



AWS
re:Invent

CON 329

Running Kubernetes clusters at scale: Square

Geoff Flarity

Engineering Manager
Square

Jay Estrella

Software Engineer
Square

Omar Lari

Business Development
Manager
Amazon Web Services

Agenda

Amazon Elastic Kubernetes Service (Amazon EKS) and how it helps customers

About that Cash App

Cash App cloud platform

Platform impact on velocity

Open source


AWS best practices for building modern applications

- Create a culture of innovation by **organizing into small DevOps teams**
- Continually evaluate your security posture by **automating security**
- Componentize applications using **microservices**
- Update applications & infrastructure quickly by **automating CI/CD**
- Standardize and automate operations by **modeling infrastructure as code**
- Simplify infrastructure management with **serverless technologies**
- Improve application performance by **increasing observability**

What is Kubernetes?



Open source container
management platform



Helps you run
containers at scale



Gives you primitives
for building
modern applications

How are customer using Amazon EKS?




Microservices



Platform as a service



**Enterprise App
Migration**



Machine Learning



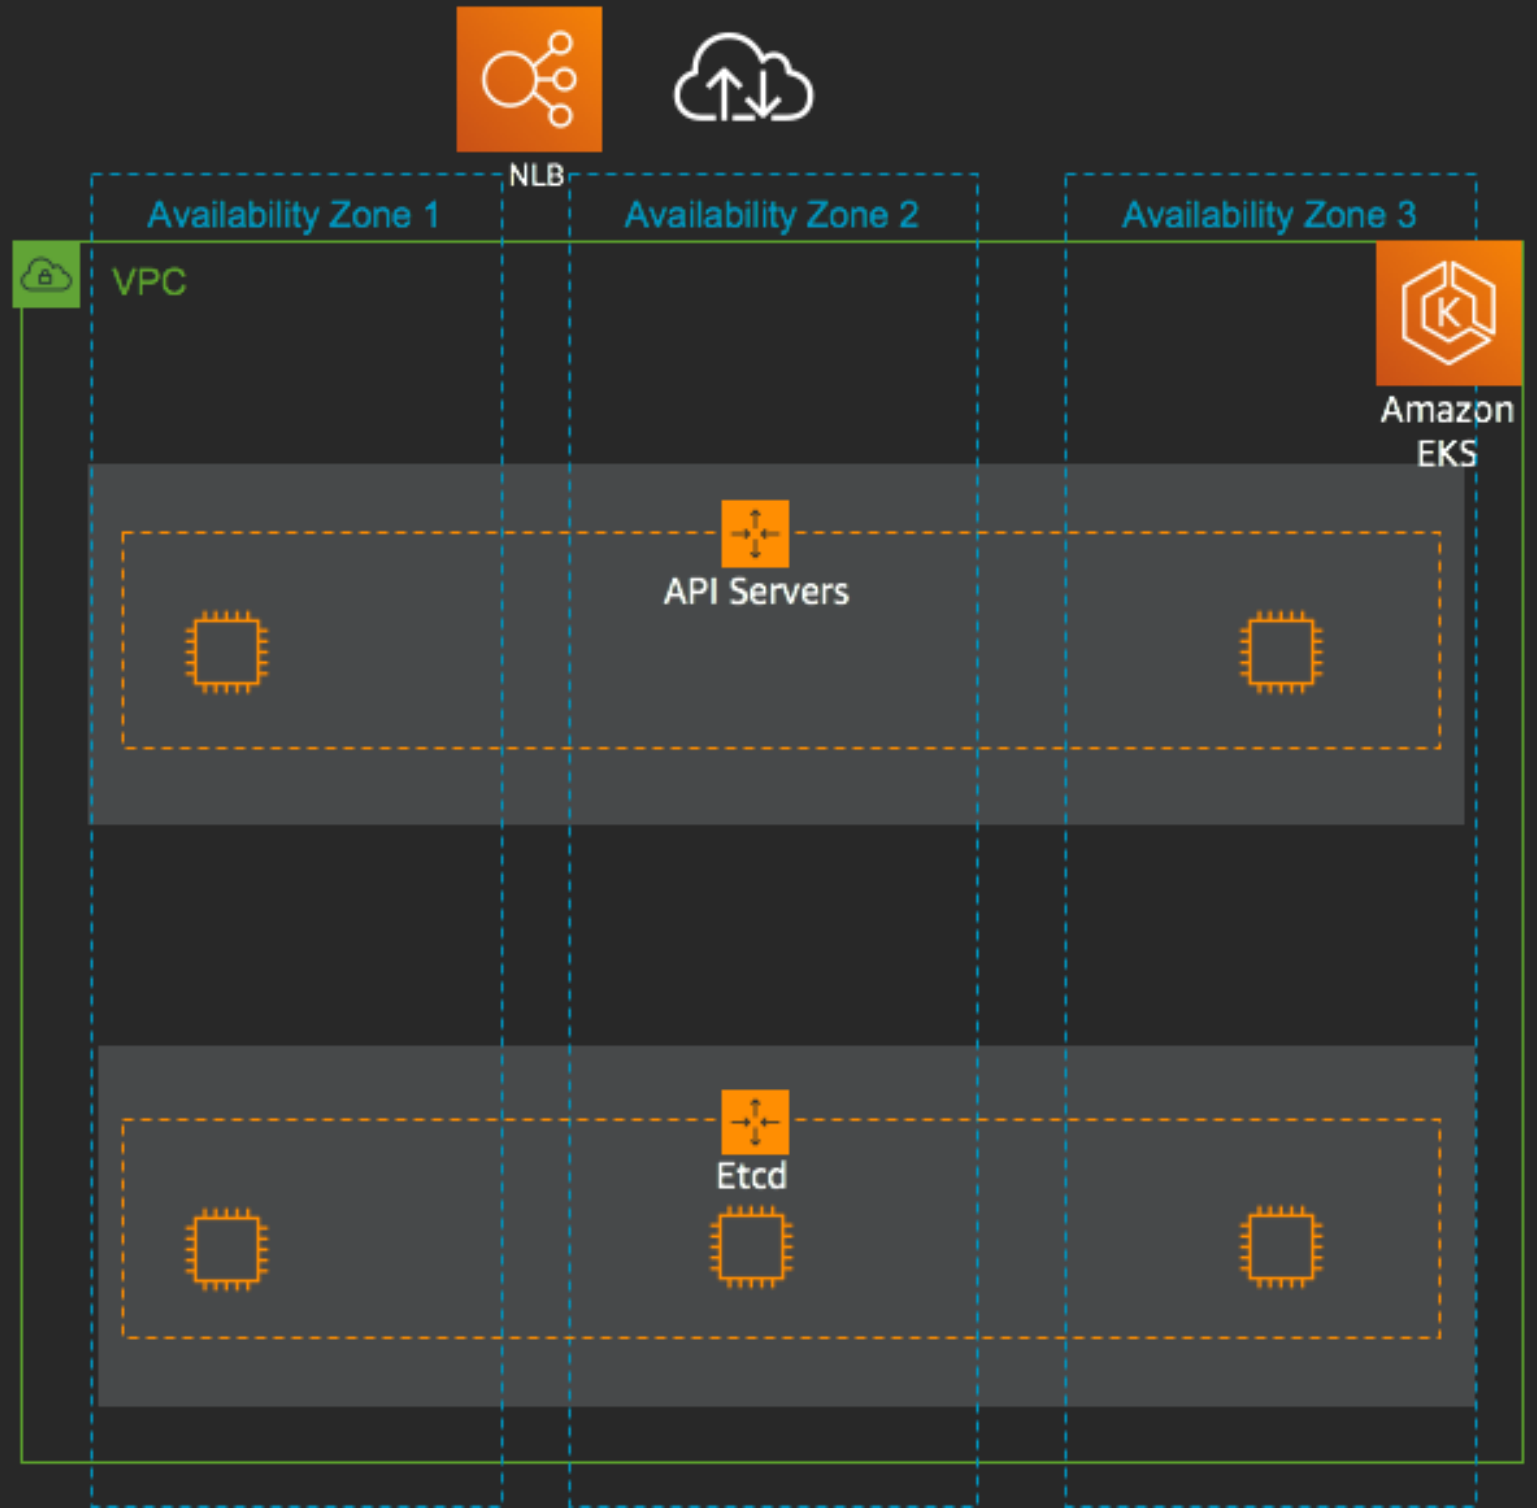
Amazon Elastic Kubernetes Service

Kubernetes control plane

Highly available and single
tenant infrastructure

All “native AWS” components

Fronted by a Network Load
Balancer



Our tenets

1. **Amazon EKS is a platform to run production-grade workloads.** Security and reliability are our first priority. After that we focus on doing the heavy lifting for you in the control plane, including life cycle-related things like version upgrades.
2. **Amazon EKS provides a native and upstream Kubernetes experience.** Amazon EKS provides vanilla, un-forked Kubernetes. In keeping with our first tenant, we ensure the Kubernetes versions we run have security-related patches, even for older, supported versions as quickly as possible. But there's no special sauce and no lock in.
3. **If you want to use additional AWS services, integrations are as seamless as possible.**
4. **The Amazon EKS team in AWS actively contributes to the upstream Kubernetes project** and the wider CNCF activities, both on the technical level as well as community, from communicating good practices to participation in SIGs and working groups.

Amazon EKS, a year in review

June – December 2018

Amazon EKS achieves K8s conformance, HIPAA-eligibility, Generally available

Amazon EKS AMI build scripts and AWS CloudFormation templates available in GitHub.

Support for GPU-enabled EC2 instances, support for HPA with custom metrics.

Amazon EKS launches in Dublin, Ireland

Amazon EKS simplifies cluster setup with update-kubeconfig CLI command

Amazon EKS adds support for Dynamic Admission Controllers (Istio), ALB Support with the AWS ALB ingress controller

Amazon EKS launches in Ohio, Frankfurt, Singapore, Sydney, and Tokyo

Amazon EKS adds Managed Cluster Updates and Support for Kubernetes Version 1.11, CSI Driver for Amazon Elastic Block Store (Amazon EBS)

2019

Amazon EKS launches in Seoul, Mumbai, London, and Paris

Amazon EKS achieves ISO and PCI compliance, announces 99.9% SLA, cluster creation limit raised to 50

API Server Endpoint Access Control, AWS App Mesh controller

Windows support (preview), Kubernetes version 1.12,

