



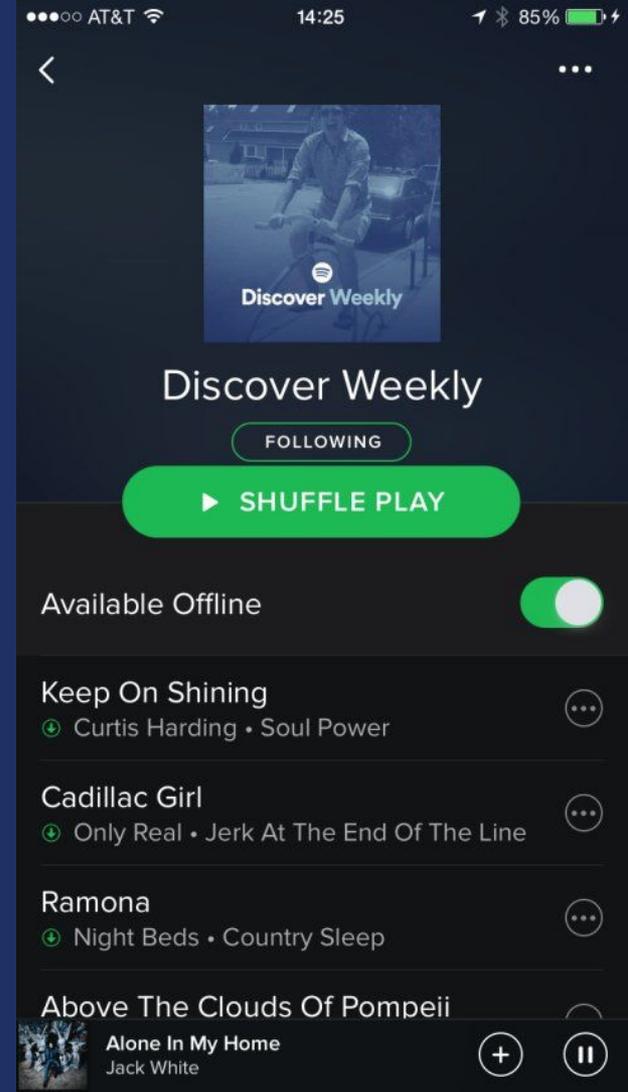
# Making an Internal Kubernetes Offering Generally Available

James Wen - [jameswen@spotify.com](mailto:jameswen@spotify.com)  
@rochesterinnyc

