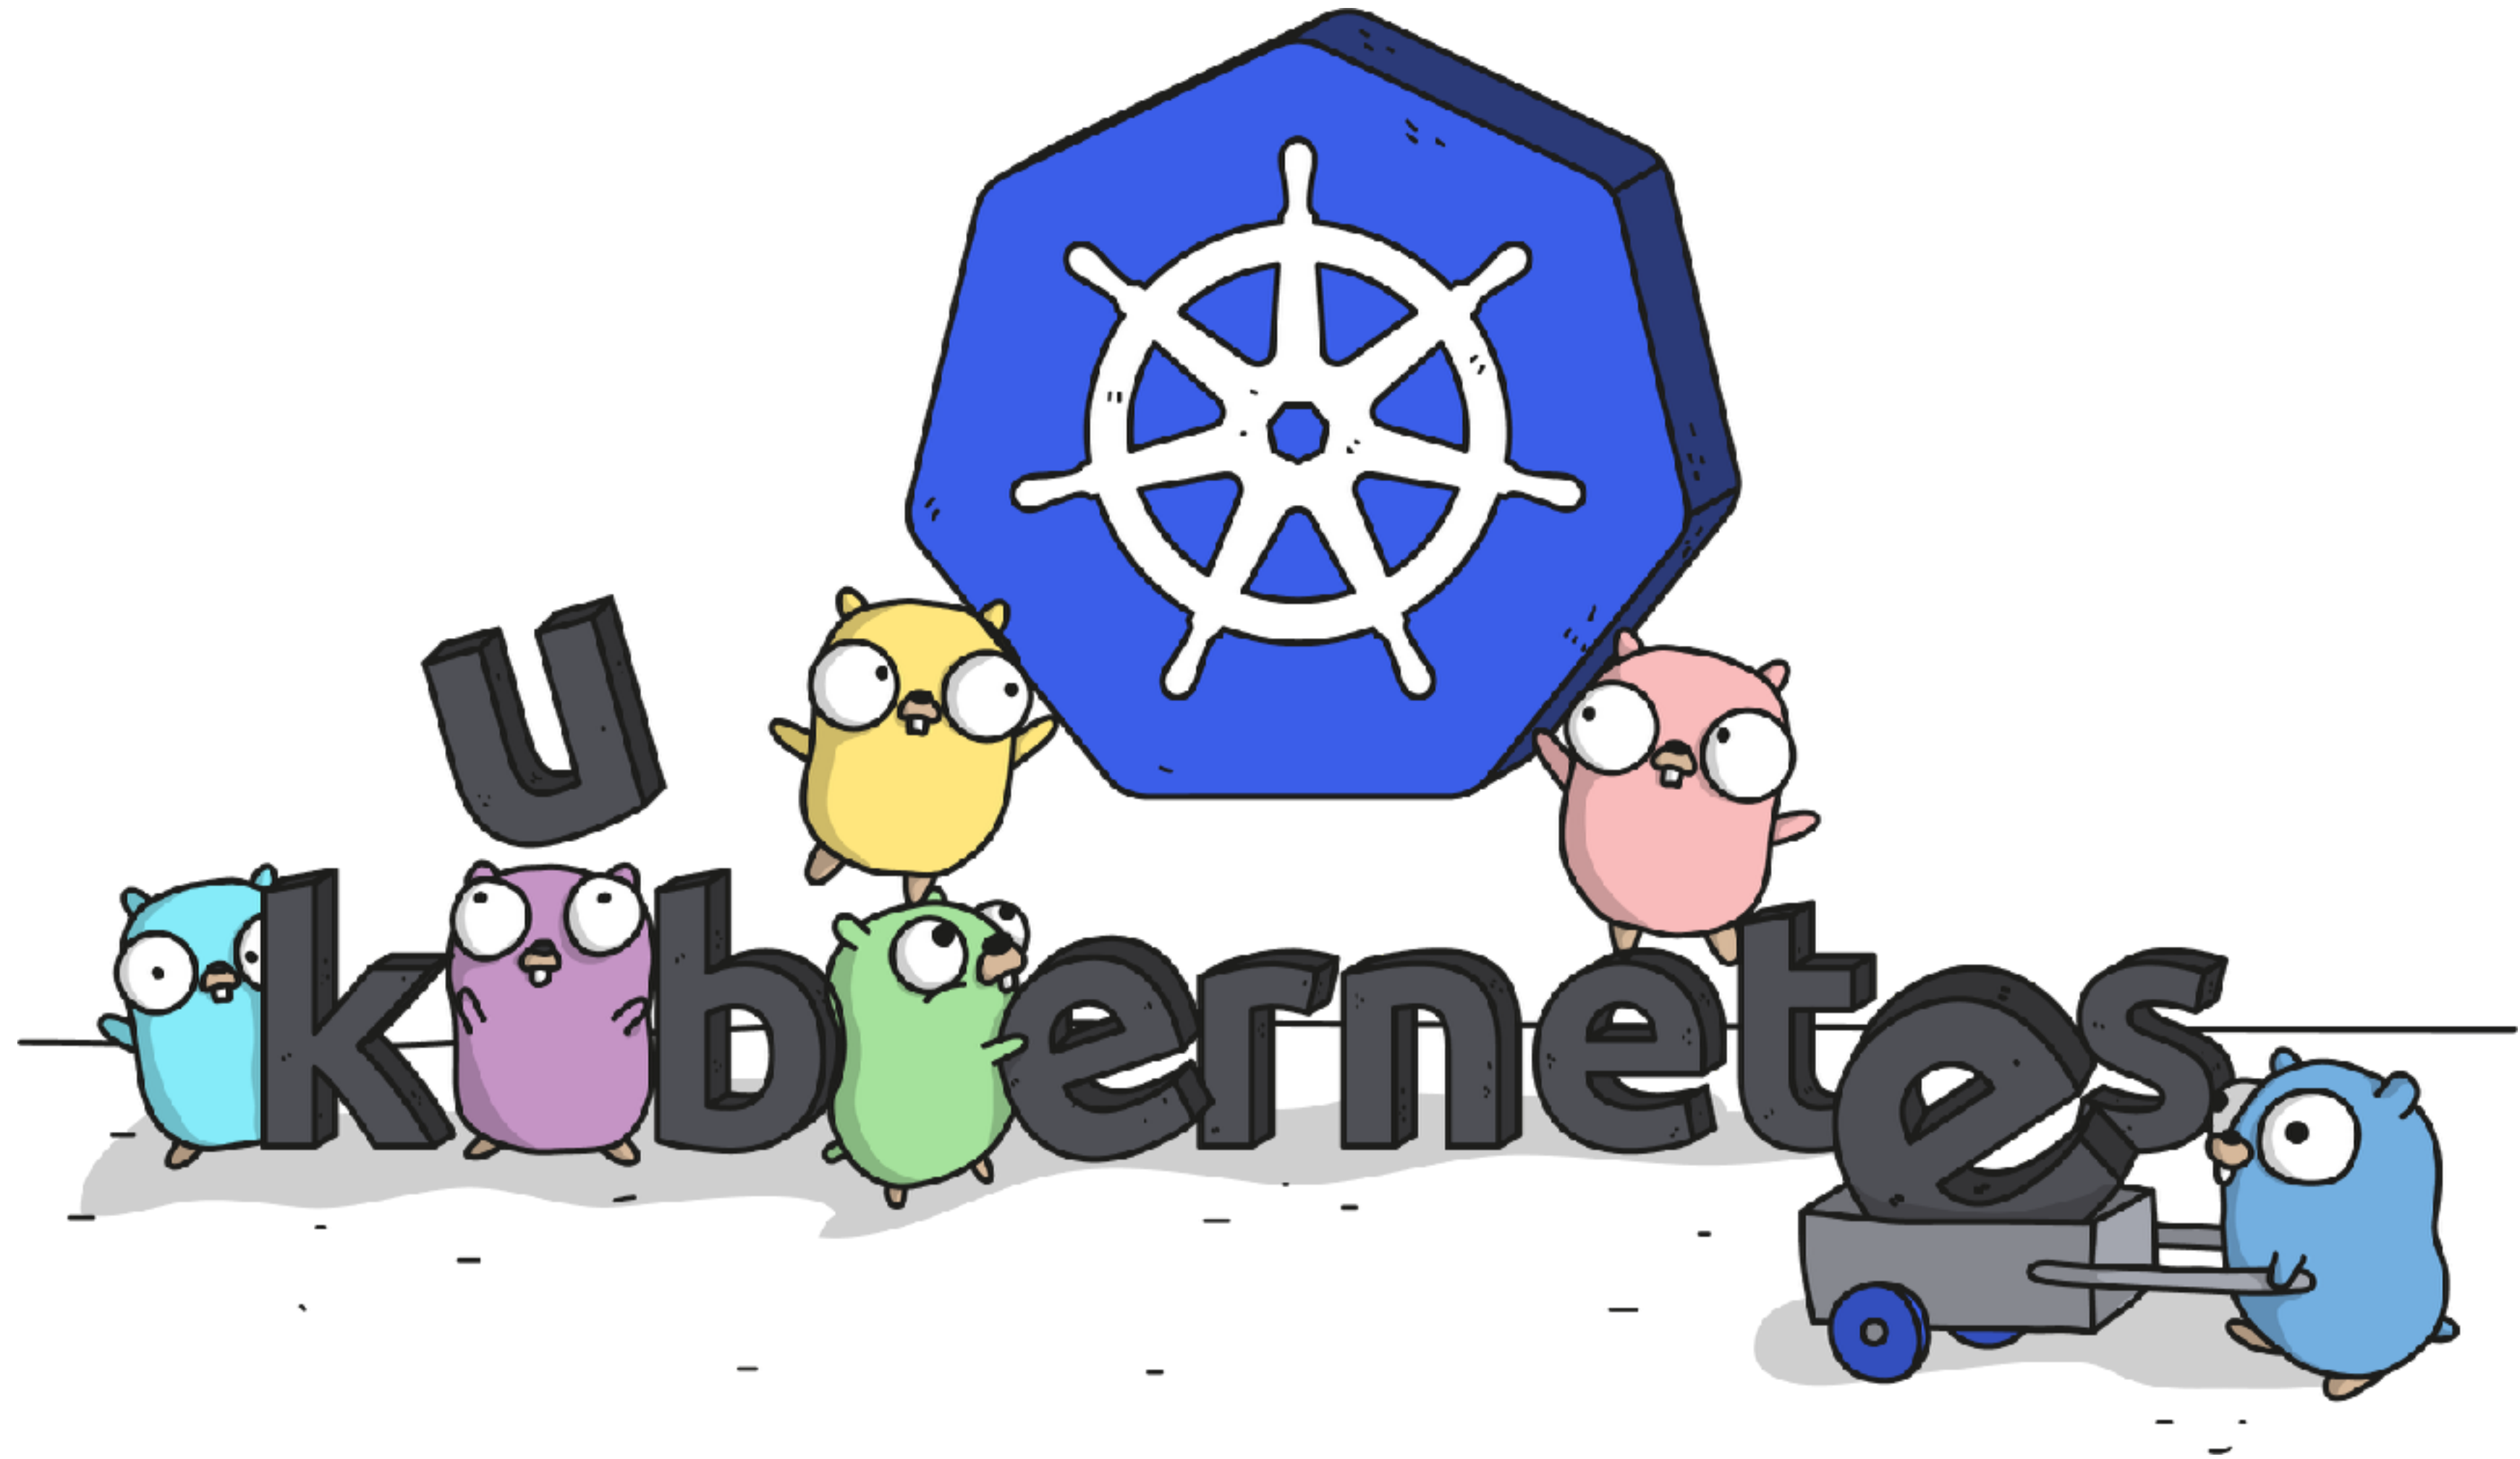
- A highly available Kubernetes cluster across 3 availability zones
- Route traffic from the internet to the cluster
- Monitoring with Prometheus
- Log collection with fluentd
- Log management with Humio

... and hopefully Thomas will build upon this setup!

A large container ship is docked at a port. The ship's hull is dark blue with the letters 'UASC' visible. The deck is filled with stacks of colorful shipping containers. Several large red cranes are positioned along the ship's length, with the word 'RDKAI' visible on their booms. The background shows a clear sky and the silhouettes of other port infrastructure.

Container Orchestration with Kubernetes

<https://www.pexels.com/photo/shallow-focus-photography-of-black-ship-1095814/>



Kubernetes, what is it?