CSI Drivers for Amazon Elastic File System (Amazon EFS), Amazon FSx for Lustre, Control Plane Logs, A1 (ARM) instance support (preview)

Deep Learning Benchmark Utility, Public IP Address Support,

Simplified cluster authentication, SOC compliance, Kubernetes 1.13, PodSecurityPolicies,

Container Insights, CNI 1.5.0, Amazon ECR, AWS PrivateLink Support

Open-source roadmap

<https://github.com/aws/containers-roadmap/>

The screenshot displays the GitHub repository for the AWS Containers Roadmap. The repository is titled "aws / containers-roadmap" and has 365 watches, 1,370 stars, and 37 forks. The main navigation bar includes links to Code, Issues (213), Pull requests (0), Projects (1), Security, and Insights. The repository is updated an hour ago.

The issues are categorized into five columns, each representing a different stage of development:

- Researching (8 results):** Issues include "[EKS]: Managed Cluster Addons", "[EKS] [Requesting Feedback] Support for Deploying to Kubernetes from CloudFormation", "[EKS] Install AWS-Service-Operator on master nodes", "[EKS] Install AWS EBS CSI Driver as Part of EKS cluster creation", "[EKS] [request]: allow to configure ipv6 kube-proxy mode", "[EKS] [request]: Security Groups per Pod", "[EKS] [request]: Ability to configure pod-eviction-timeout", and "[EKS] Enable HPA with CloudWatch metrics and alarms".
- We're Working On It (18 results):** Issues include "#166 opened by jaxxstorm", "[EKS] [Security]: Allow restricting EKS API Access via Security Groups", "[EKS] CloudFormation support for control plane logging and endpoint access control", "EKS-Optimized AMI Metadata SSM Parameter", "[EKS] [request]: EKS Support for Kubernetes 1.14", "New EKS Region : Beijing", "DNS resolution for EKS Private Endpoints", "Support for PodSecurityPolicy Admission Controller", and "Fargate for EKS".
- Coming Soon (6 results):** Issues include "[EKS] : Kubernetes v1.10 Deprecation", "[EKS]: Release CNI v1.5.0", "[EKS]: Service Linked Role for Amazon EKS", "EKS Support for Kubernetes 1.13", "New EKS region: São Paulo", and "New EKS Region: Canada Central".
- Developer Preview (2 results):** Issues include "EKS Windows Nodes (preview)" and "[EKS]: Support for Arm Nodes - EC2 A1 Instances".
- Just Shipped (31 results):** Issues include "SOC compliance for EKS", "EKS: Get-Token CLI Subcommand", "Support for Public IP space in VPC with EKS", "[EKS] [request]: Release CNI Plugin 1.4 for EKS", "EKS / Kubernetes: Add support for using AWS Fleet to atlasian/escalator", "Control Plane Metrics Endpoint", "Amazon EKS: Deep Learning Benchmarking Utility", "EKS: Documentation for using Kubeflow on AWS", and "EKS Control Plane Logs".

Each issue card shows the issue number, title, and the user who opened it. Some issues are marked with "EKS" or "Proposed" labels.

Amazon EKS services roadmap: Highlights

Shipped

- **Amazon EKS control plane logs**
- **IAM Roles for Service Accounts**
- **Region Expansion – Seoul, Mumbai, London, Paris, Hong Kong, Bahrain**
- **Managed Node Groups**
- **EKS Support for K8s version 1.13 + ECR AWS PrivateLink**

Coming soon

- **Service linked role for Amazon EKS**
- **KMS Encryption Provider**
- **Fargate for EKS**
- **New Amazon EKS Regions: Beijing, Ningxia**

Working on it

- **Managed add-ons**
- **DNS resolution of Amazon EKS private endpoints**
- **Improved console experience**
- **Next-generation CNI plugin**
- **Regional Expansion**