KubeCon + CloudNativeCon North America 2019  
November 19th, 2019

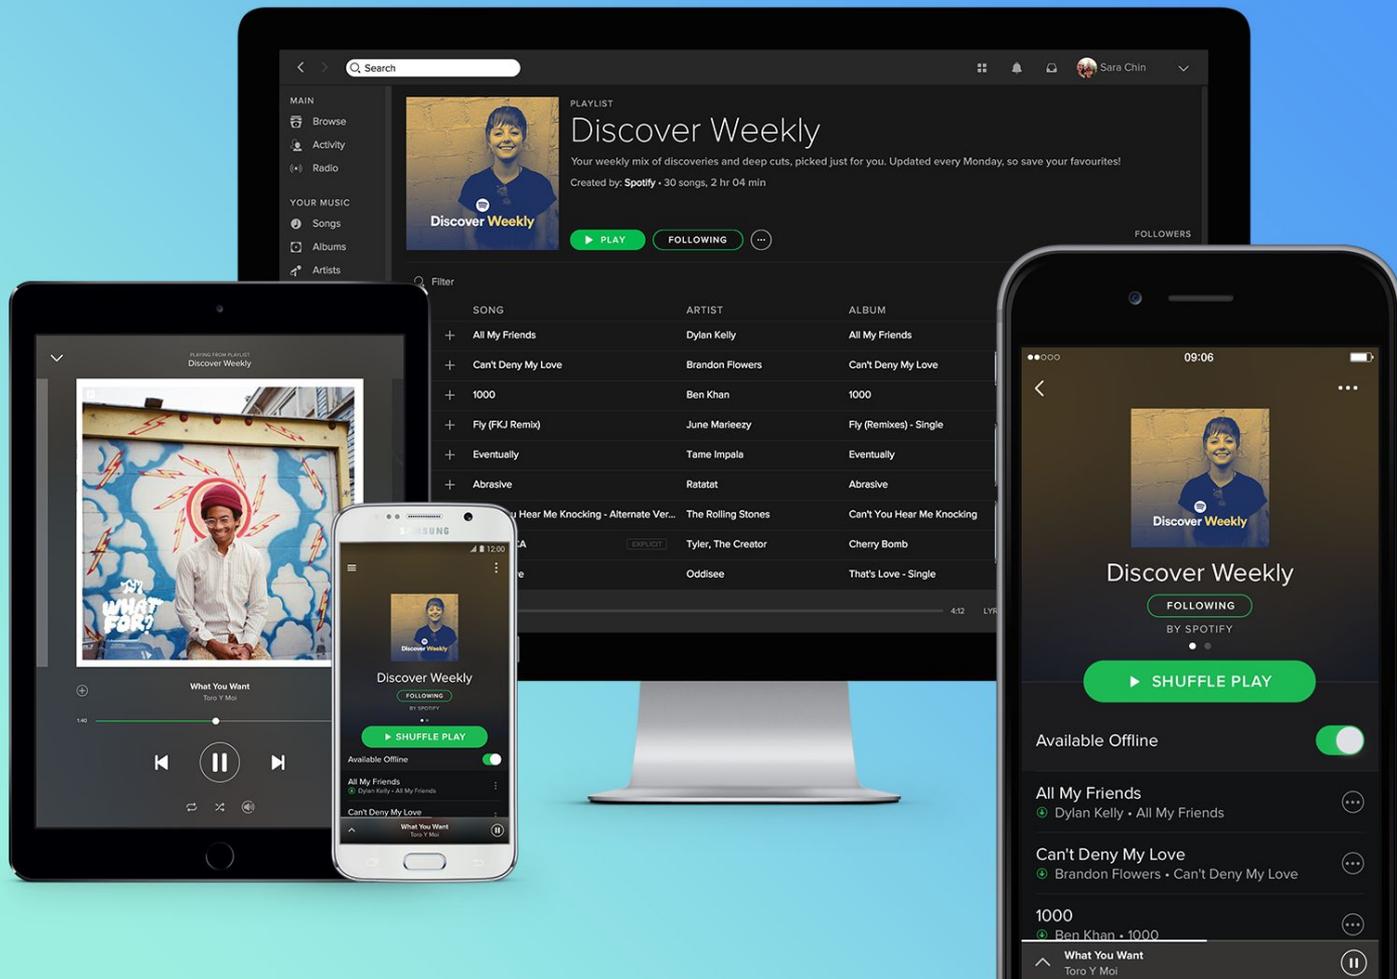
# About Me

- Senior Site Reliability Engineer
- Once Provisioning + DNS
- Core Infra: Deployment + Kubernetes Migration