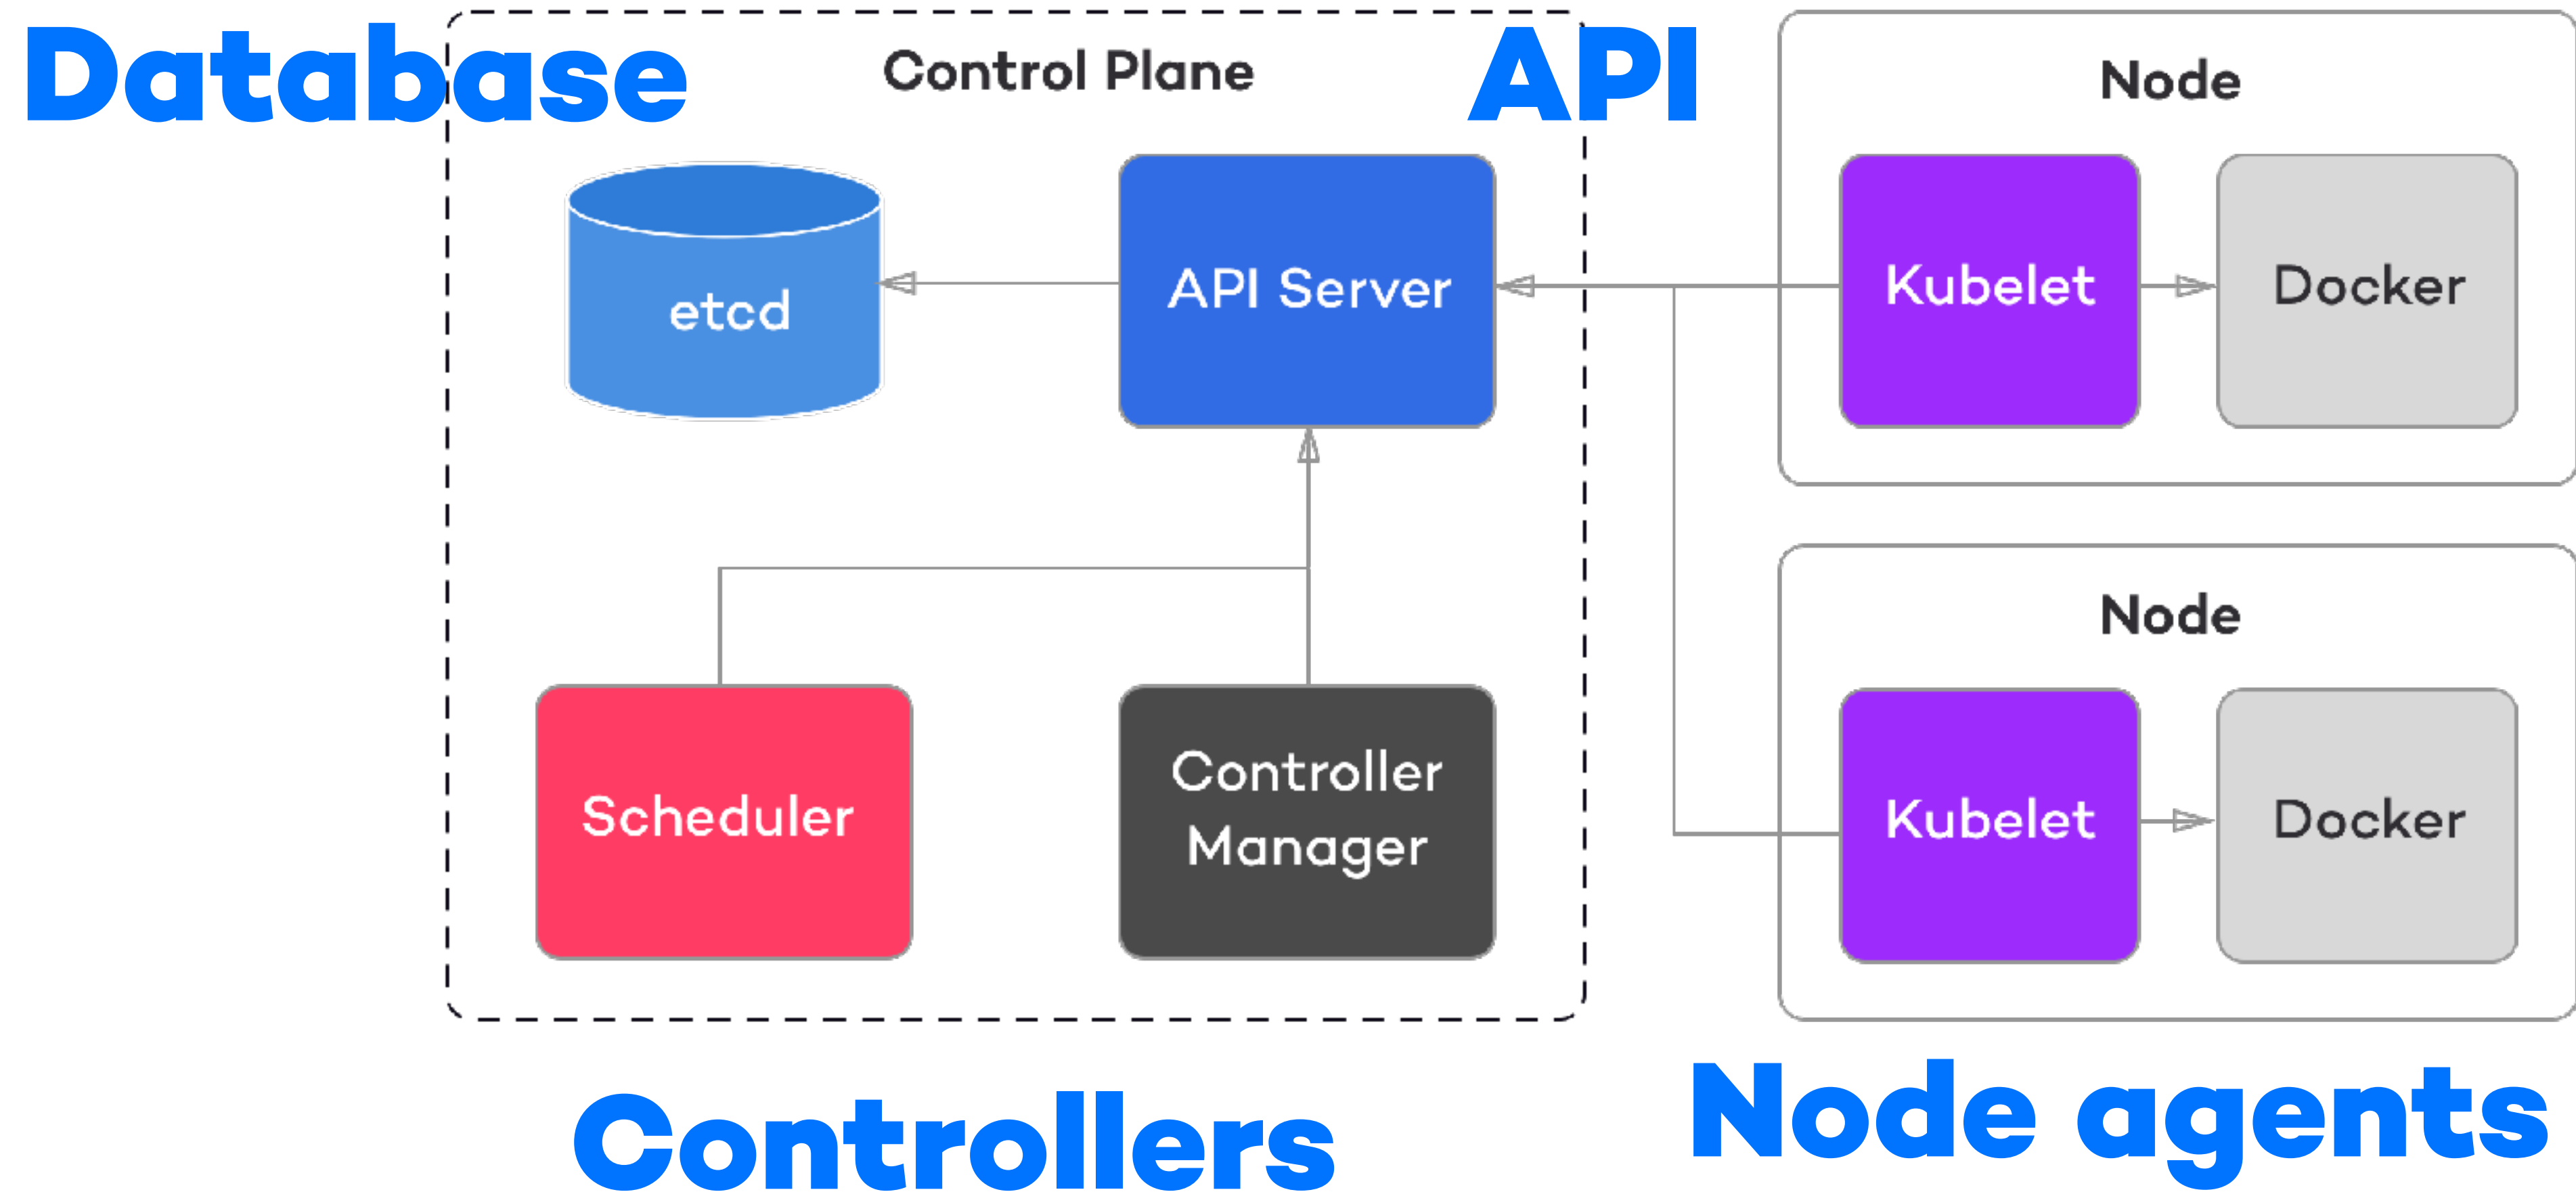
A Cluster Orchestrator

- Makes decisions on which nodes containers should be started
- Provides Service Discovery
- Integrate with networking