Cash App




Cash App



Cash App



Cash App




Cash App




Send Cash



Free Visa Debit
A free Visa debit card with custom signature




Cash Boost
Instant discounts at your favorite merchants



Cash from ATMs
Use your Cash Card to make ATM withdrawals



Cash App




Send Cash



Free Visa Debit

A free Visa debit card with custom signature



Cash Boost

Instant discounts at your favorite merchants



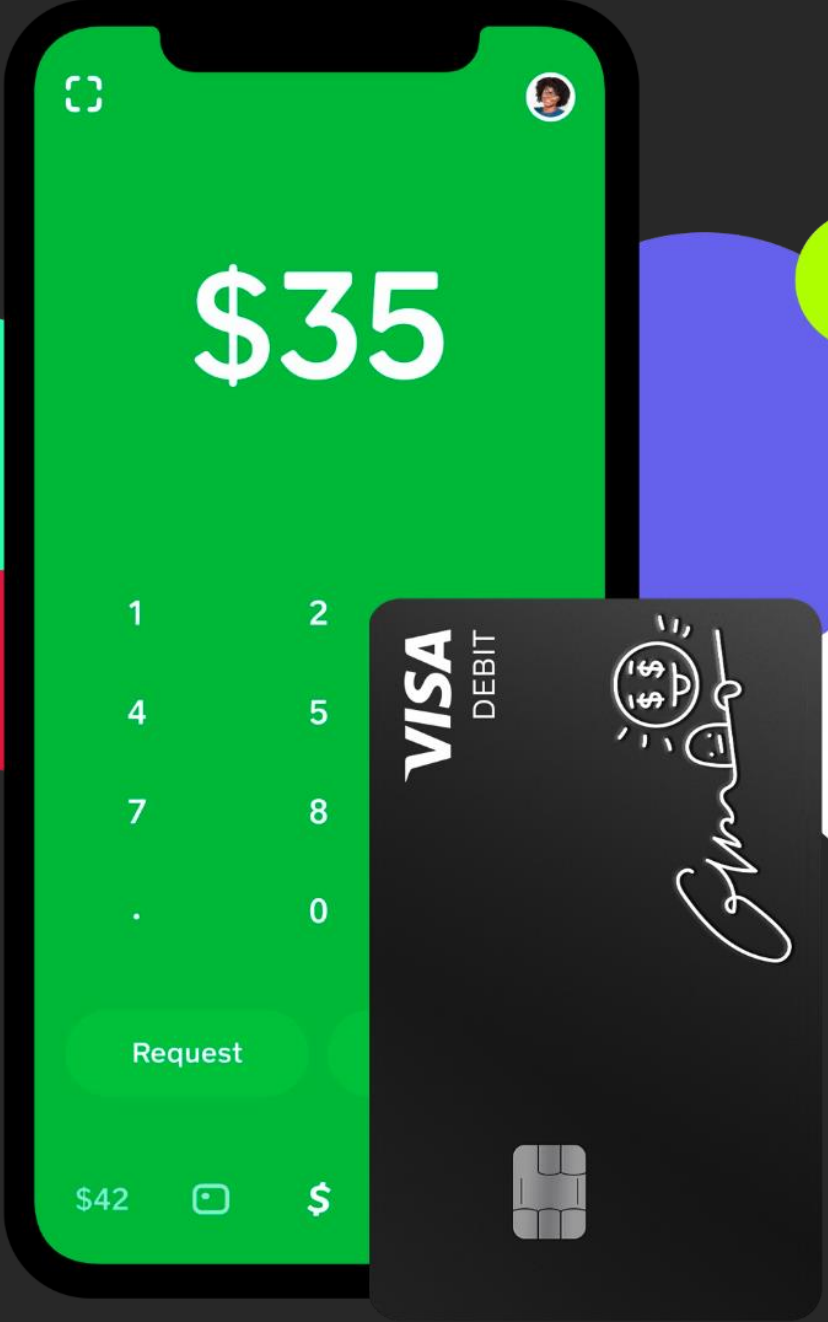
Cash from ATMs

Use your Cash Card to make ATM withdrawals




Direct Deposit

Deposit your paycheck directly into Cash App



Cash App




Send Cash