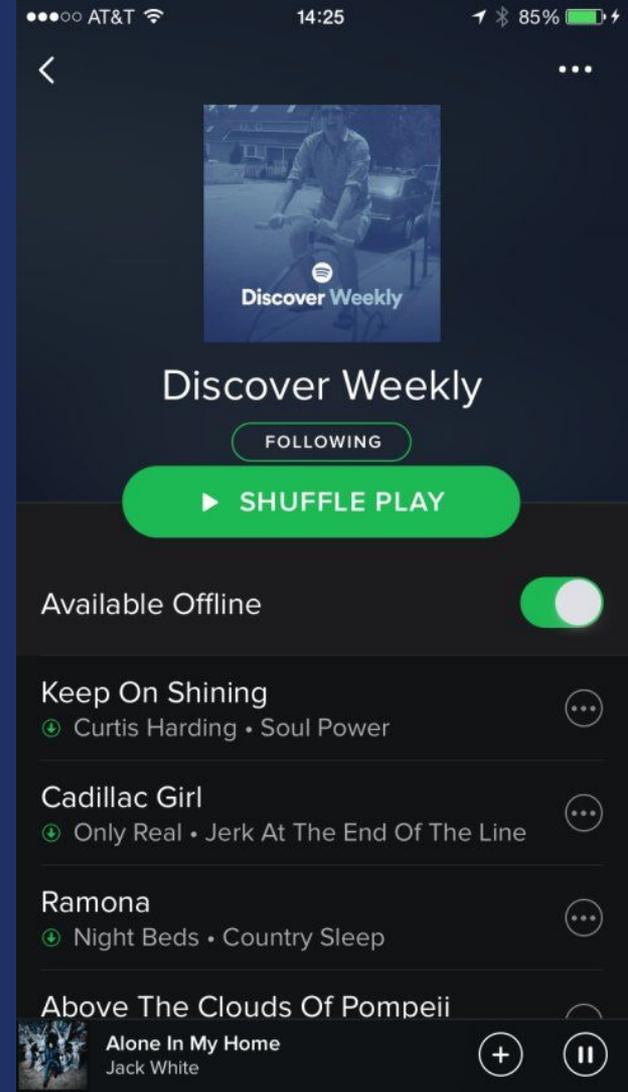
# Agenda

- Intro Spotify/Ops/Runtime
- GA Journey
- Learnings/Takeaways



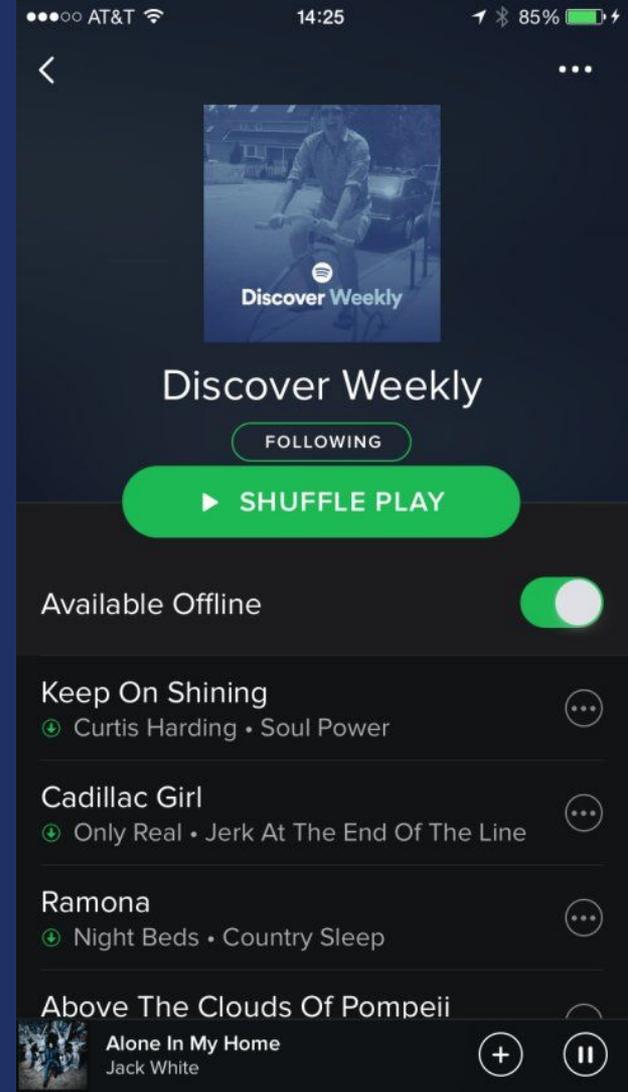
# About Us

- 248M+ active users
- 113M+ subscribers
- 50M+ songs
- 500K+ podcasts
- 3B+ playlists



# About R&D

- 280+ teams
- 1500+ engineers
- 1200+ microservices
- 8M+ client requests/s
- Half trillion events captured per day