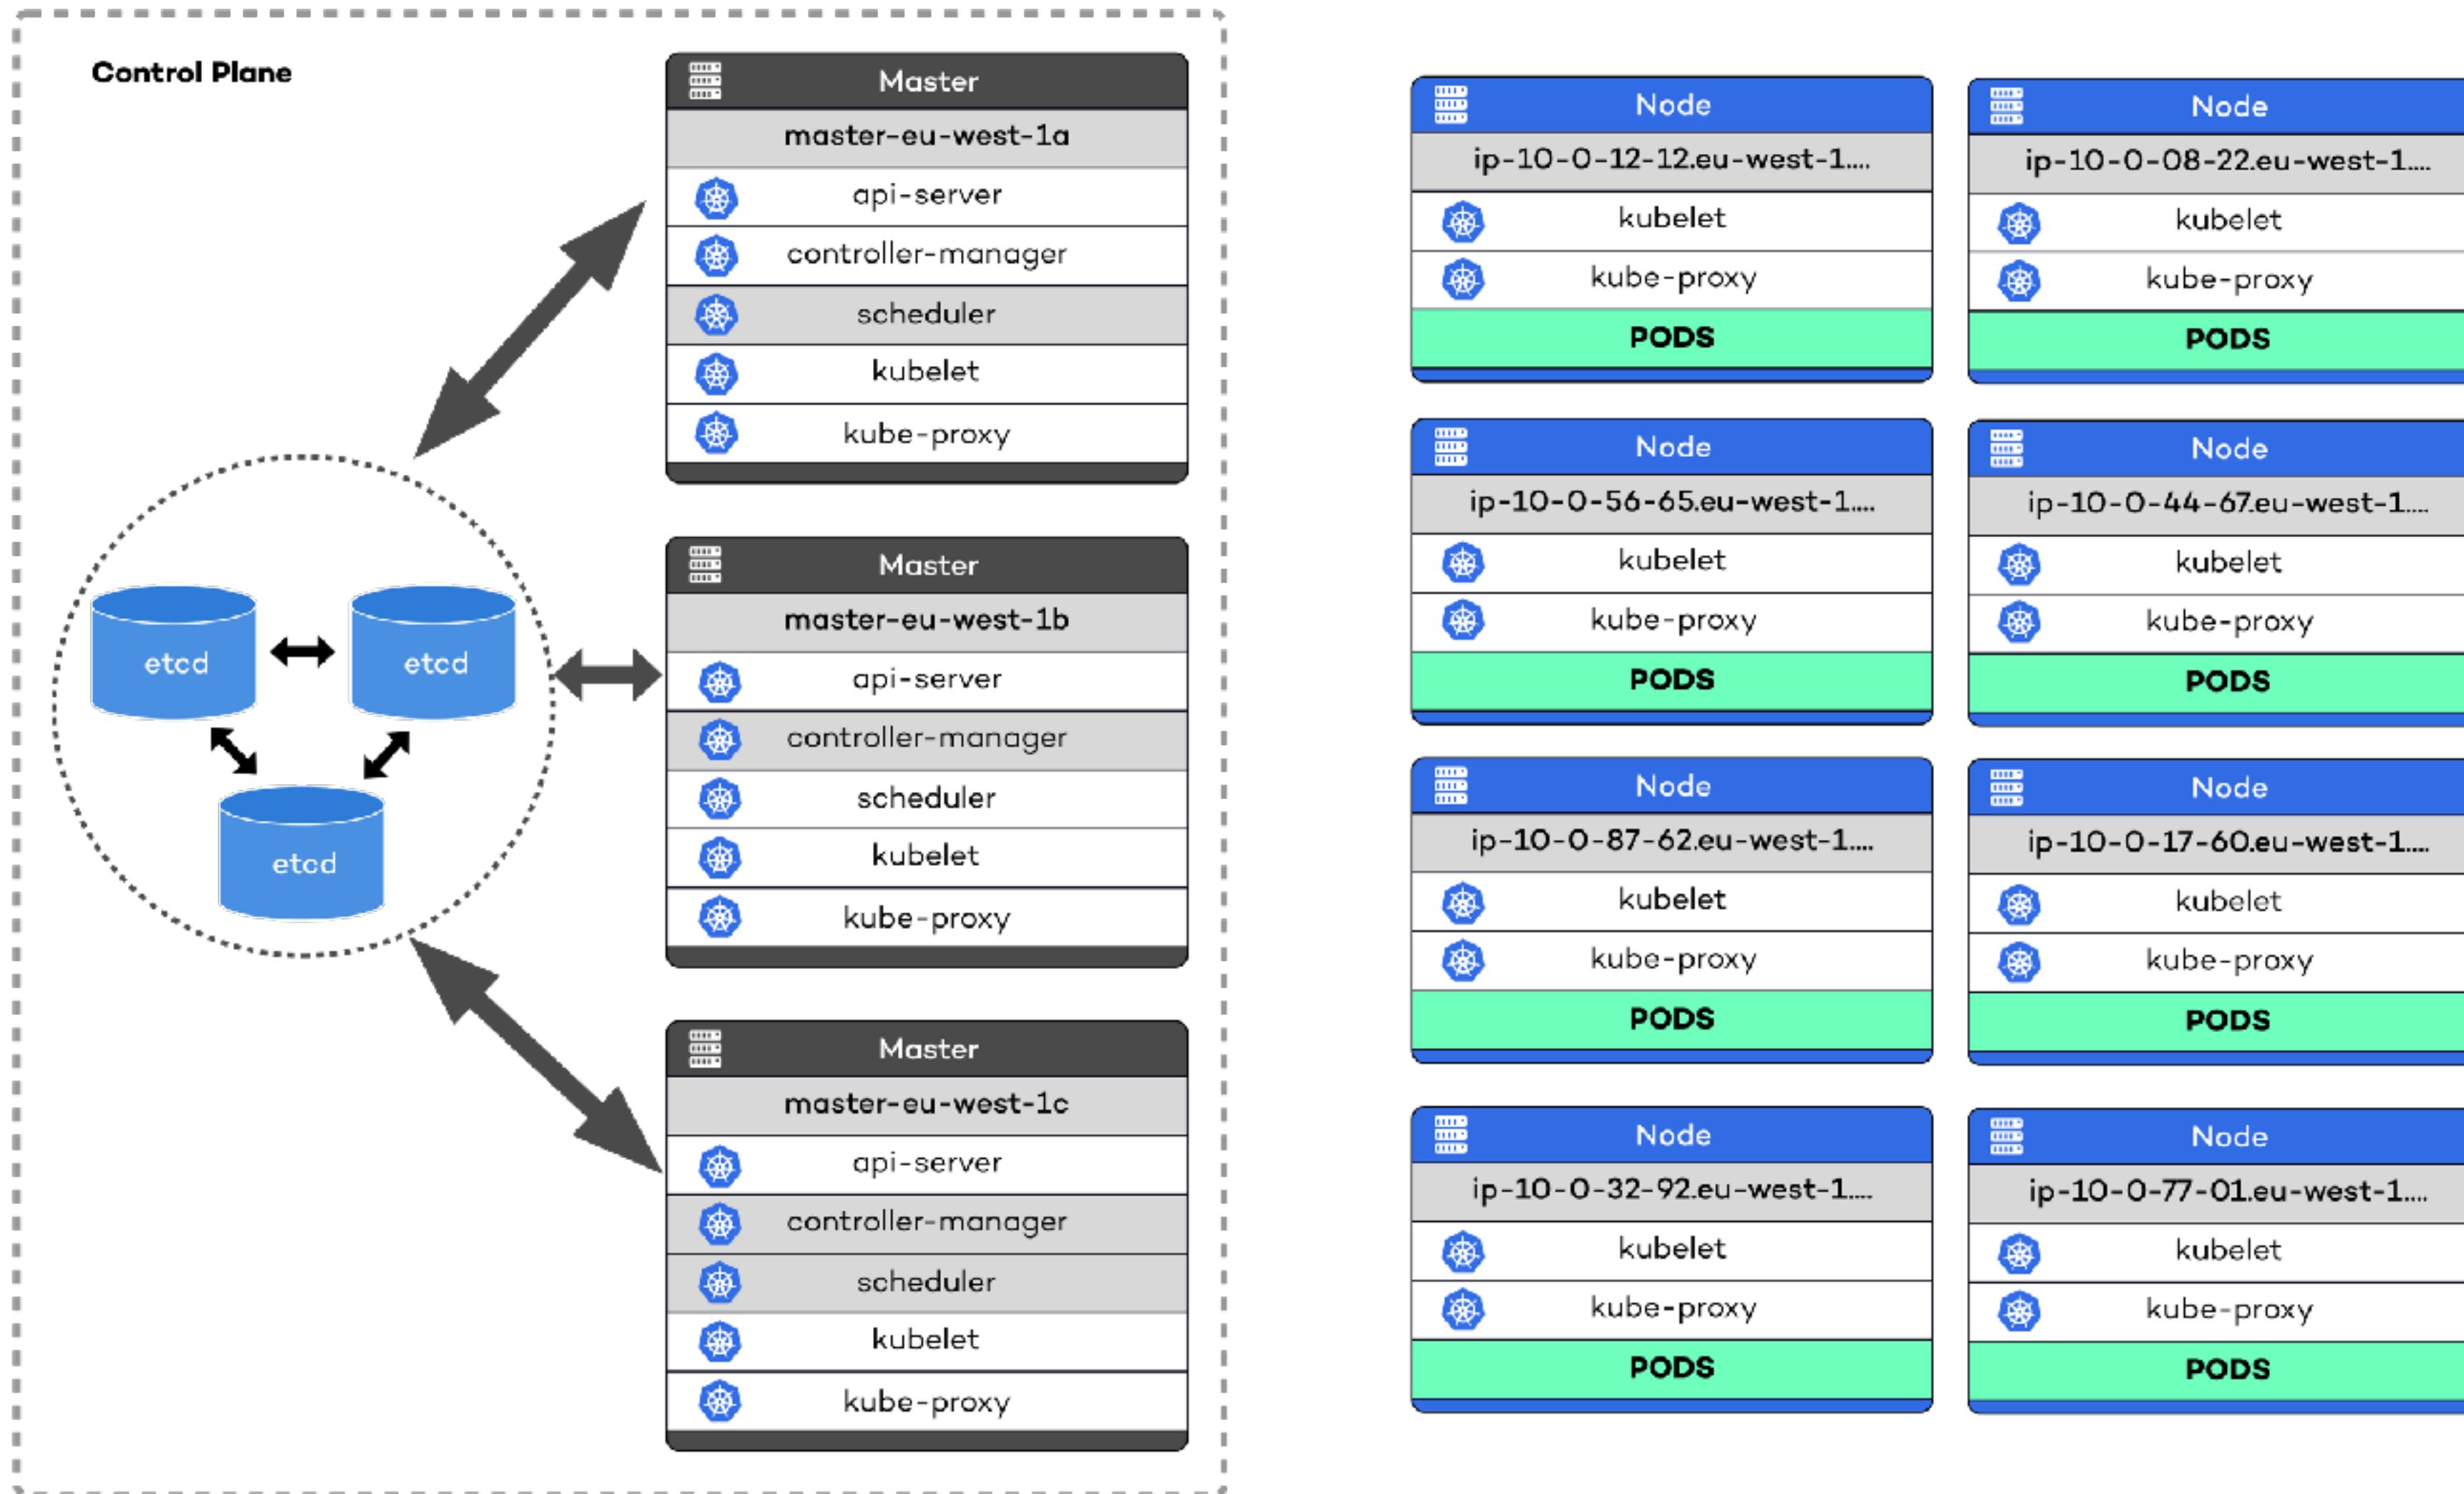
Value Proposition

- Increased developer and operator productivity
- Better resource utilisation (bin packing)
- Resilience