Free Visa Debit

A free Visa debit card
with custom signature



Cash Boost

Instant discounts at your
favorite merchants




Cash from ATMs

Use your Cash Card to
make ATM withdrawals




Direct Deposit

Deposit your paycheck
directly into Cash App




Transact Bitcoins

Buy/Sell Bitcoins from Cash
App



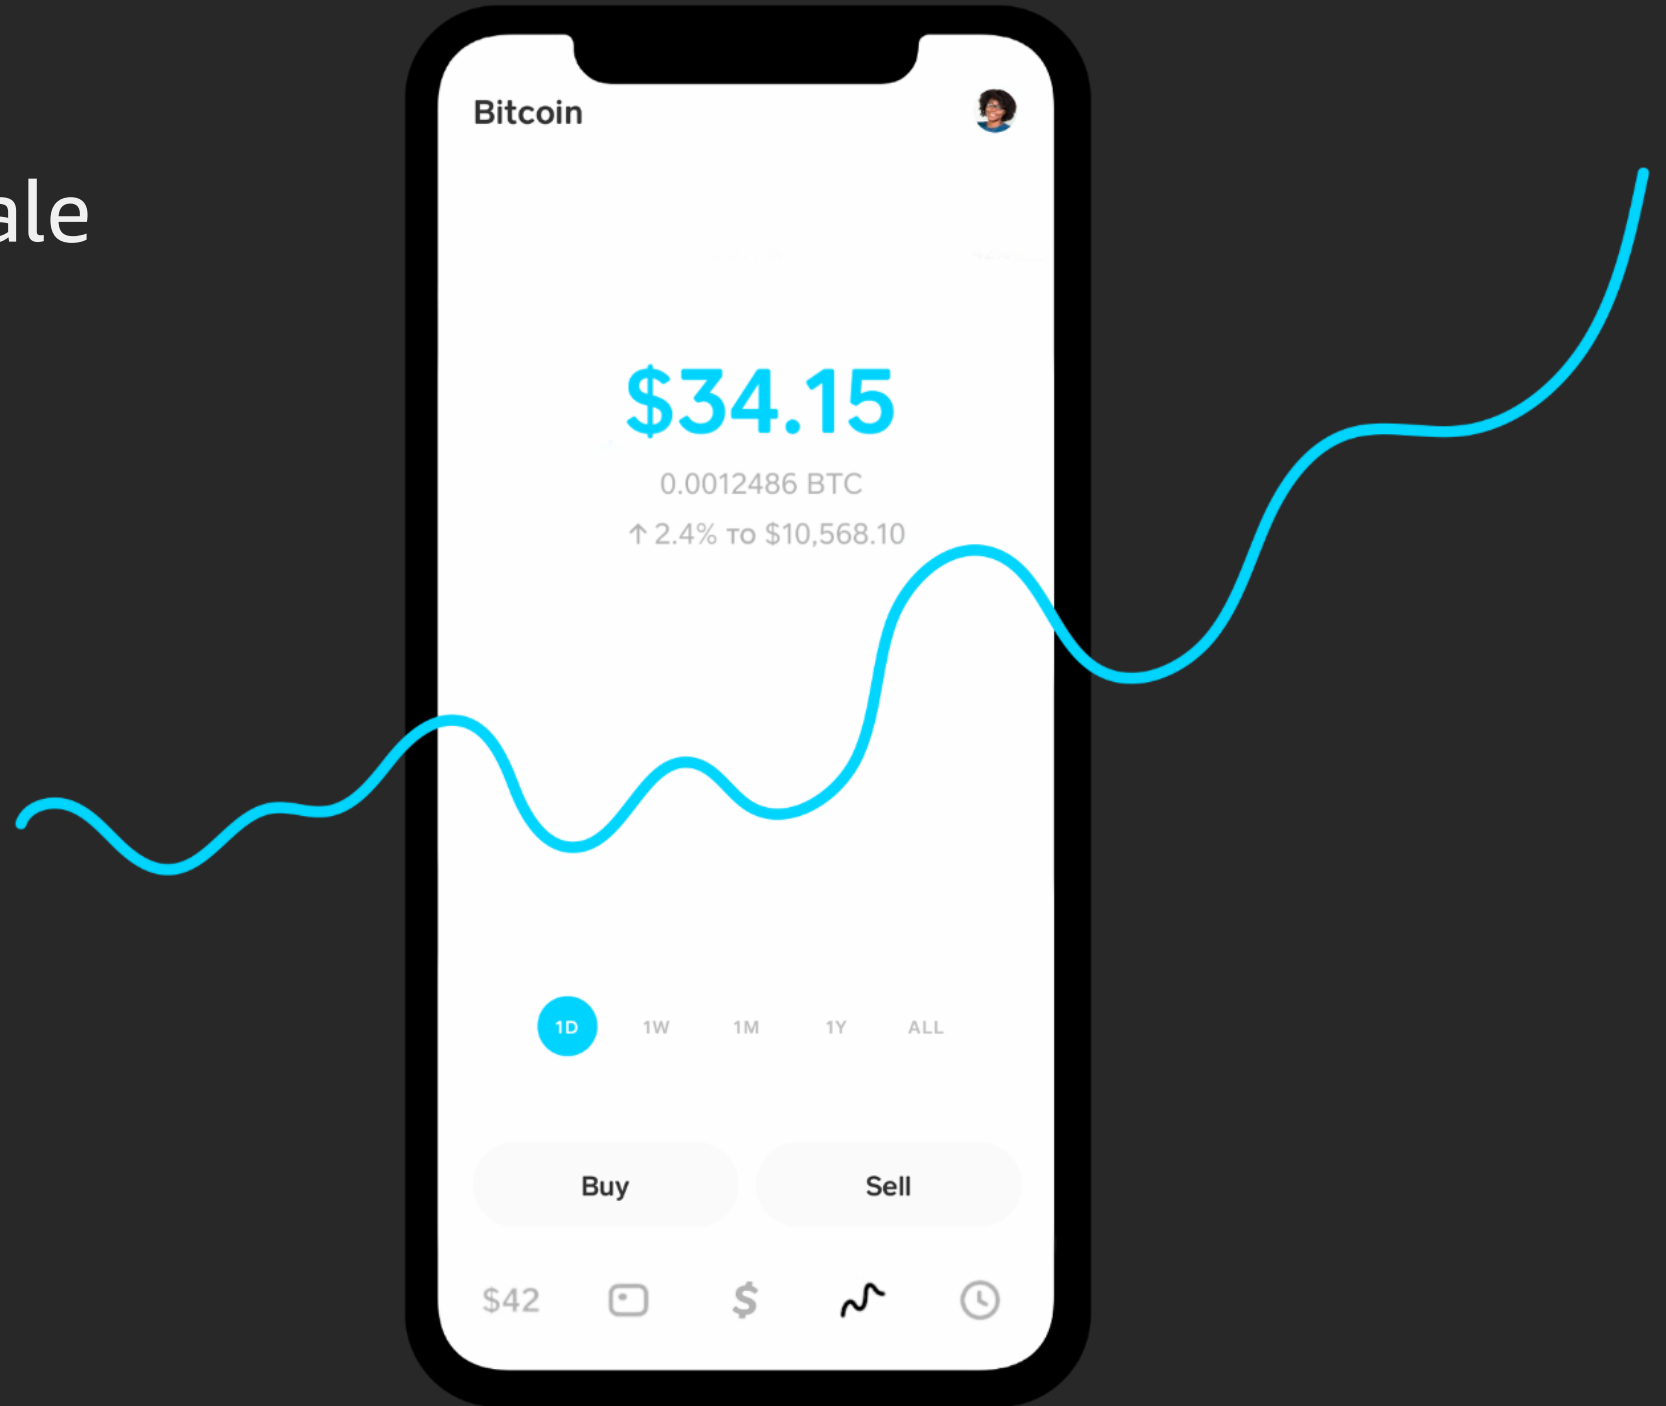
Cash App

- Each month, over 15 million users make or receive payments through the Cash App
- We're the No. 1 financial app in the Apple AppStore
- This year we saw 2.5x year-over-year growth in Q1 and booked over \$260 million in revenue for Q2