# Operations at Spotify



Complete Team  
Autonomy



Ops in Teams

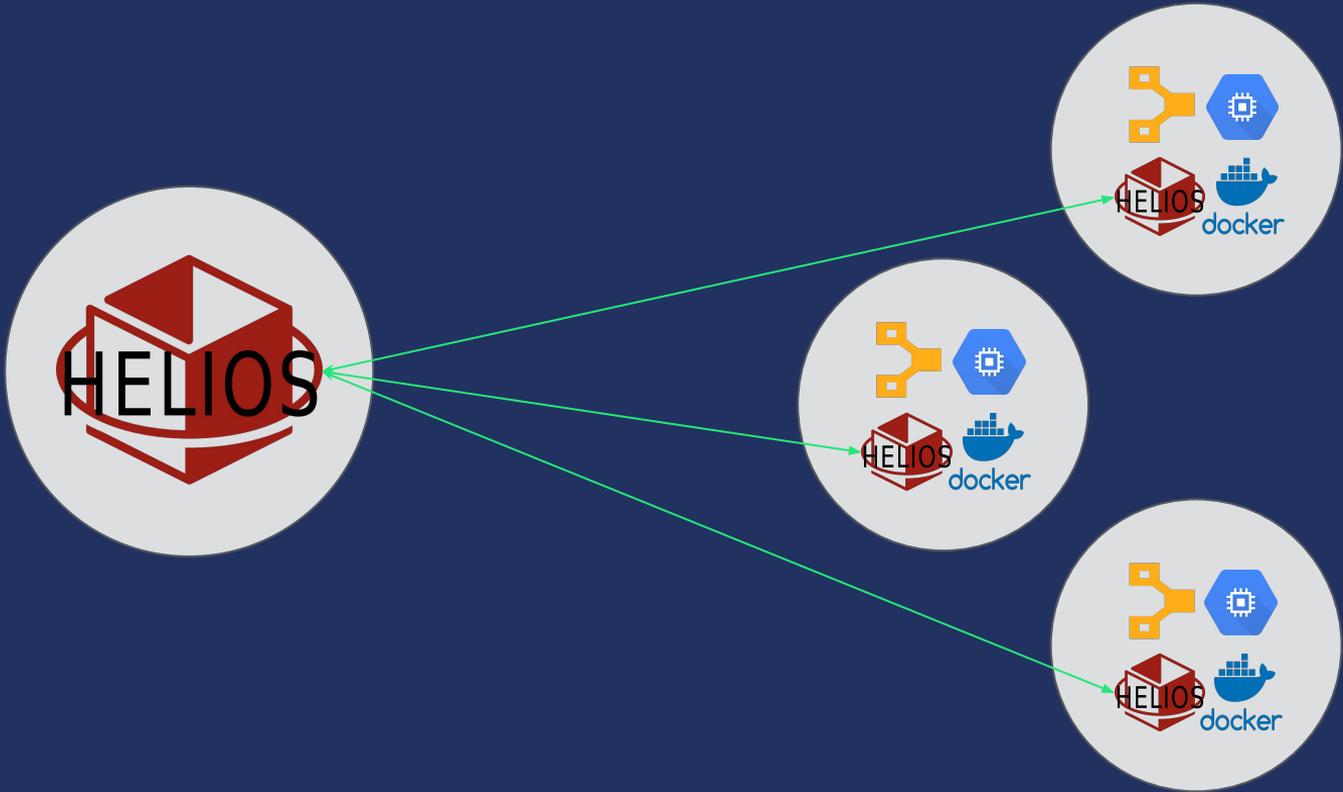
Core Infrastructure Org

Golden Path



Centralized Ops

# Old Runtime



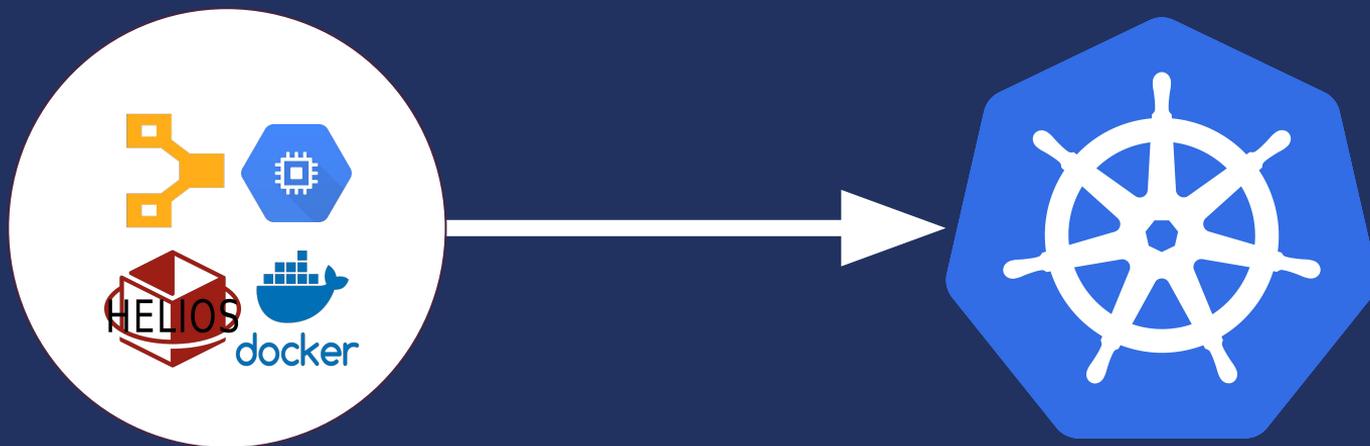
# New Runtime



# New Runtime\*



# Kubernetes: How We're Migrating



Goal/Scope: Run all stateless backend services at Spotify on Kubernetes.

Migrating Spotify's Runtime to Kubernetes: <https://www.youtube.com/watch?v=US4ttqOO-V4>

# GA Phase: Golden Path

**A** The Golden Path for Backend Engineering



# Operational Solidity

# Multicluster Strategy: Management

- Automated scripts → Declarative
- Terraform
- How Spotify Accidentally Deleted All its Kube Clusters with No User Impact - David Xia:

<https://kccnceu19.sched.com/event/MQbb/keynote-how-spotify-accidentally-deleted-all-its-kube-clusters-with-no-user-impact-david-xia-infrastructure-engineer-spotify>