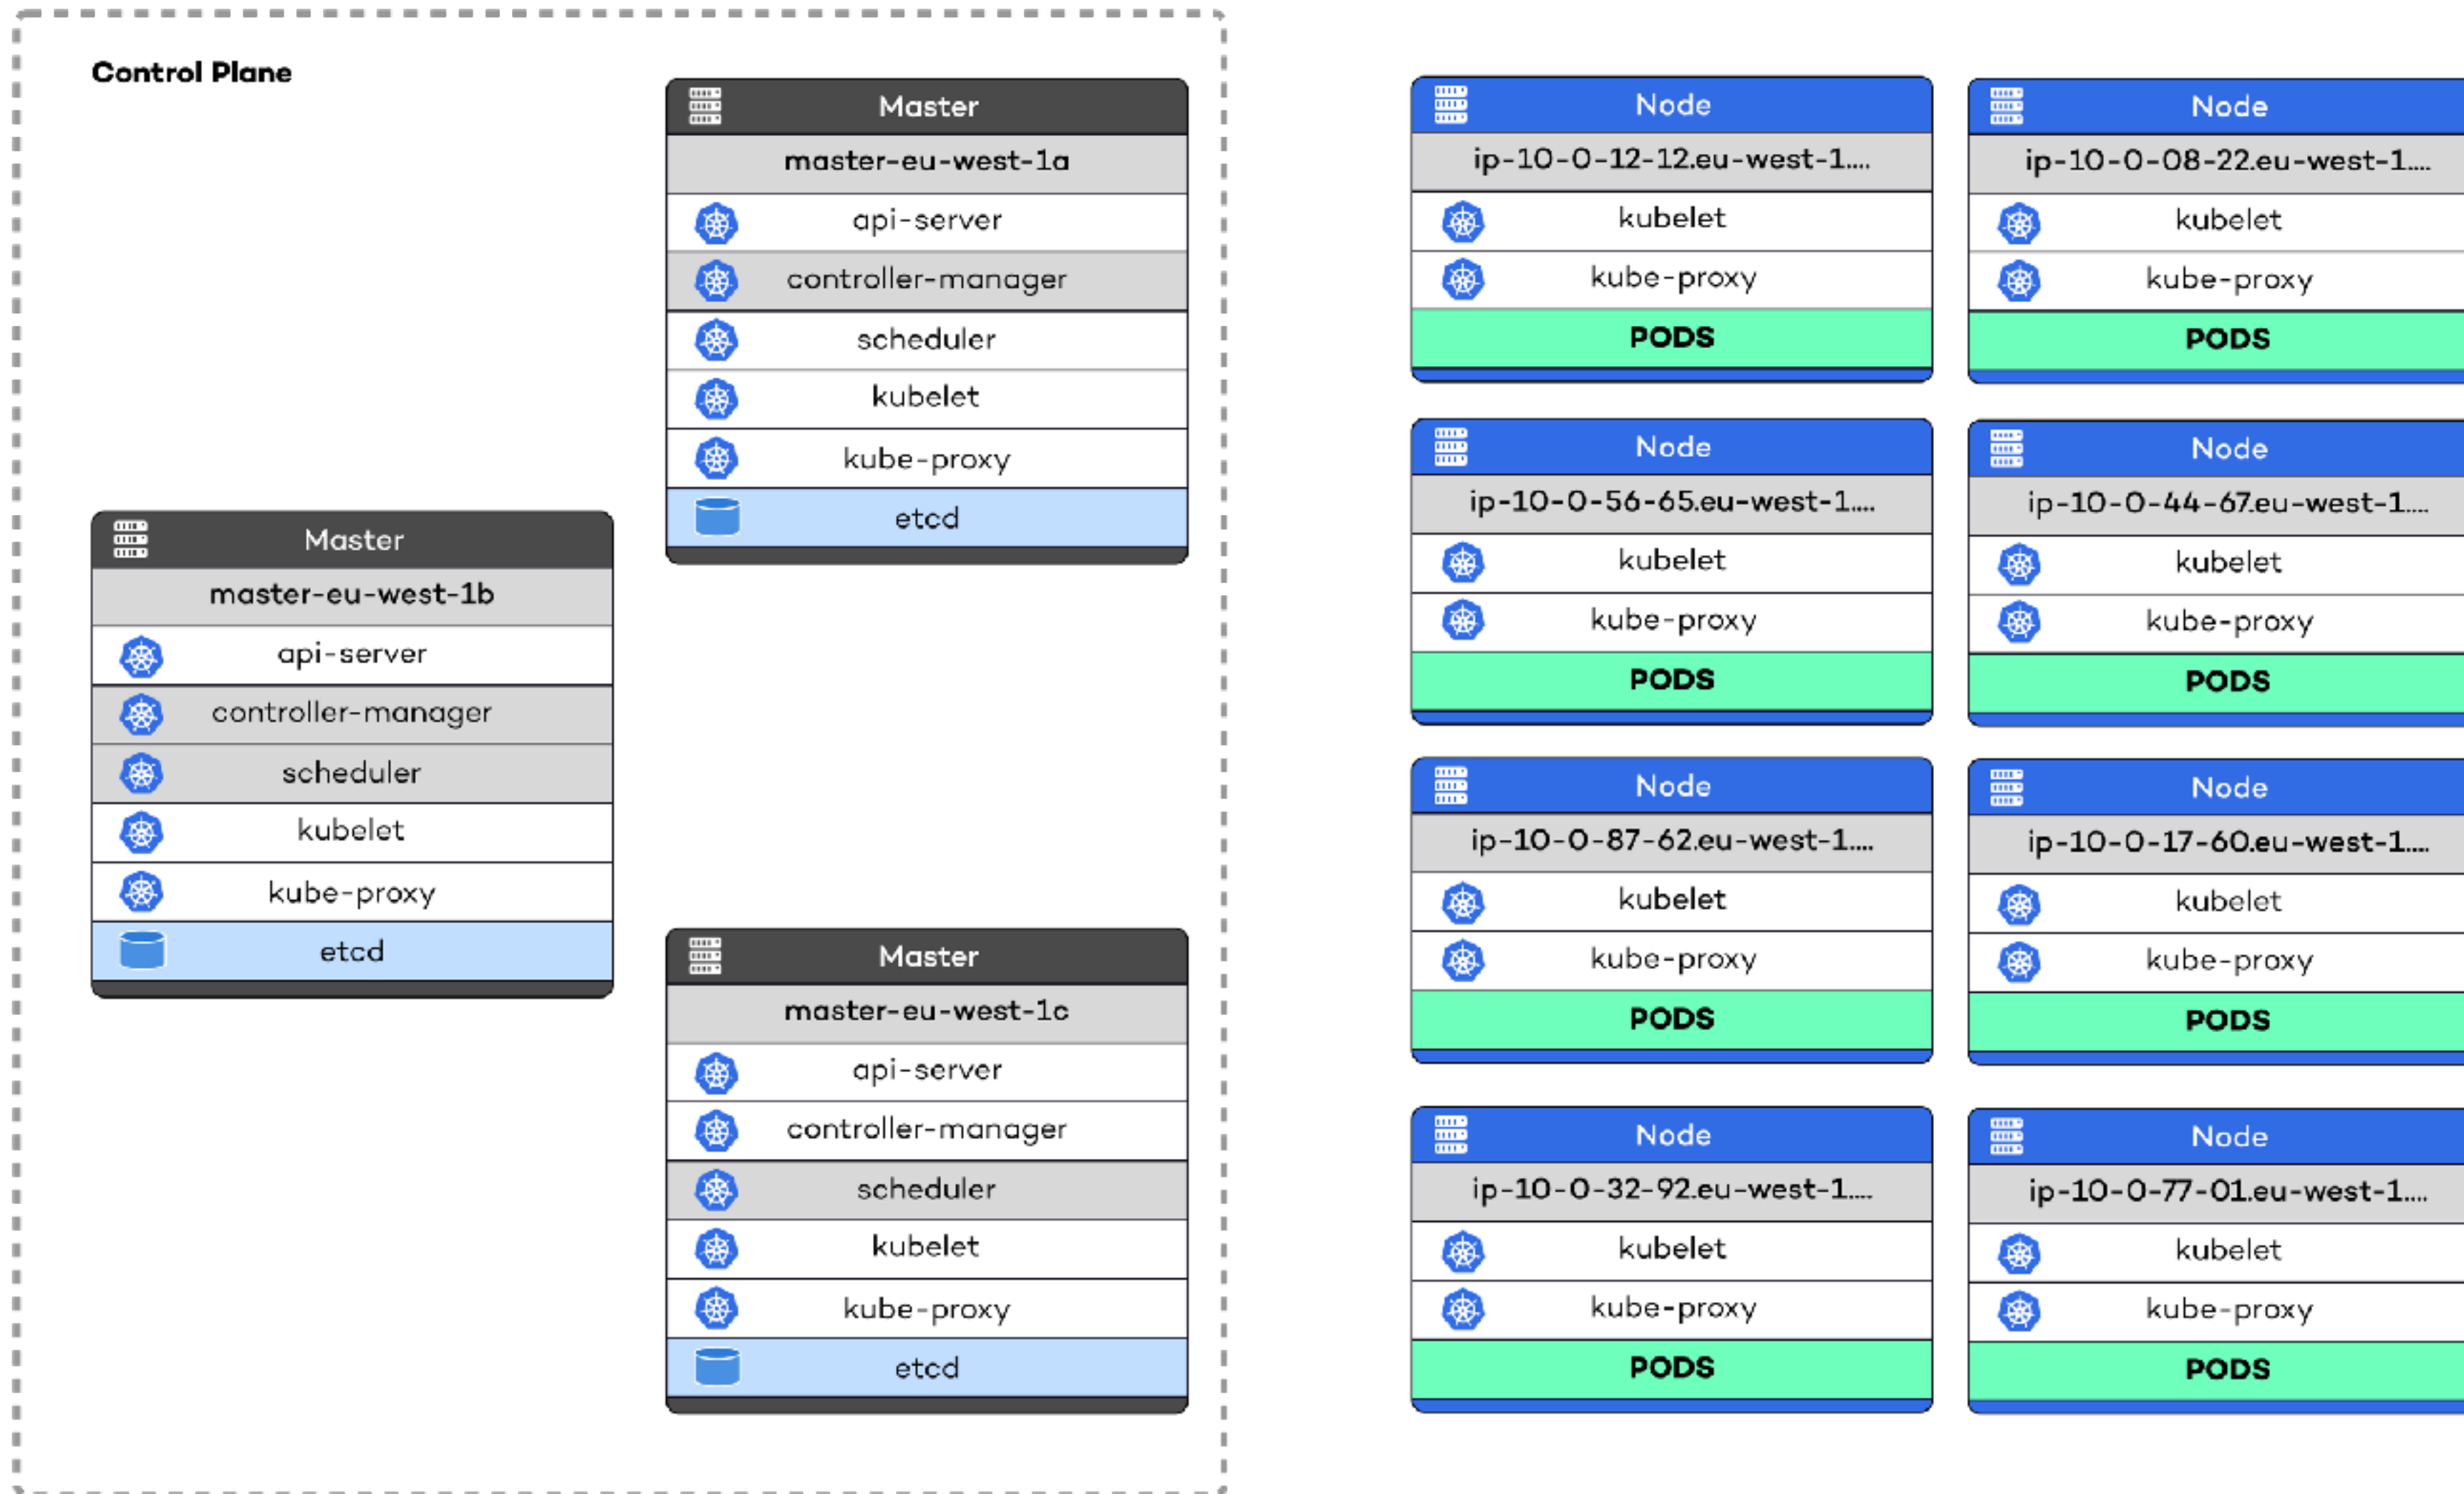
Kubernetes at a high-level



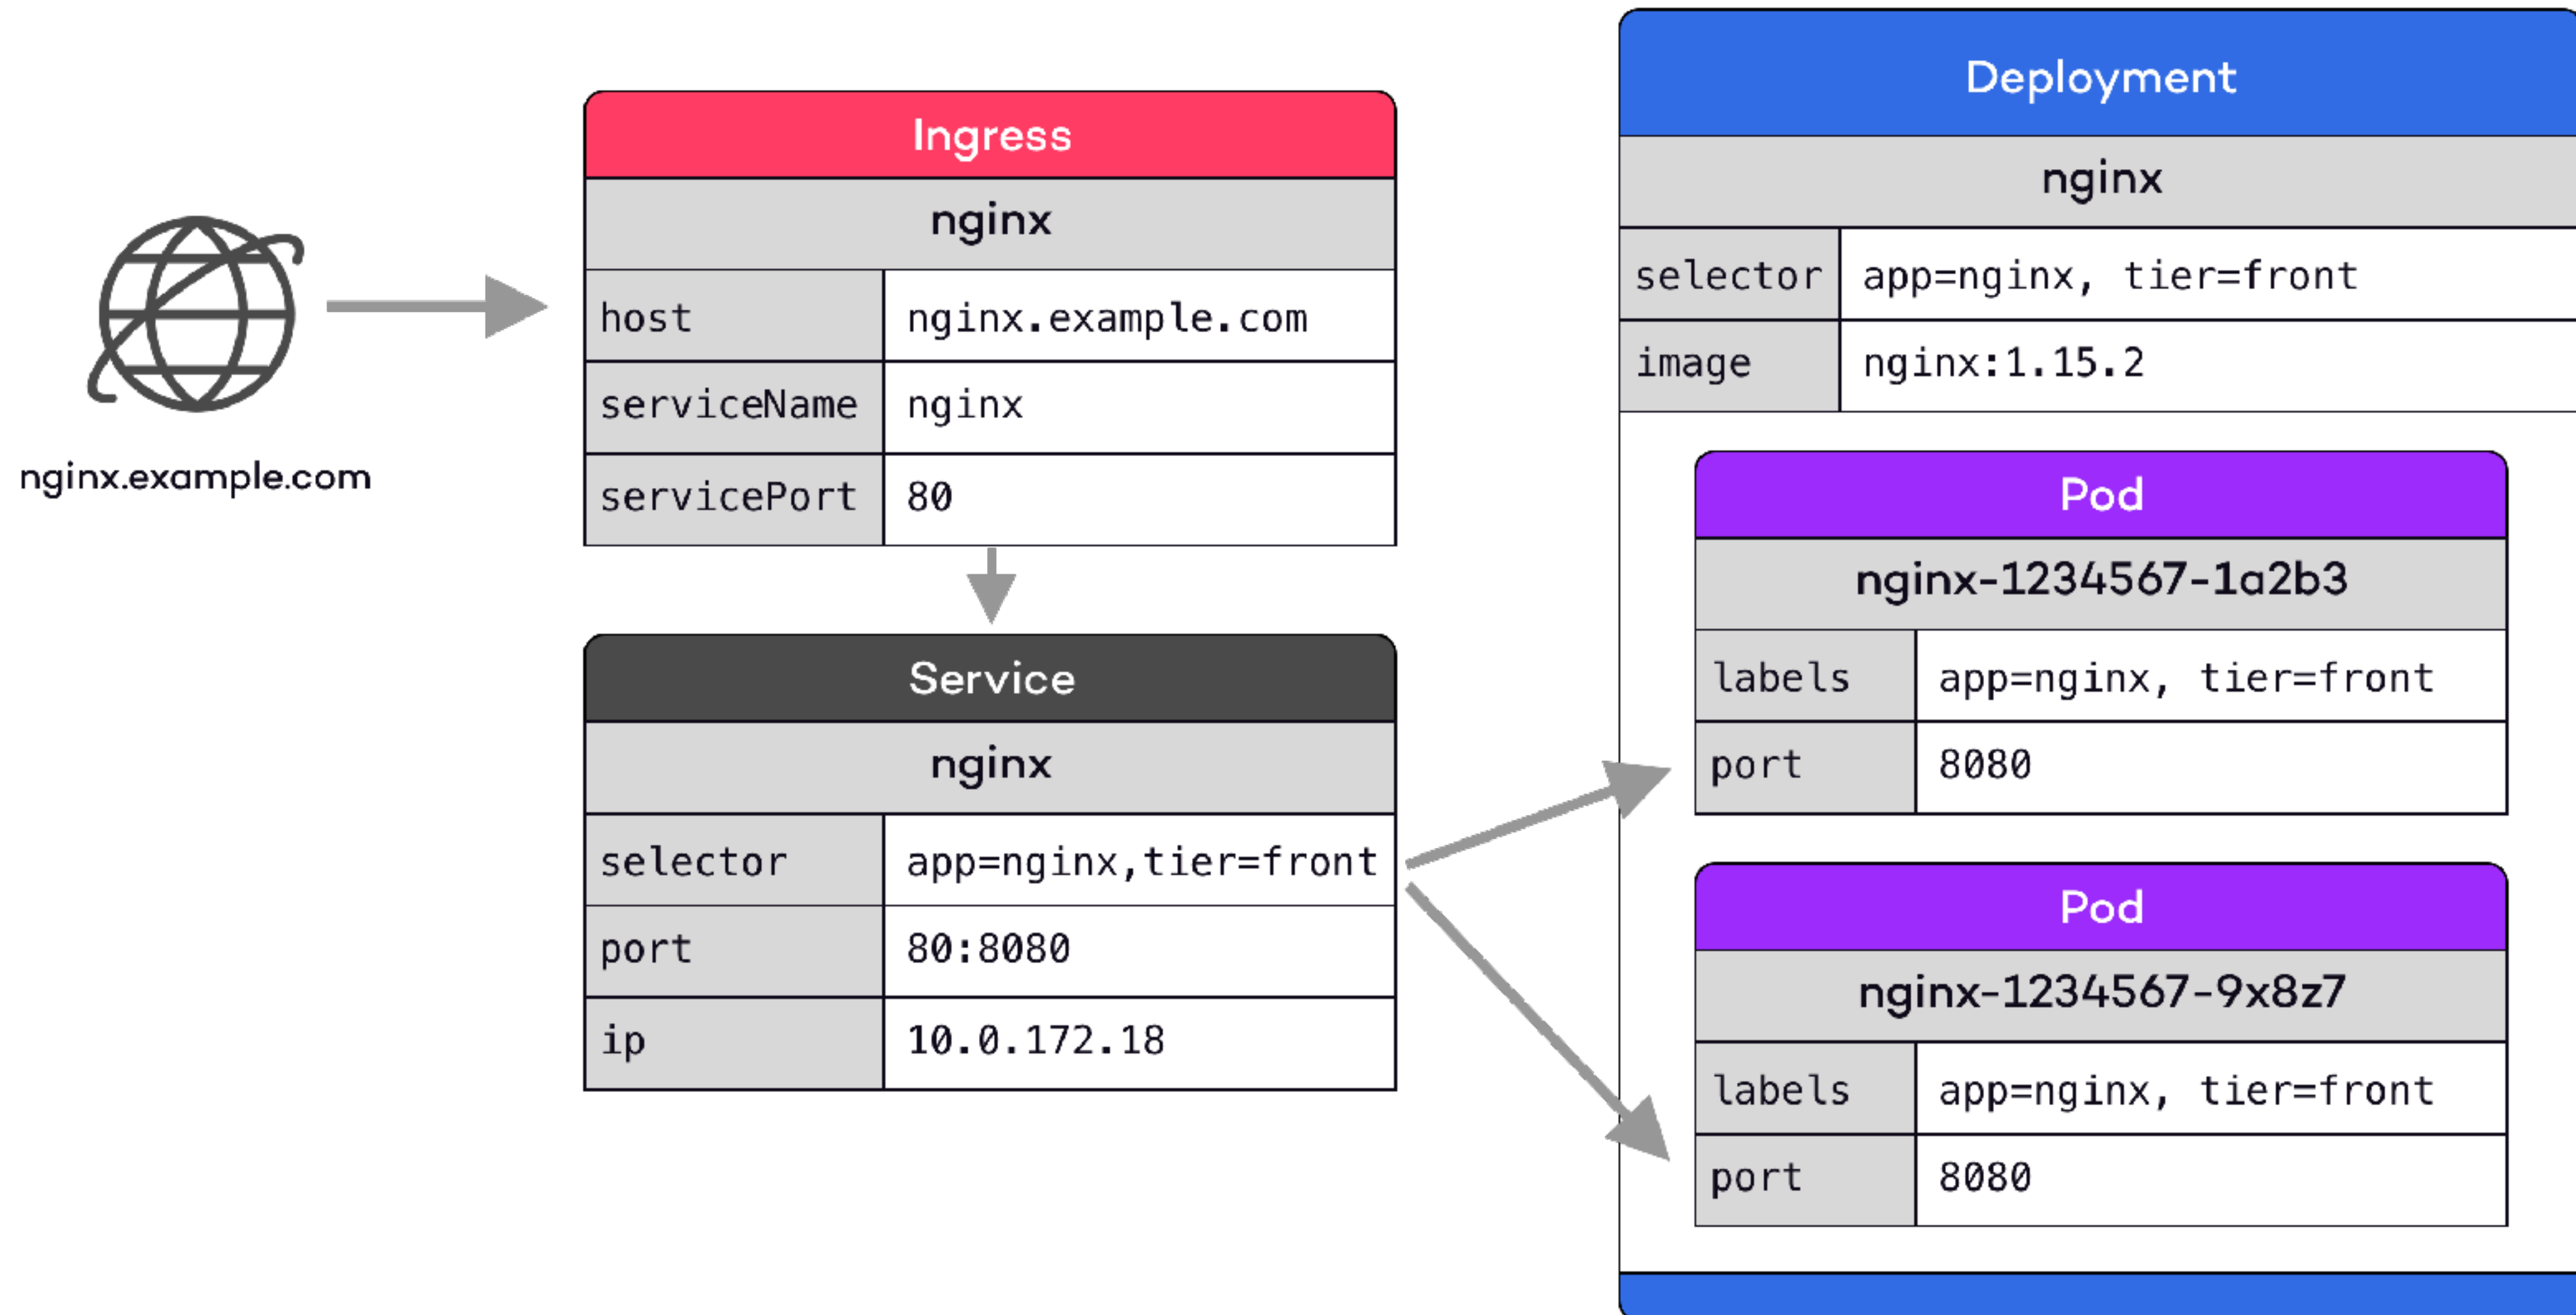
Highly Available Kubernetes



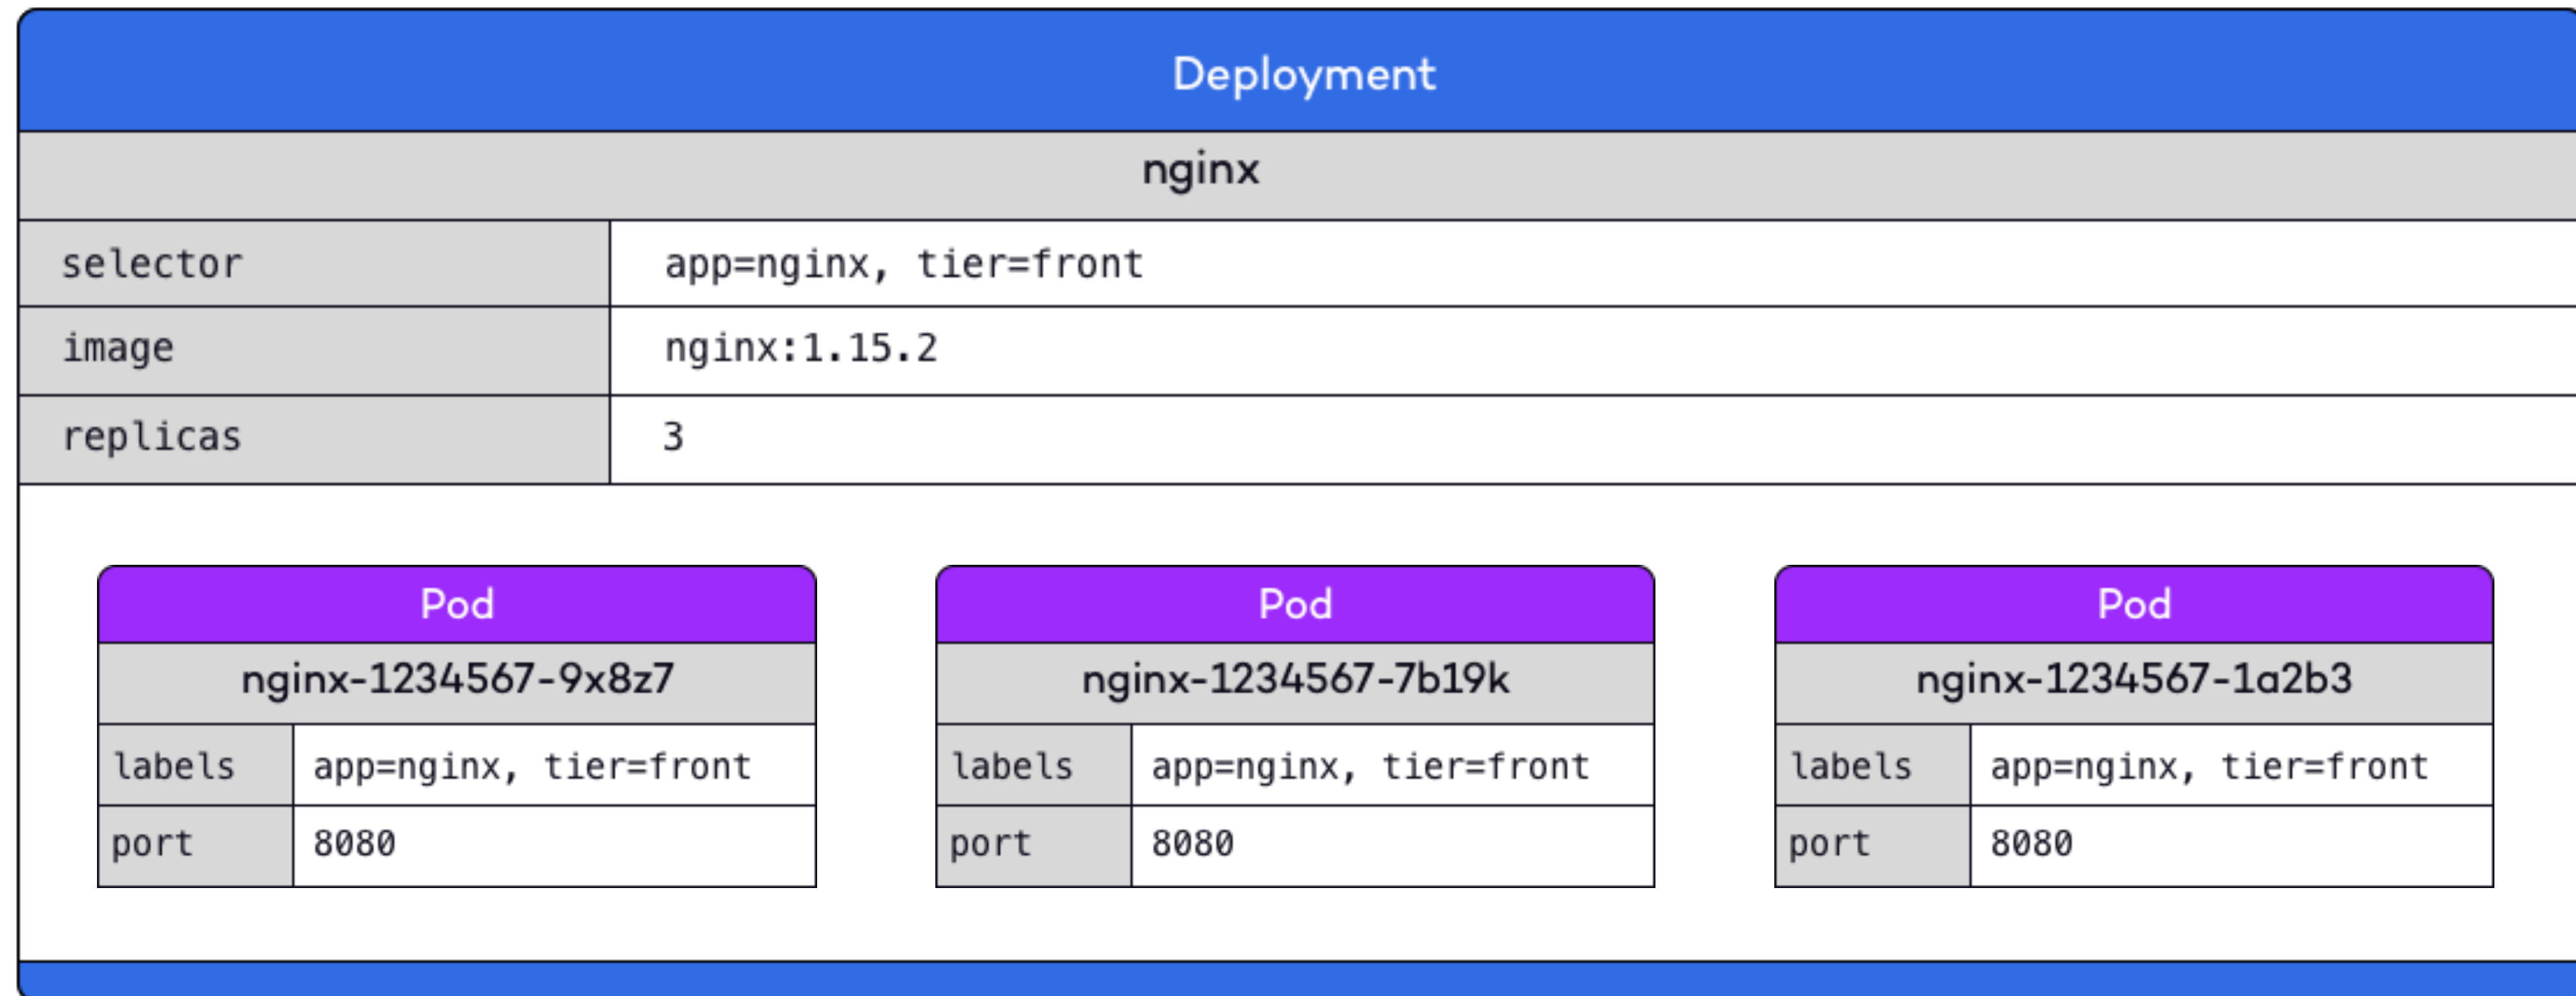
Highly Available Kubernetes



Core Concepts



Declarative Reconciliation Loop



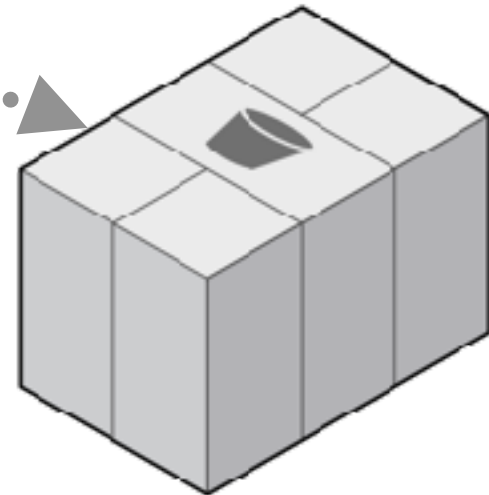
Kubernetes at Lunar Way

- Kubernetes in production since March 2017
- Three clusters at the moment (dev, staging, prod)
- Kubernetes Operations (Kops) with quite a lot of configuration
- Production environment is a multi-master highly available cluster
- Started at Kubernetes 1.5 and are now at 1.9.6

Kops

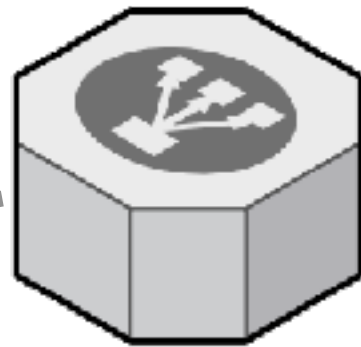
```
kops create cluster \  
  --name dev.example.com \  
  --state s3://some-s3-bucket \  
  --node-count 3 \  
  --zones eu-west-1a,eu-west-1b,eu-west-1c \  
  --master-zones eu-west-1a,eu-west-1b,eu-west-1c \  
  --dns private \  
  --node-size m4.large \  
  --master-size m4.large \  
  --topology private \  
  --networking weave \  
  --yes
```