Cash App

- Popularity necessitates scale



Cash App platform

Cash App platform: Core values

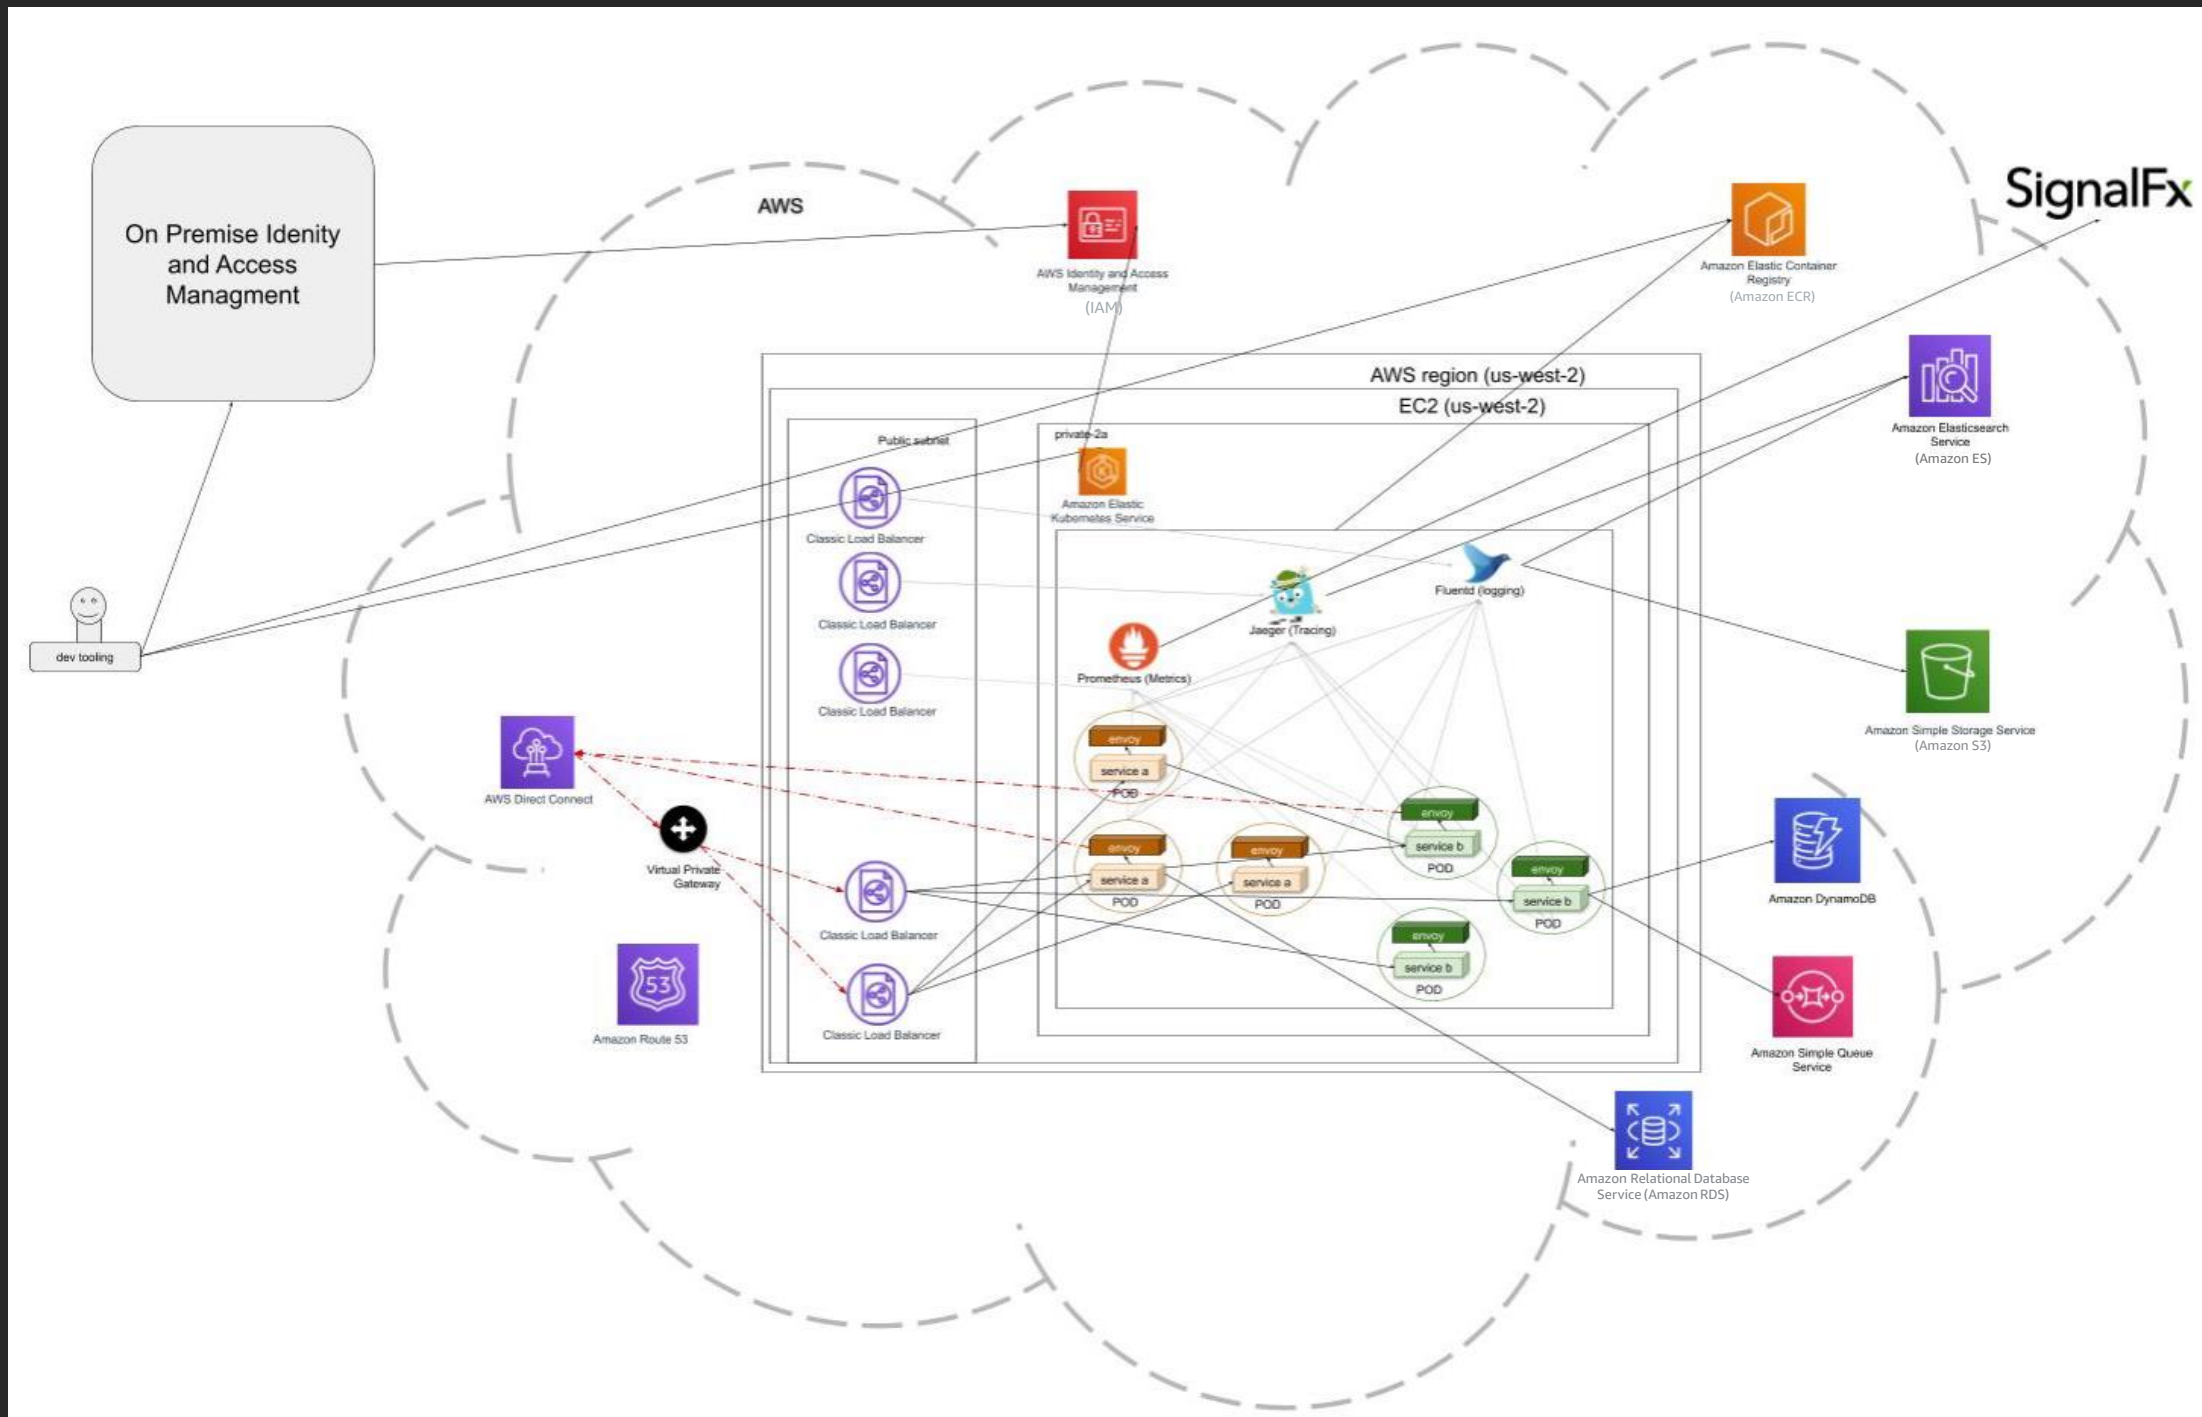
- Maximize engineering velocity
 - Hiring is hard, make your current engineers more productive
- Security
 - Security is every casheir's responsibility, and that's reflected in the platform
- Leverage-managed free/libre open source
 - Spend more time building product, find help/how-to's easily, limit lock-in to the operational
- Infrastructure as code
 - Code reviewed and self service

Cash App platform: Feature highlight

Key Features:

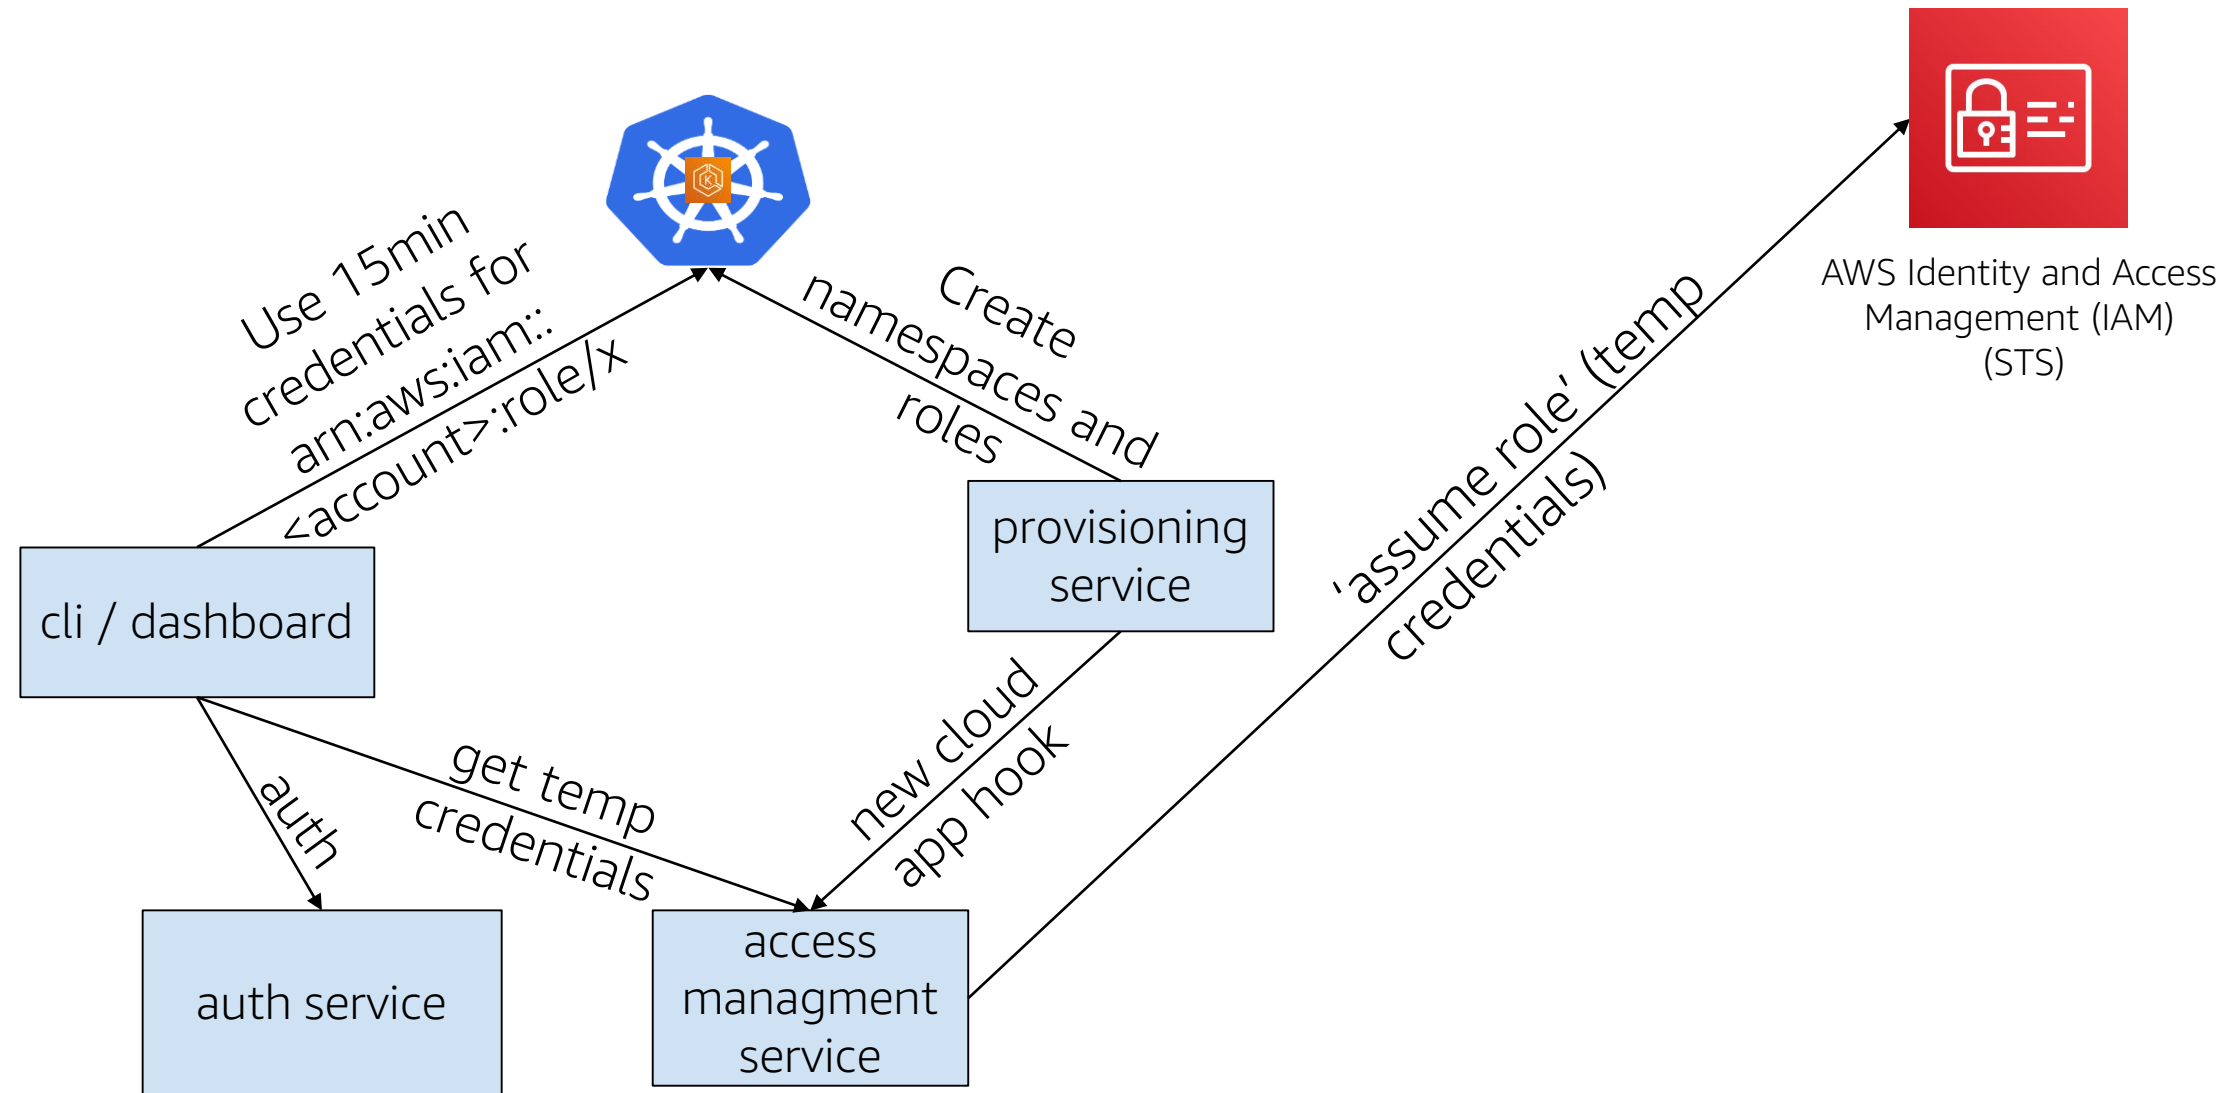
1. Security
2. Autoscaling
3. Metrics & alerting
4. Logging & tracing
5. Self service everything