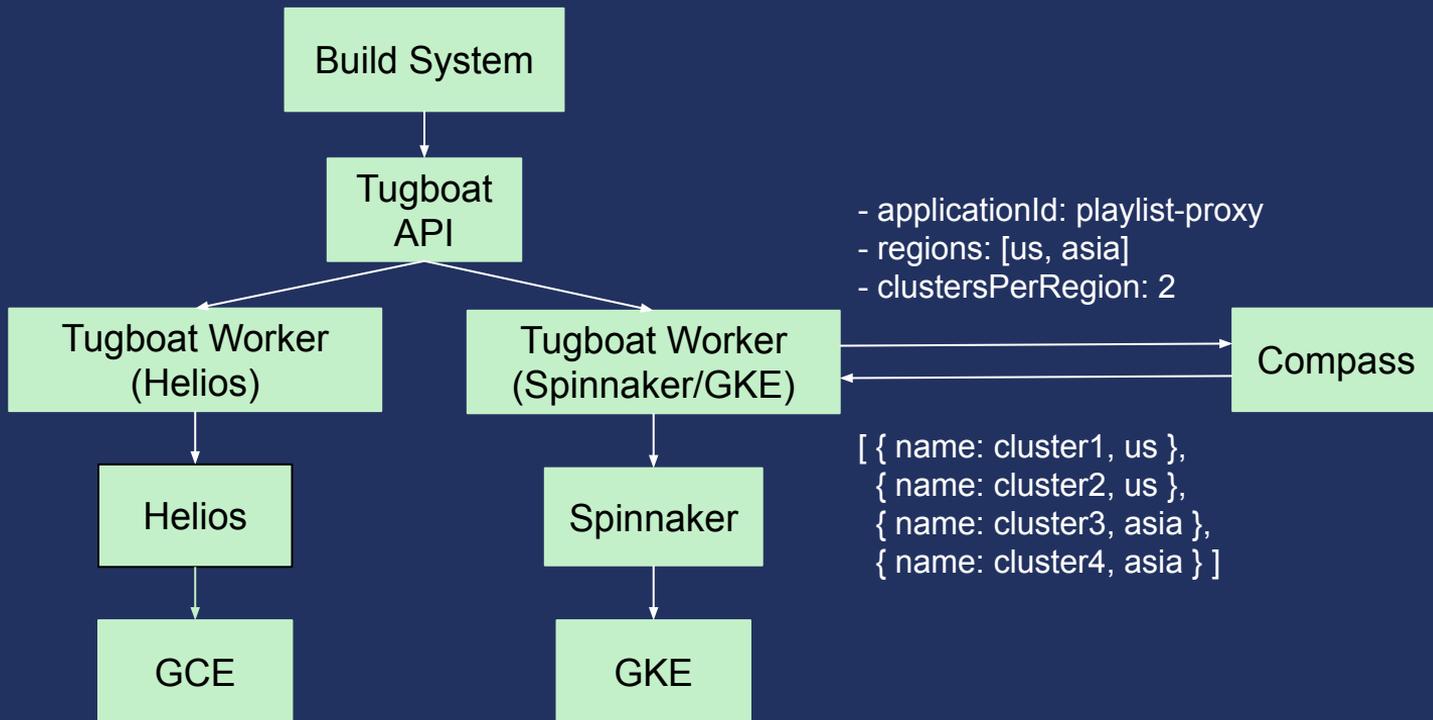


# Multicloud Strategy: Management

**Offload cluster  
management if possible.**

# Multicluster Strategy: Deployment

Tugboat (deployment layer):



# Multicloud Strategy: Compass

- Sticky assignment of apps —> clusters
- Current behavior:
  - Selection of clusters based on available capacity
  - Ability to deploy an app to N clusters per region

# Multicloud Strategy: Deployment

**Keep it simple for  
developers to reason about.**

# Multicluster: Other Talks

## To Infinite Scale and Beyond: Operating Kubernetes Past the Steady State

- Austin Lamon, Spotify & Jago Macleod, Google
- Wednesday (Nov 20) from 3:20pm - 3:55pm

<https://kccncna19.sched.com/event/UaYk/to-infinite-scale-and-beyond-operating-kubernetes-past-the-steady-state-austin-lamon-spotify-jago-macleod-google>

# Operational Metrics/Alerts

- Monitoring on integration components
  - error
  - latency
  - saturation
- Smoke test → periodic deploys
- Metrics for:
  - available cluster capacity
  - status of backups
  - scheduling

# Operational Metrics/Alerts

**Establish trust  
through monitoring.**

# DiRTing

- Disaster Recovery Testing
- Planned and executed DiRTs
  - Cluster deletions
  - Master being down
- Validation and improvement of:
  - Metrics + Alerting
  - Playbooks
  - Backups