kops

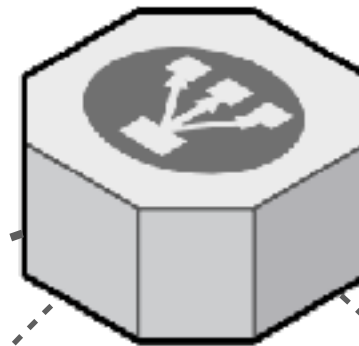


s3://k8s-test-state-store

bastion.k8s.test.lunarway.com



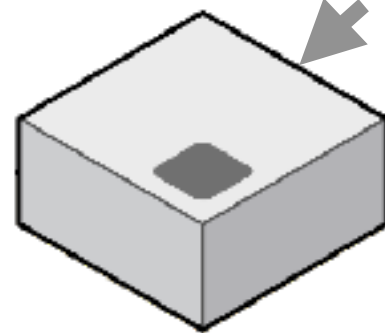
api.k8s.test.lunarway.com



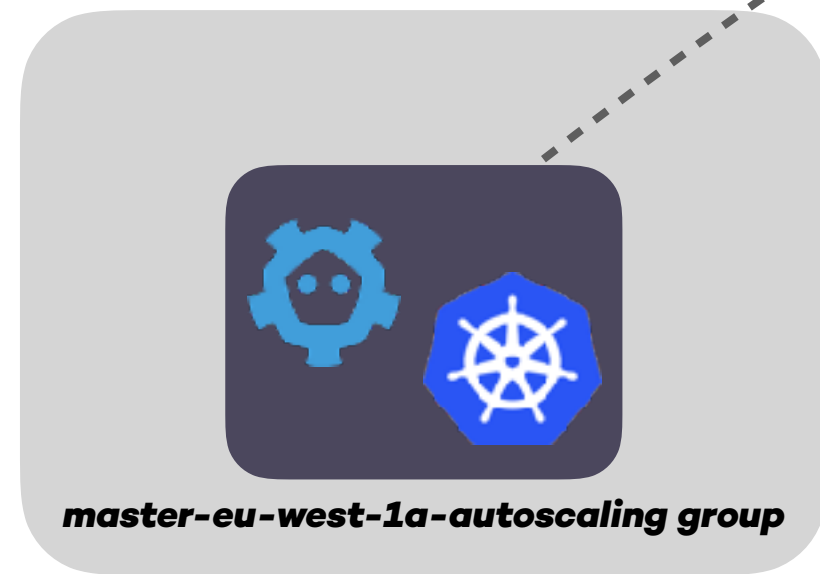
kubectl

VPC

public



private



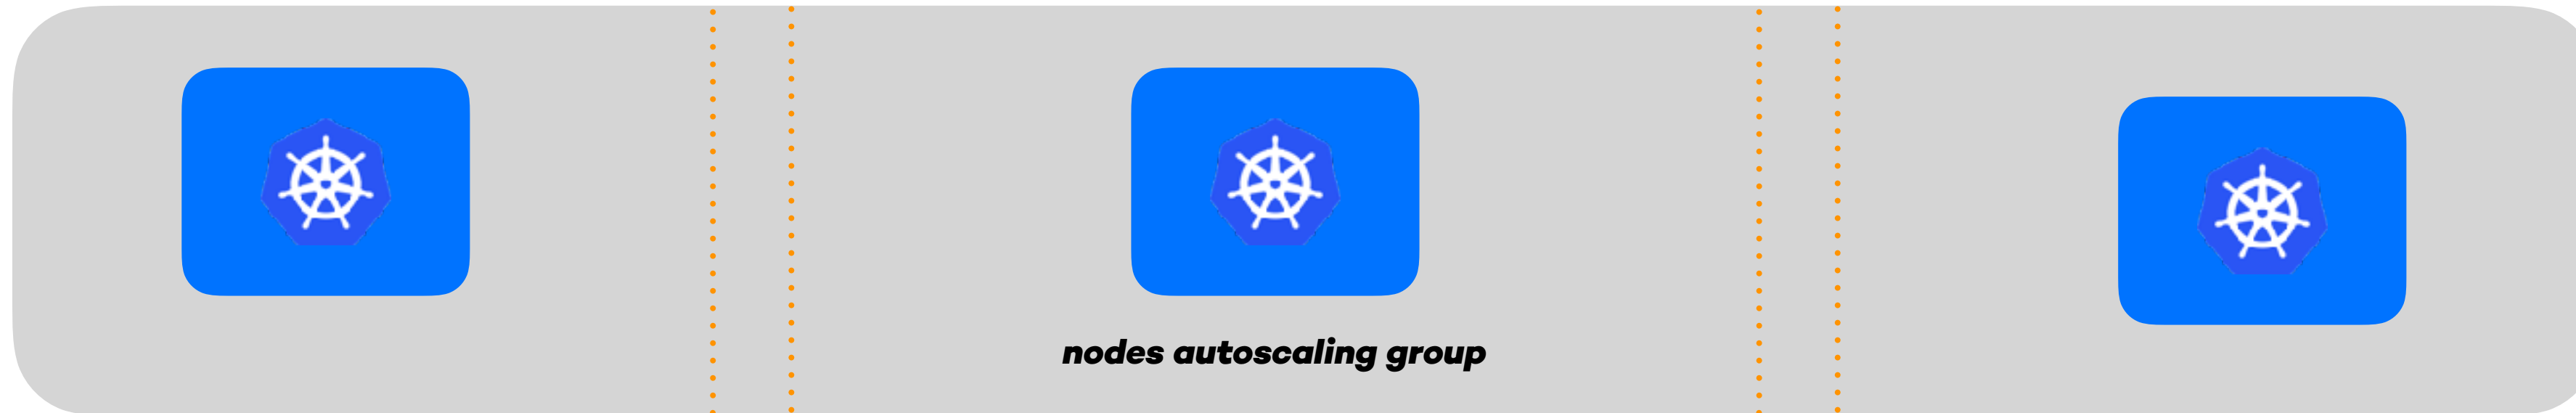
master-eu-west-1a-autoscaling group



master-eu-west-1b-autoscaling group



master-eu-west-1c-autoscaling group



nodes autoscaling group

eu-west-1a

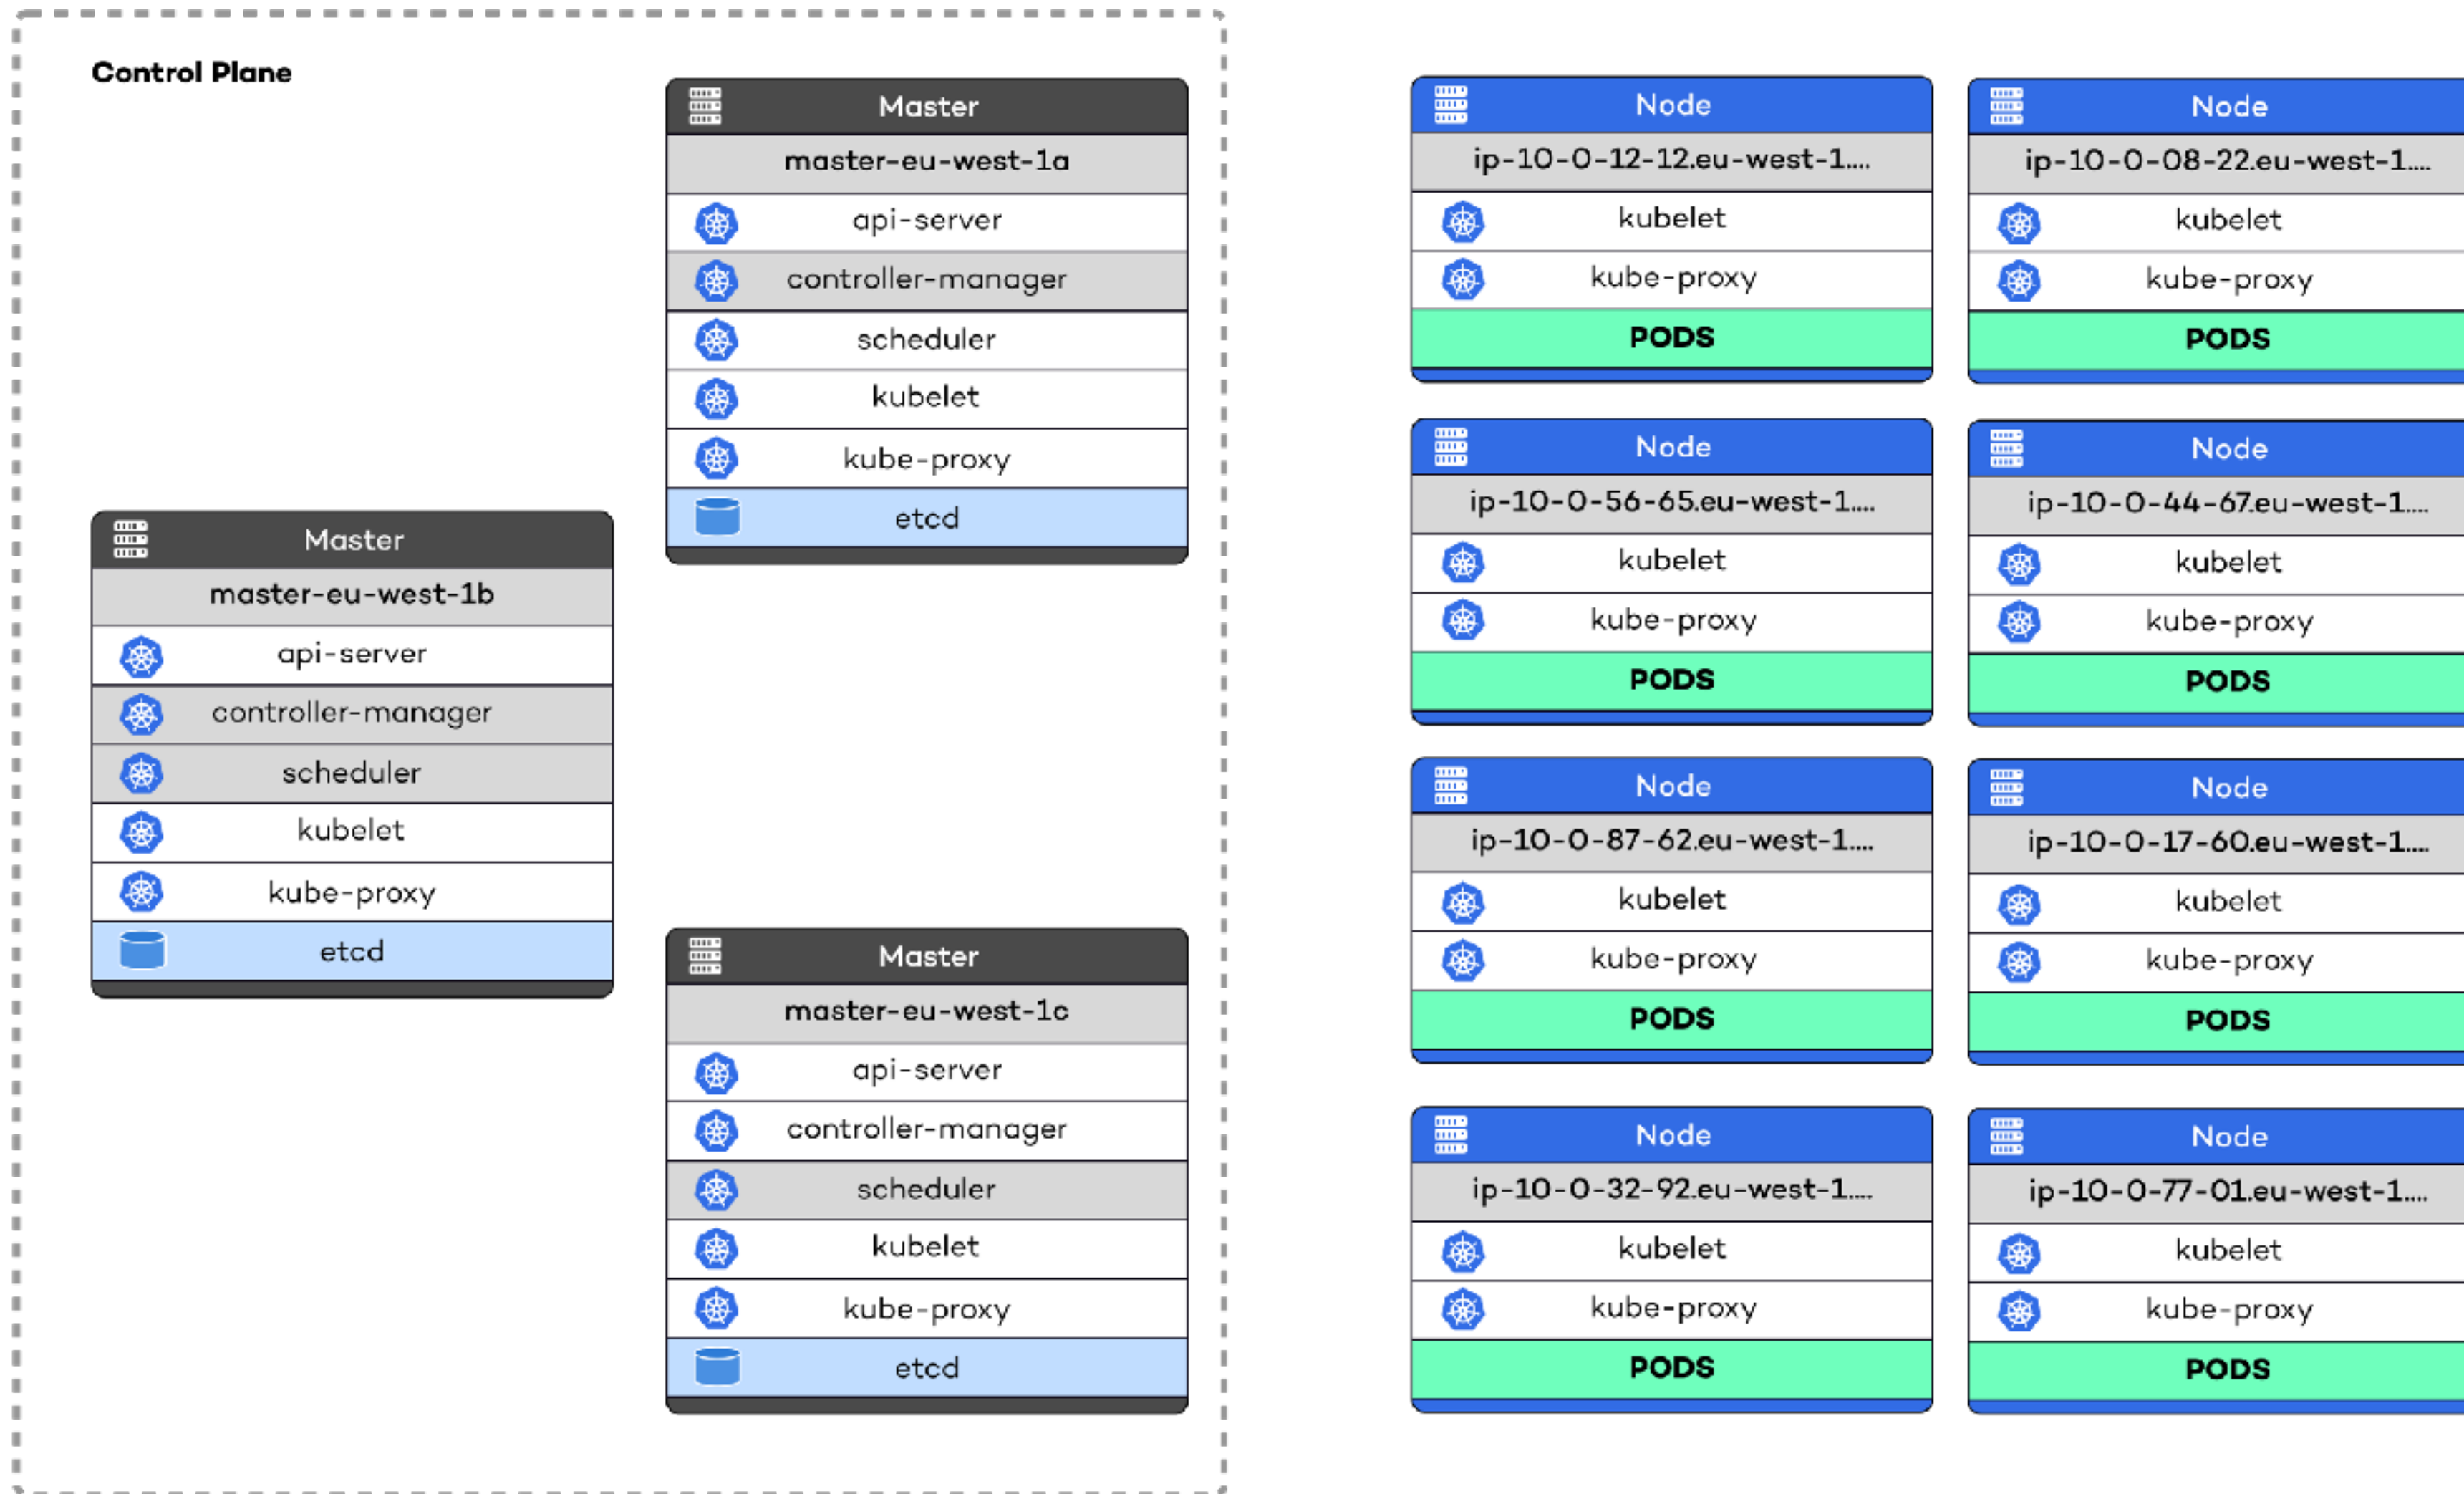
eu-west-1b

eu-west-1c

**Let's spin up
a cluster!**



Highly Available Kubernetes



Declarative Cluster Spec

```
apiVersion: kops/v1alpha2
kind: Cluster
metadata:
  name: k8s.test.lunarway.com
spec:
  api:
    loadBalancer:
      type: Public
  authorization:
    rbac: {}
  channel: stable
  cloudProvider: aws
  configBase: s3://somebucket/k8s.test.lunarway.com
  dnsZone: DNSZONE
  etcdClusters:
  - etcdMembers:
    - instanceGroup: master-eu-west-1a
      name: a
      name: main
  - etcdMembers:
    - instanceGroup: master-eu-west-1a
      name: a
      name: events
  kubeAPIServer:
    auditLogMaxAge: 30
    auditLogMaxBackups: 10
```

Kops - Pros/Cons

Pros

- Very easy to spin up clusters
- Highly configurable
- Declarative cluster specifications
- Possible to output to terraform if needed

Cons

- Previously pretty bad defaults
- Release cadence is lacking a couple of months behind upstream Kubernetes

Kubernetes - Pros/Cons

Pros

- Allow us to easily deploy services independently
- Rolling back is fast
- Provides us with resilience
- Makes management of services easy and immutable

Cons

- Very complex system
- Sometimes too high velocity

Other interesting tales

- First upgrades of the production clusters caused a 30 min outage, because of network congestion fetching pods
- Some services were just moved to this dynamic environment, not handling termination very well
- Problems with kubelet increasing resource consumption