# DiRTing

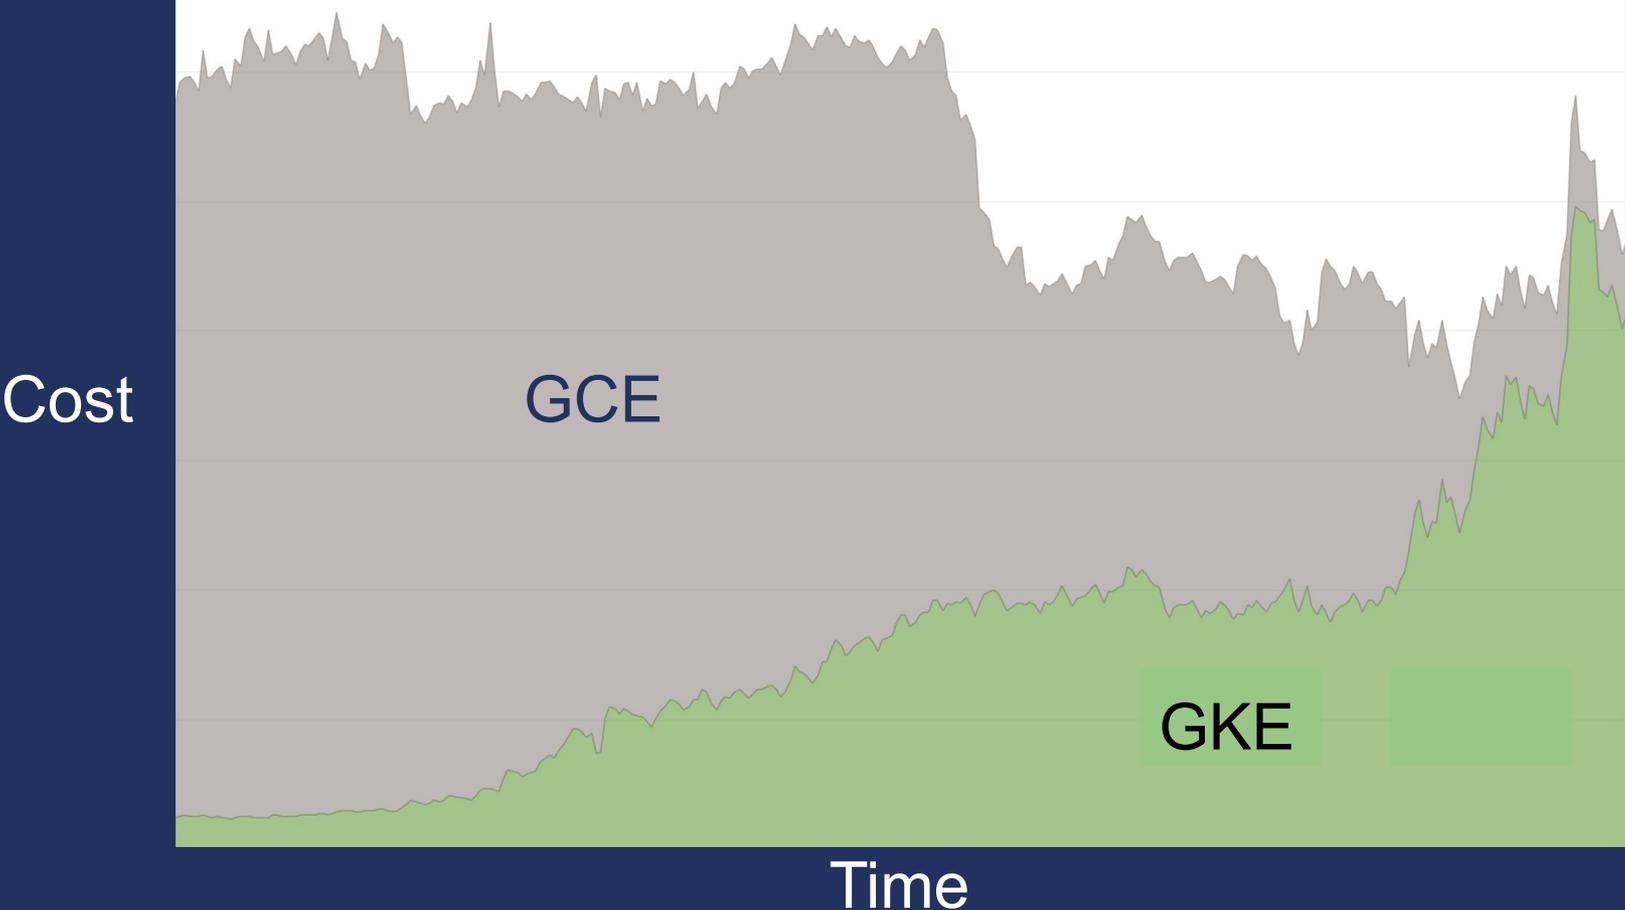
**Plan for and  
simulate failures.**

# Business Impact

# Metrics Tracking

- Metrics/Dashboard to track for us:
  - Resource Utilization
  - Cost
  - % apps migrated
  - Developer productivity