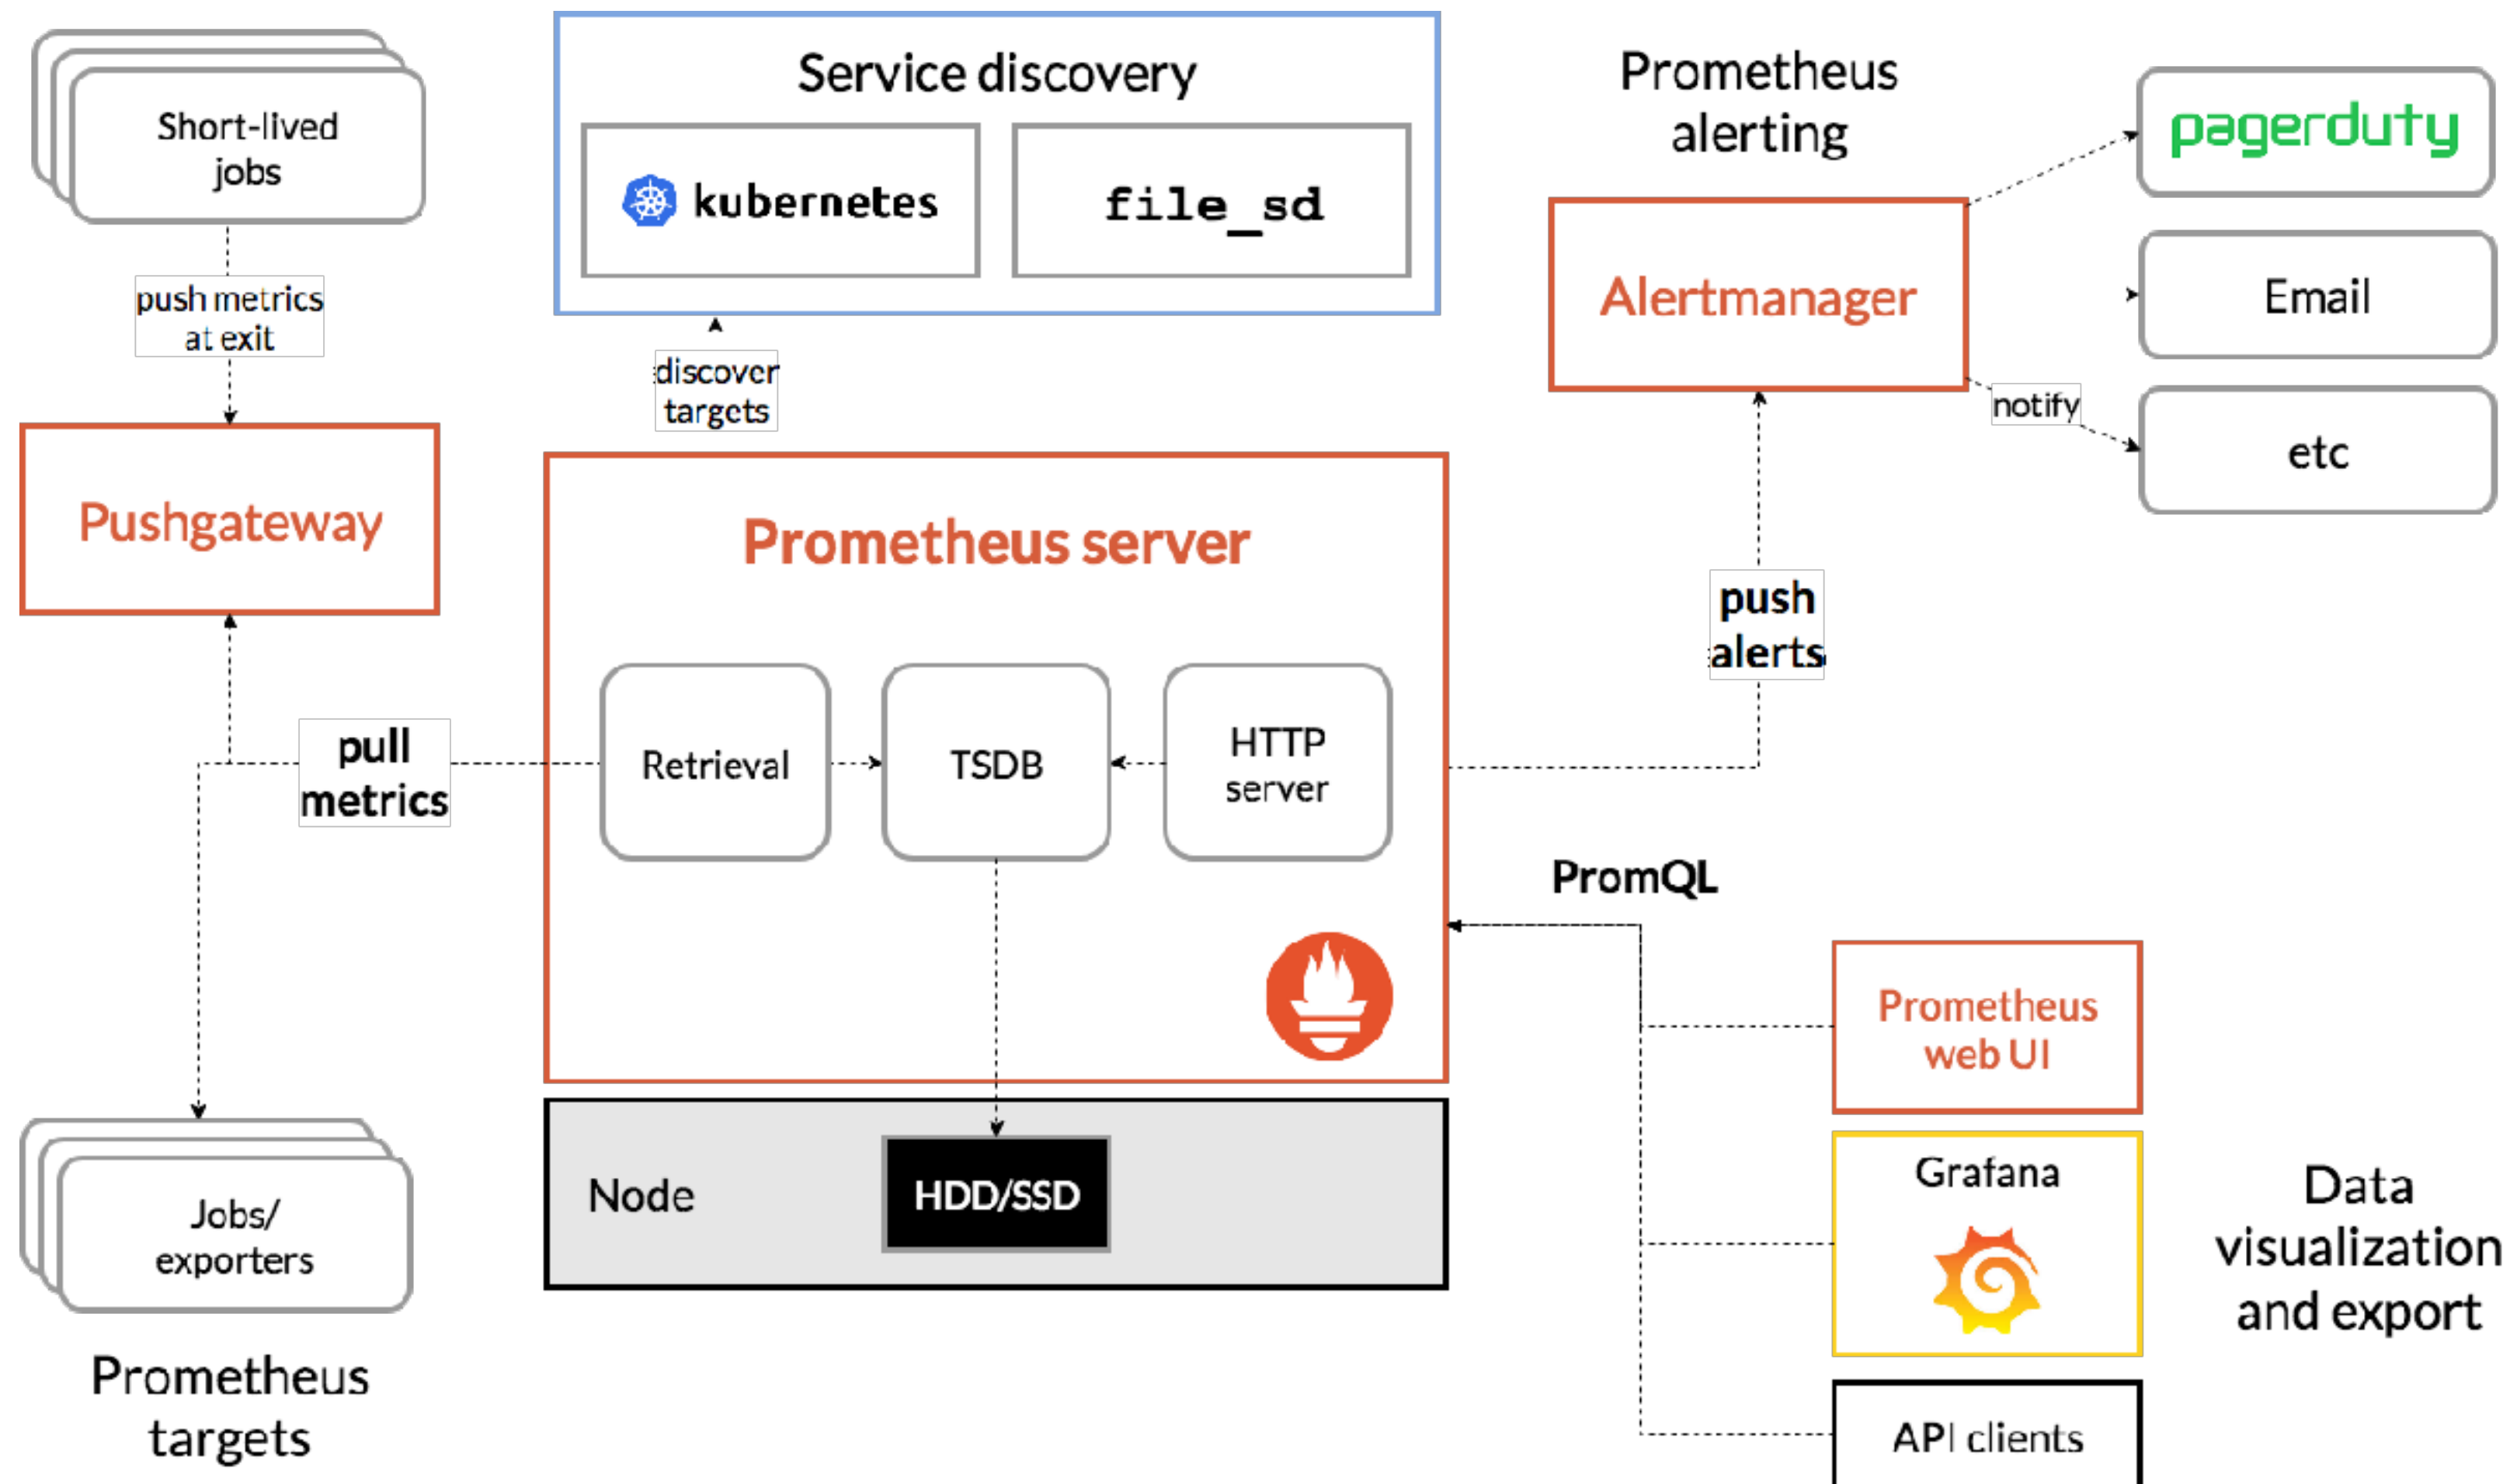


Monitoring with Prometheus

What is Prometheus?

- Monitoring system and Timeseries Database
- Instrumentation
- Metrics collection and storage
- Querying
- Alerting
- Dashboard / Graphing / Trending

Prometheus, an overview



Why Monitor?

- Analysing long-term trends
- Comparing over time or experiment groups
- Alerting
- Building dashboards to gain insights
- Conducting ad hoc retrospective analysis

**Basically, being able to find out what is broken and why...
and ... even better... know it before it impacts customers..**

**Let's set up some
monitoring of our
cluster!**



Why do we need it?

- We need insights into all the moving parts of our infrastructure
- We need to monitor error rates, saturation, latency, etc.
- We need a dynamic and powerful query language
- We need a way to get alerted when things go wrong

Prometheus at Lunar Way

- Prometheus provides insights in to all our different systems running, whether it's services running in Kubernetes or outside
- We currently run Prometheus in our Kubernetes cluster
- It runs as a simple deployment, and managed with our own set of yaml configurations

Prometheus - Pros/Cons

Pros

- Provides great insights to all of our services
- Makes it easy for developers to instrument their services
- Integrates well with many different services

Cons

- Prometheus do not support clustered setup

Other interesting tales

- We've had many internal discussions on when to use logs and when to use metrics
- Before Prometheus 2.0 we had a lot of difficulties with high memory consumption
- We write exporters internally for monitoring external partners

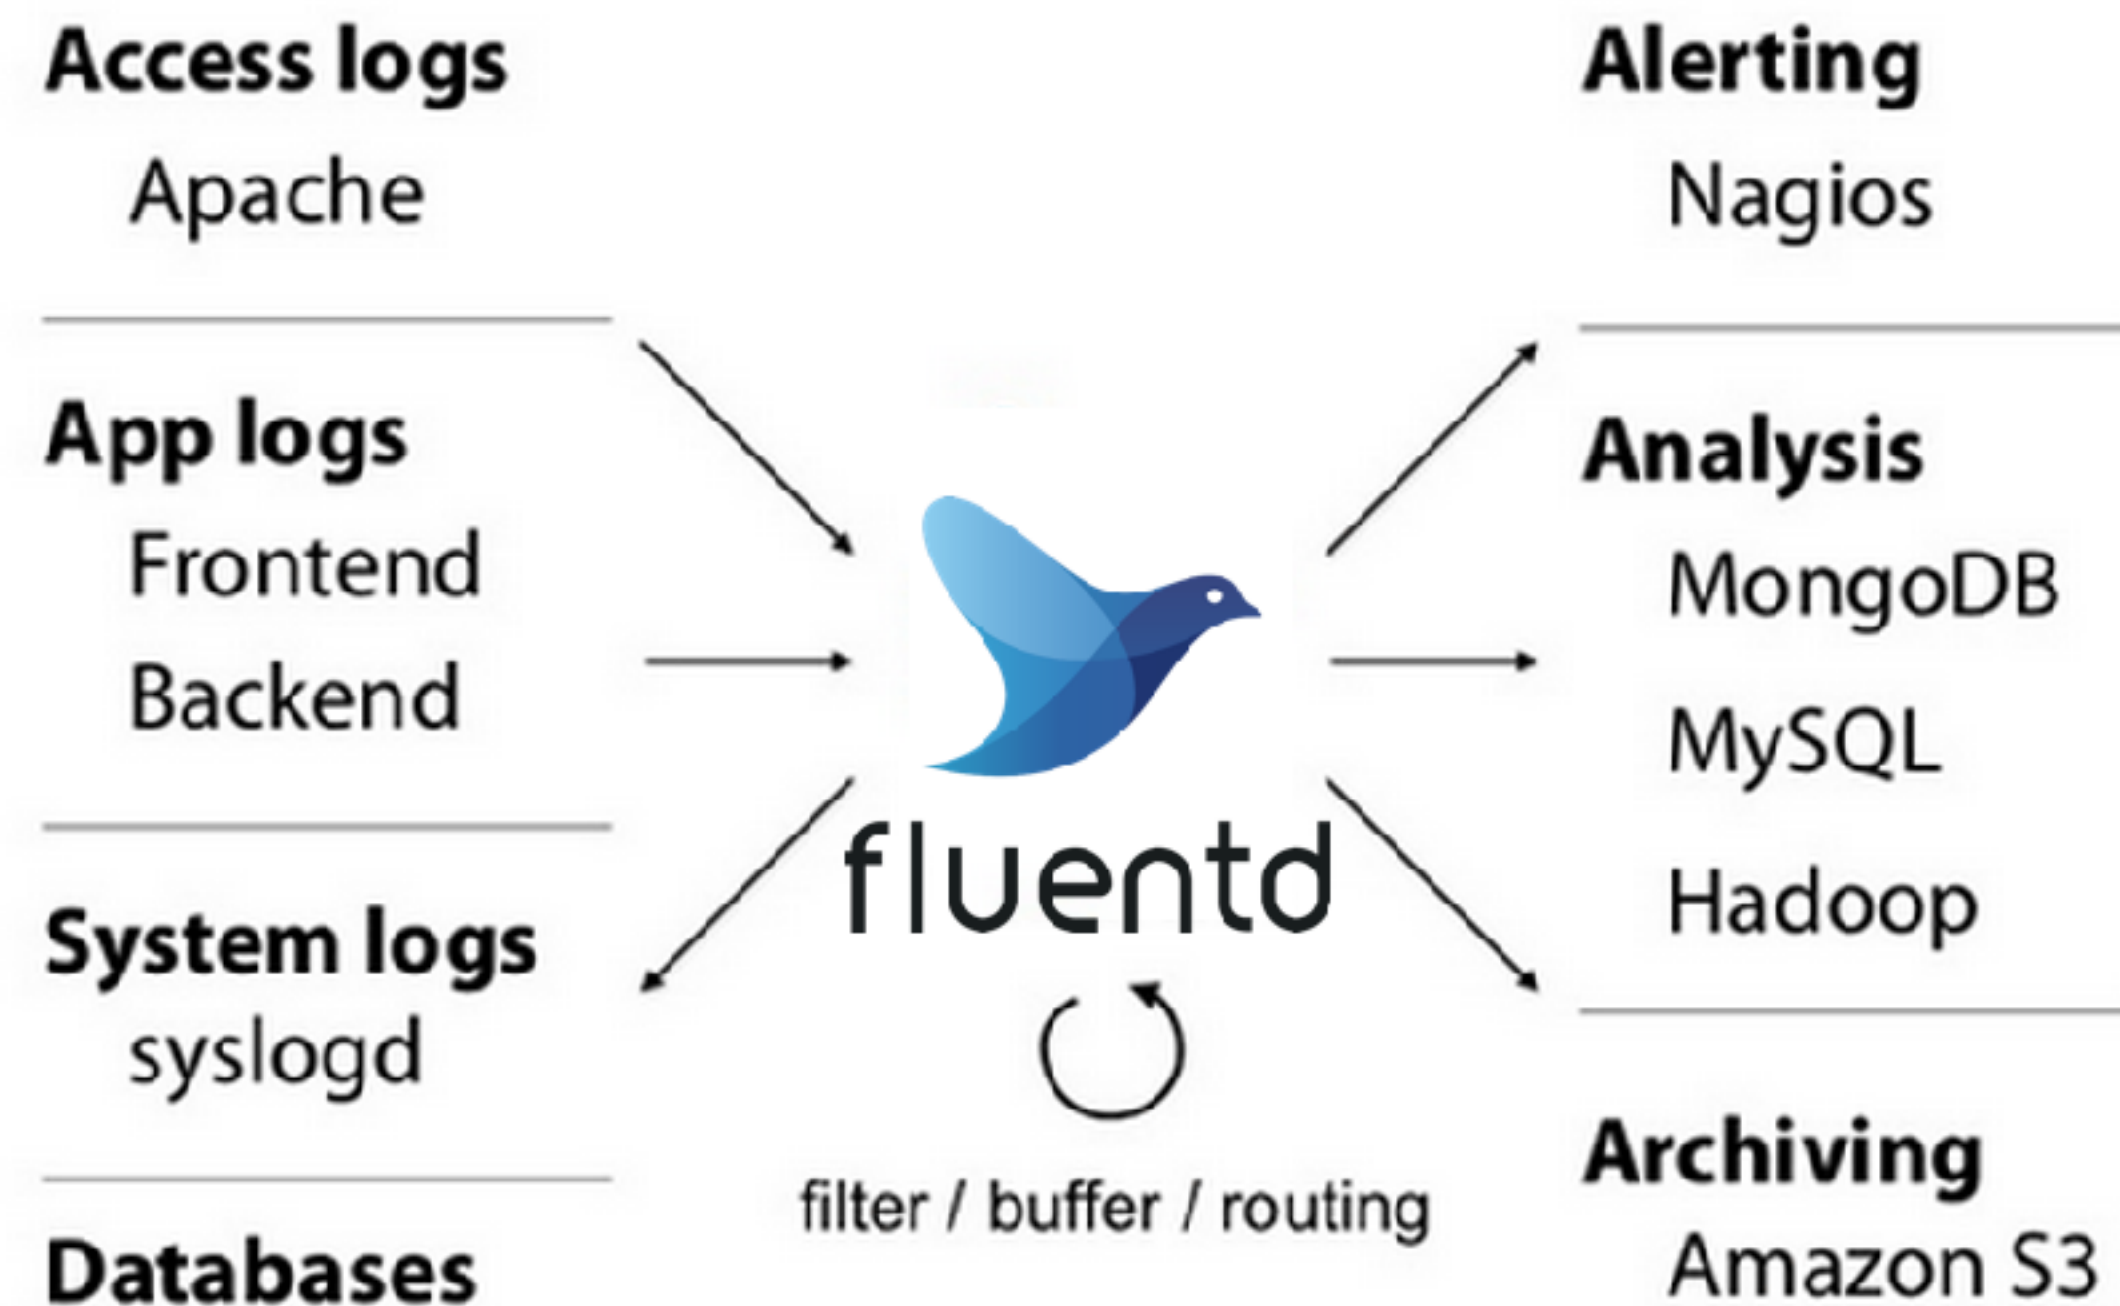
<https://www.pexels.com/photo/art-classic-contemporary-design-345046/>



Log
Management
with
fluentd and Humio

fluentd, what is it?