# Cost Tracking



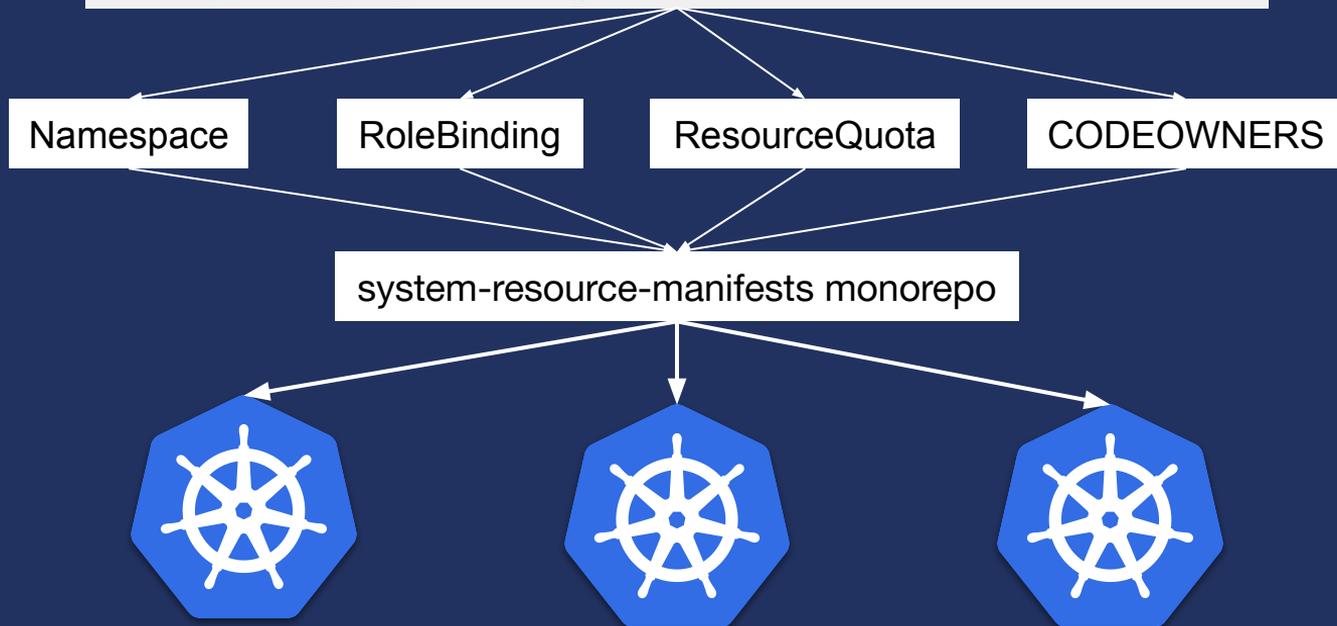
# Business Impact

**Track enough data to  
answer questions.**

# Self Service

# Self-Service Experience: Namespaces

```
# generate namespace manifests and supplementary resources
$ docker pull gcr.io/[REDACTED]/gke-onboard
$ docker run -v $(pwd):/generated gcr.io/[REDACTED]/gke-onboard \
  --namespace <namespace> \
  -l <logging_gcp_project> \
  --team <team-name> \
```



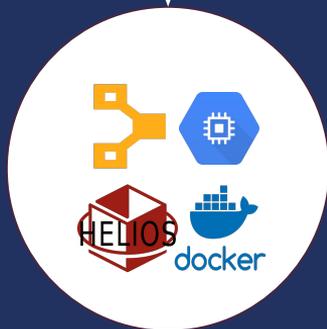
# Self-Service Experience: Namespaces

**Take complexity for  
your developers.**

# Self-Service Experience: Apps

Helios config file

```
version=$(tugboat version create)  
tugboat set-target-version --target-version=$version
```



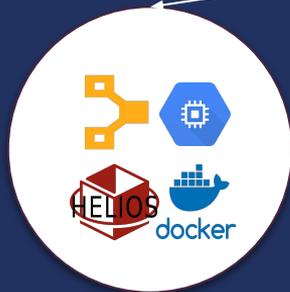
# Self-Service Experience: Apps

Helios config file

K8s resources

```
---  
version: 1  
components:  
- componentId: myservice  
  tugboat:  
    backends:  
    - backend: helios  
    - backend: kubernetes
```

```
version=$(tugboat version create)  
tugboat set-target-version --target-version=$version
```

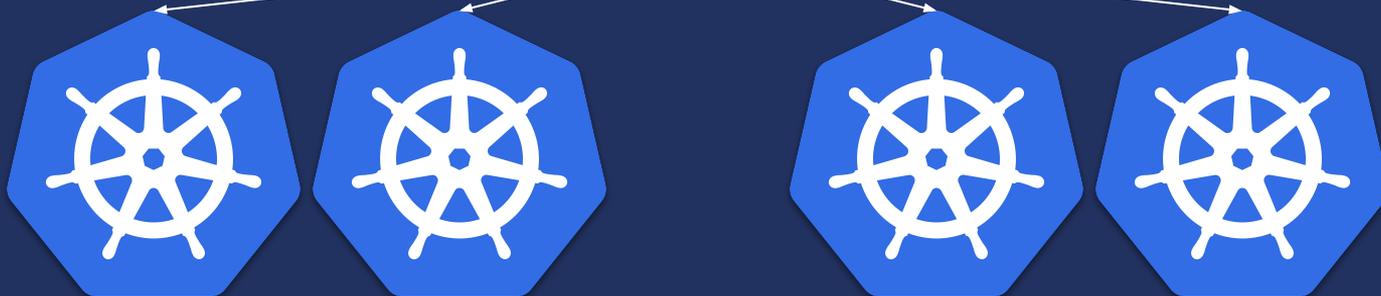


# Self-Service Experience: Apps

K8s resources

```
---  
version: 1  
components:  
- componentId: myservice  
  tugboat:  
    backends:  
      - backend: kubernetes  
        clustersPerRegion: 2  
        regions: [us, europe]
```

```
version=$(tugboat version create)  
tugboat set-target-version --target-version=$version
```



# Self-Service Experience: Apps

**Use existing  
tools/abstractions.**

# Documentation

- Internal Documentation
- Stackoverflow Enterprise (vast majority of questions on there)
- Backend Golden Path

## GKE @ Spotify

Home

Introduction

Getting Started

Roadmap and Use Cases

FAQ

Definitions

Known issues

Feedback

Migration Best Practices

Developers ^

Accessing GCP

Connection Draining

CronJobs

Dealing with Disruptions

Deployment

Troubleshooting Deployments

Event Delivery

ITGC in-scope systems

Java on Kubernetes

Logging

Metrics

Operations

Reference

Scaling

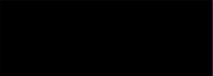
Secrets

# Documentation

**You as owners are  
not the audience.**

# Shared Support

- Developers will always have questions
- US East Coast + Europe
- Support shifts covering both
- Ex. Topic for support Slack channel:

For [backend services on GKE](http://docs/gke): <http://docs/gke>. Save time and check  <http://go/so/questions/tagged/gke> | NYC-goalie1: @   
| NYC-goalie2: @  | STO-goalie: @ 

# Shared Support

**Support can be  
mutually beneficial.**

# Feature Parity

# Deployment (Canaries)

- Existing canary feature in Tugboat for Helios
- Similar feature on K8s (basically suffix where appropriate with -canary)

NAME	DESIRED	CURRENT	UP-TO-DATE	AVAILABLE	AGE
compass	2	2	2	2	276d
hoth	2	2	2	2	175d
tugboatapi	2	2	2	2	94d
tugboatapi-canary	1	1	1	1	80d
tugboatheliosmonitor	1	1	1	1	272d
tugboatheliosmonitor-canary	1	1	1	1	80d
tugboatheliosworker	3	3	3	3	272d
tugboatheliosworker-canary	1	1	1	1	80d
tugboatnotifications	1	1	1	1	272d
tugboatspinnakerworker	2	2	2	2	354d
tugboatspinnakerworker-canary	1	1	1	1	80d

# Feature Parity

**Developers need feature parity or better.**

# Our GA Requirements

- Operational Solidity
- Business Impact
- Self Service Experience
- Important Feature Parity

# Future Plans

- Stateful services?
- Data jobs Consolidation
- CronJobs as first class citizens
- Higher level abstractions

# Migration & GA Learnings

- Be mindful and deliberate in dictating terminology.
- Migrate in stages/series of goals that increase in scope.
- Devs will adopt and/or migrate to your Kubernetes offering if it is easy, complete, and compelling.
- Talk to other companies about how they do infrastructure/solve infrastructure problems.



Thanks for listening!

We're hiring!

[spotifyjobs.com](https://spotifyjobs.com)

James Wen - [jameswen@spotify.com](mailto:jameswen@spotify.com)

[@rochesterinnyc](https://twitter.com/rochesterinnyc)

KubeCon + CloudNativeCon North America 2019

November 19th, 2019