- Fluentd is a log collector
- Hosted by the CNCF

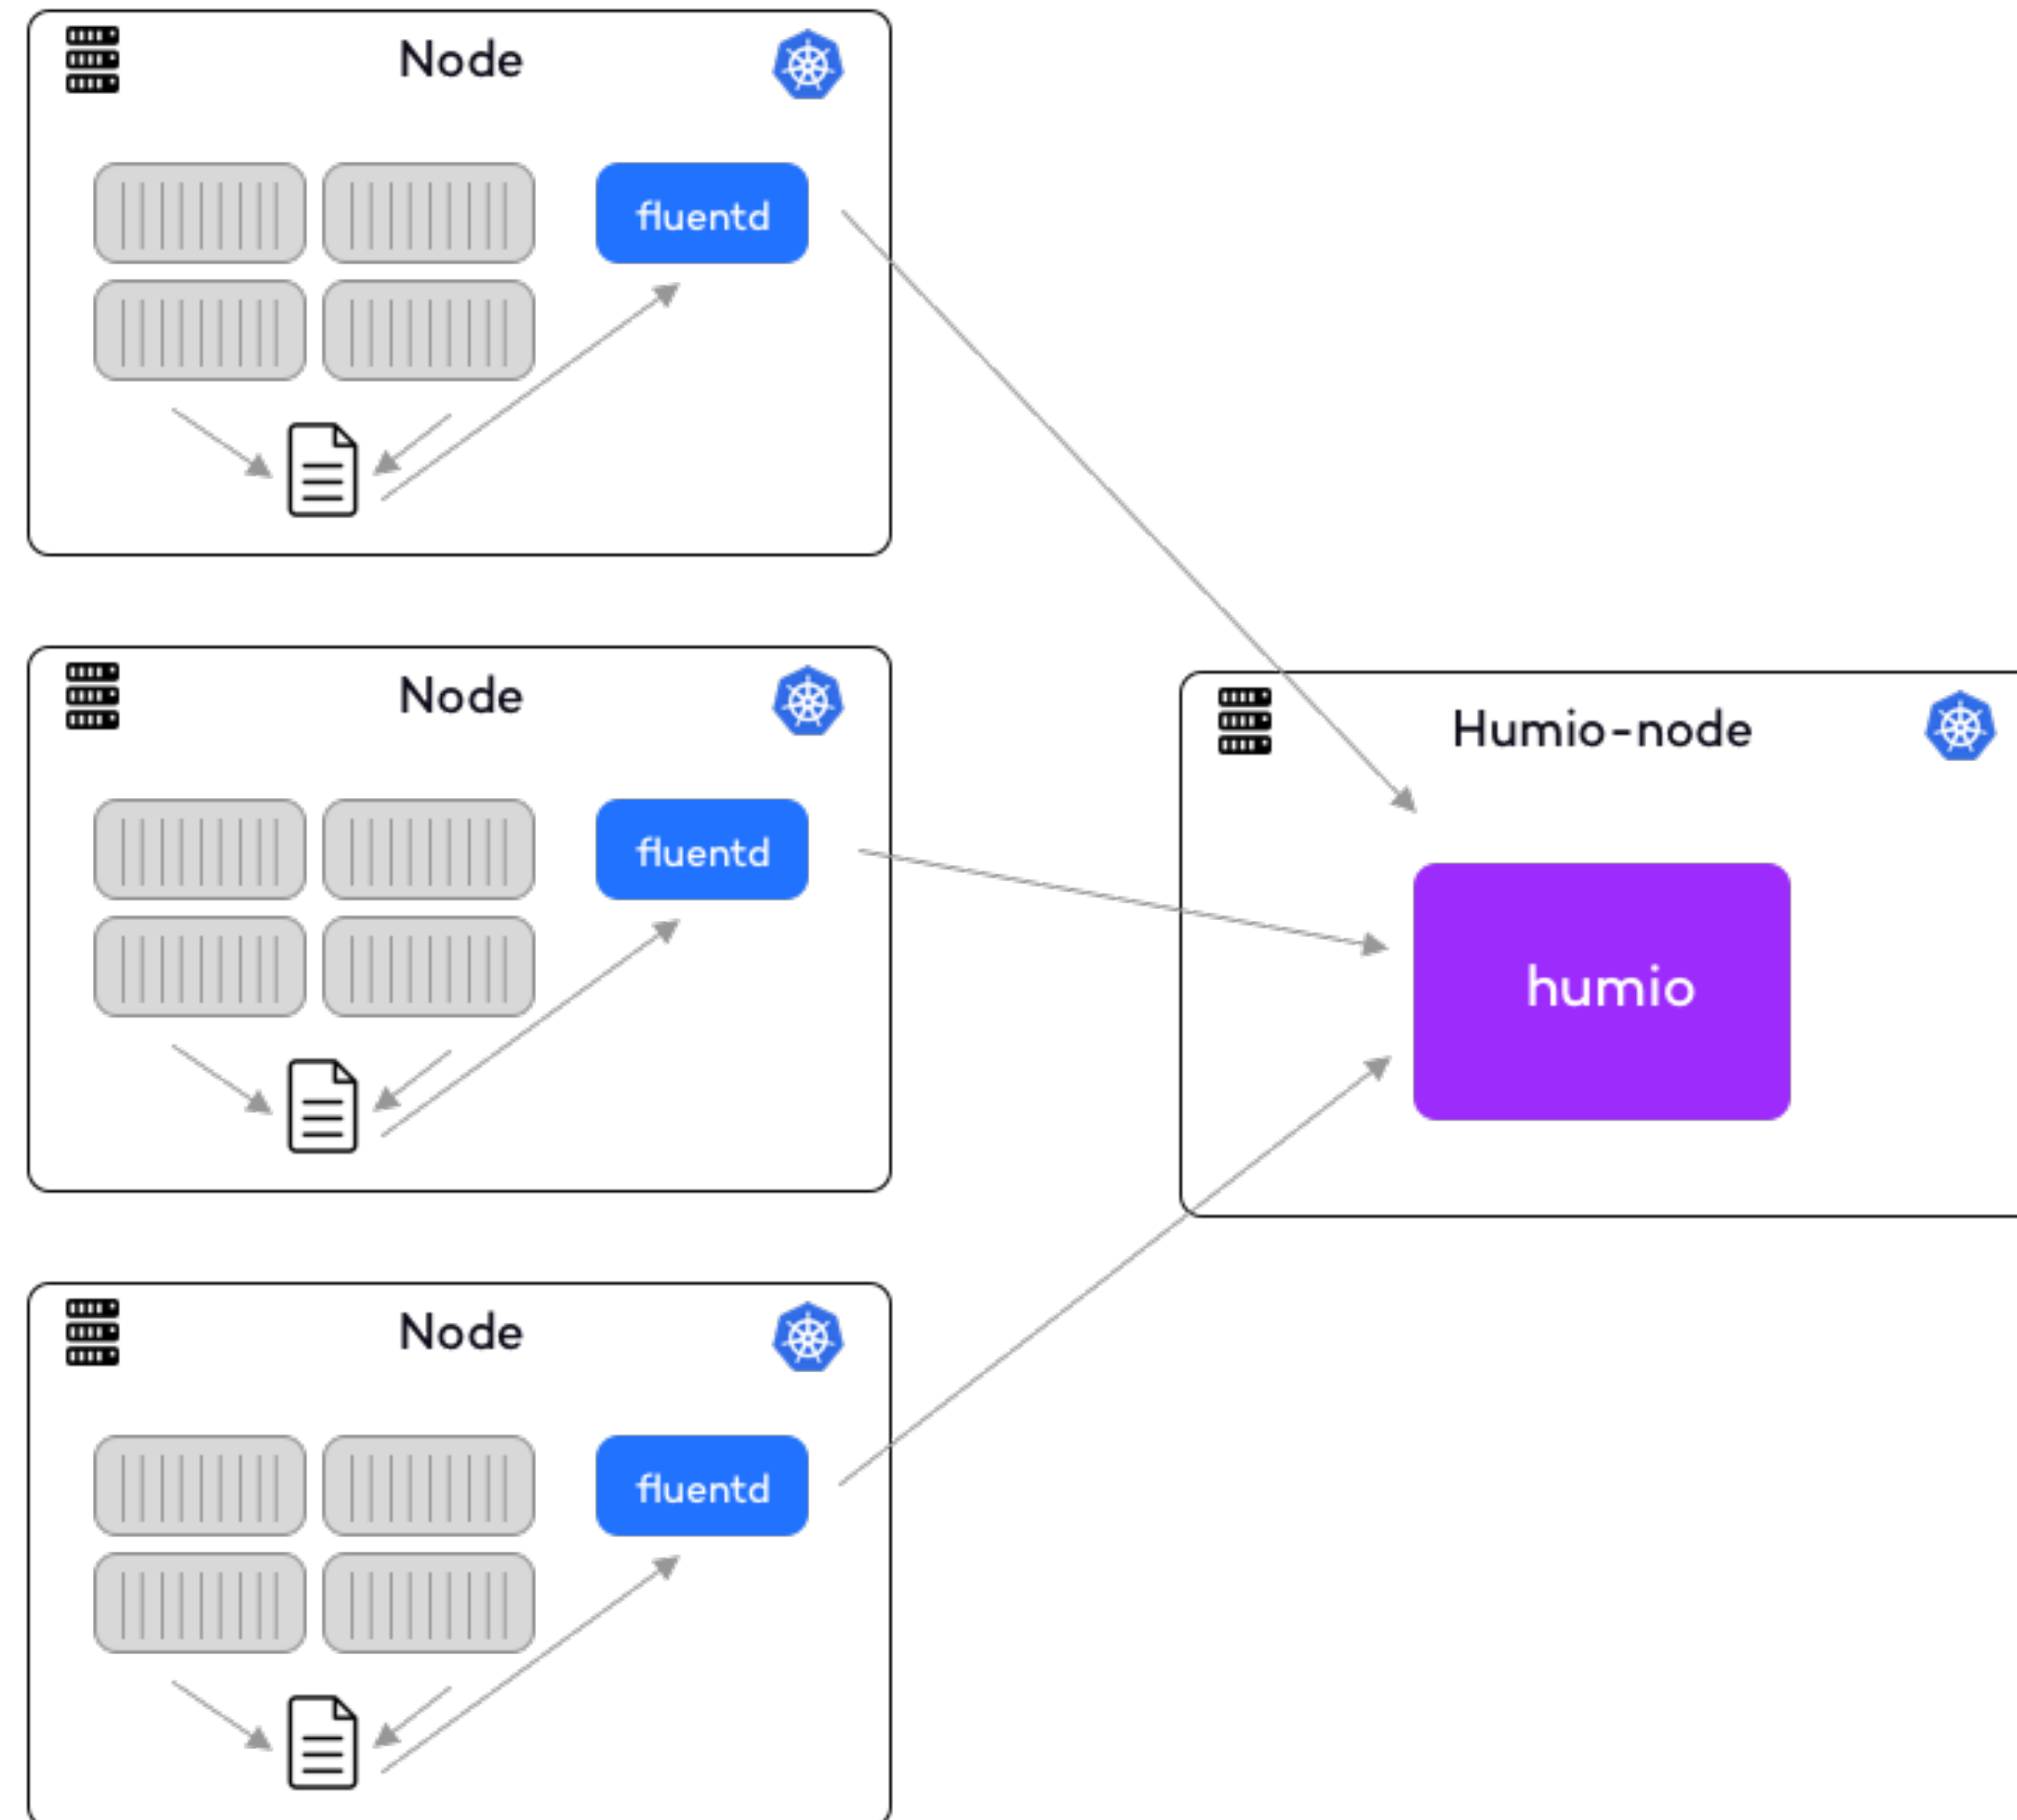


humio, what is it?

- Humio is a log management solution (unfortunately not open-source)
- Humio provides a simple and developer friendly query language for getting insights into your logs



Architecture overview



fluentd/humio - why this combo?

- fluentd is easy to get up and running
- fluentd makes it fairly easy to add additional data and parse up structured fields
- fluentd comes with a build in buffer mechanism in case of connection failures
- Humio provides an elasticsearch ingest api which works great with fluentd
- The most important reason for us choosing Humio, is it's developer friendly query language and speed

Moving towards fluent bit

- Our fluentd setup has become quite advanced, and a lot of different parsing is happening at this layer
- fluentd uses quite a lot of resources on each machine
- We have seen fluentd spike to 1.5gb memory consumption
- We wan't to move parsing further up the stack (in Humio)

For many of these reasons we are looking at fluent bit as an alternative!

**Let's start
reading our logs!**



What's next?

- Utilising Custom Resource Definitions and a controller to do Release Management
- Services Mesh
- Adopting more Operators to ease operations of e.g. Prometheus.
- Provide a FaaS on top of Kubernetes
- Extend Kubernetes with virtual-kubelet for additional serverless resources

Wrapping up

Key takeaways if entering kubernetes land

Kubernetes is complex and has a steep learning curve, but it enables so many possibilities

Monitoring and alerts are very important in such a dynamic environment, but be aware of alert fatigue

Read your logs and make them easily accessible for all your developers

***lunar
way***[®]

Questions?

Thank you!

@phennex