Cash cloud platform: High-level overview

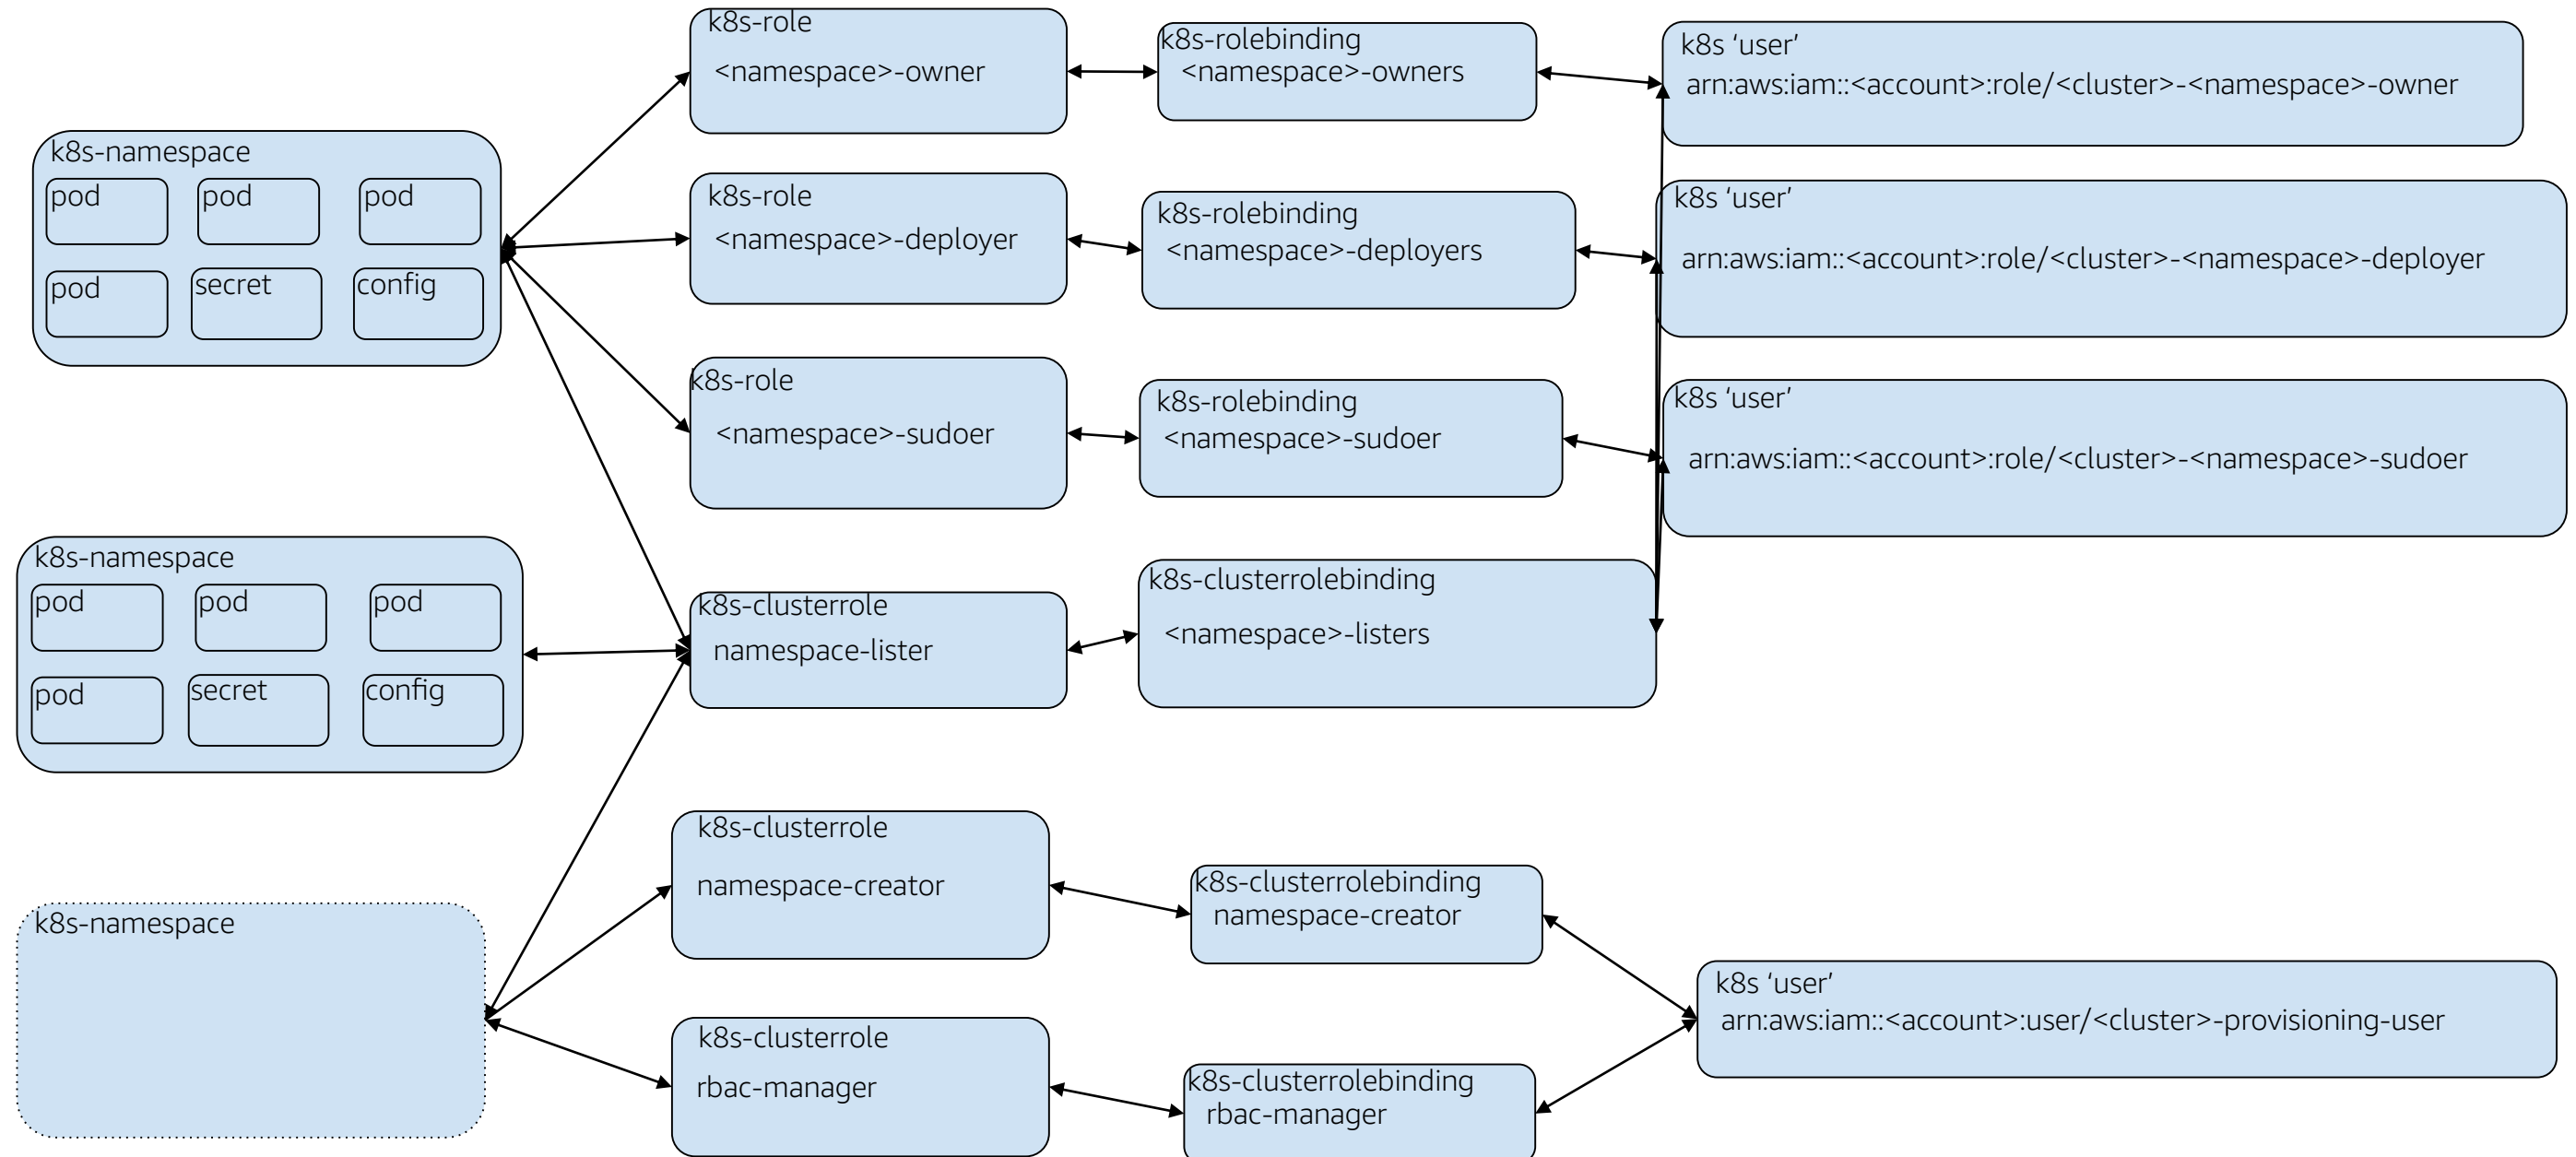


Security

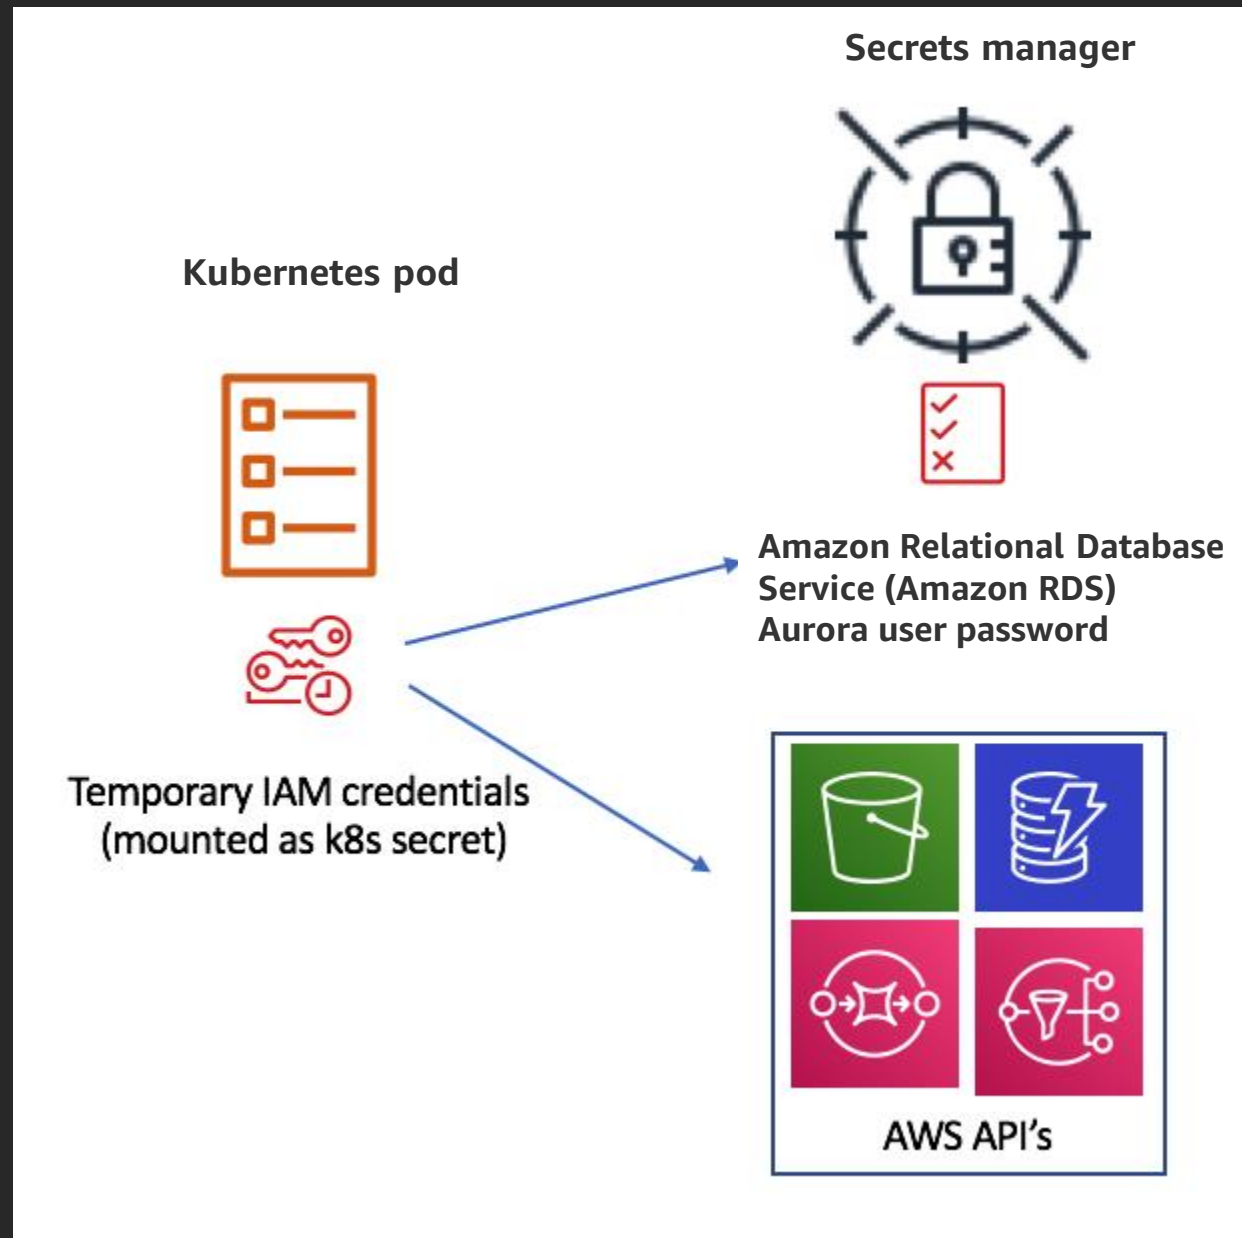
Cash App platform: Security



Cash App platform: Security



Cash platform: Security



Cash platform: Security

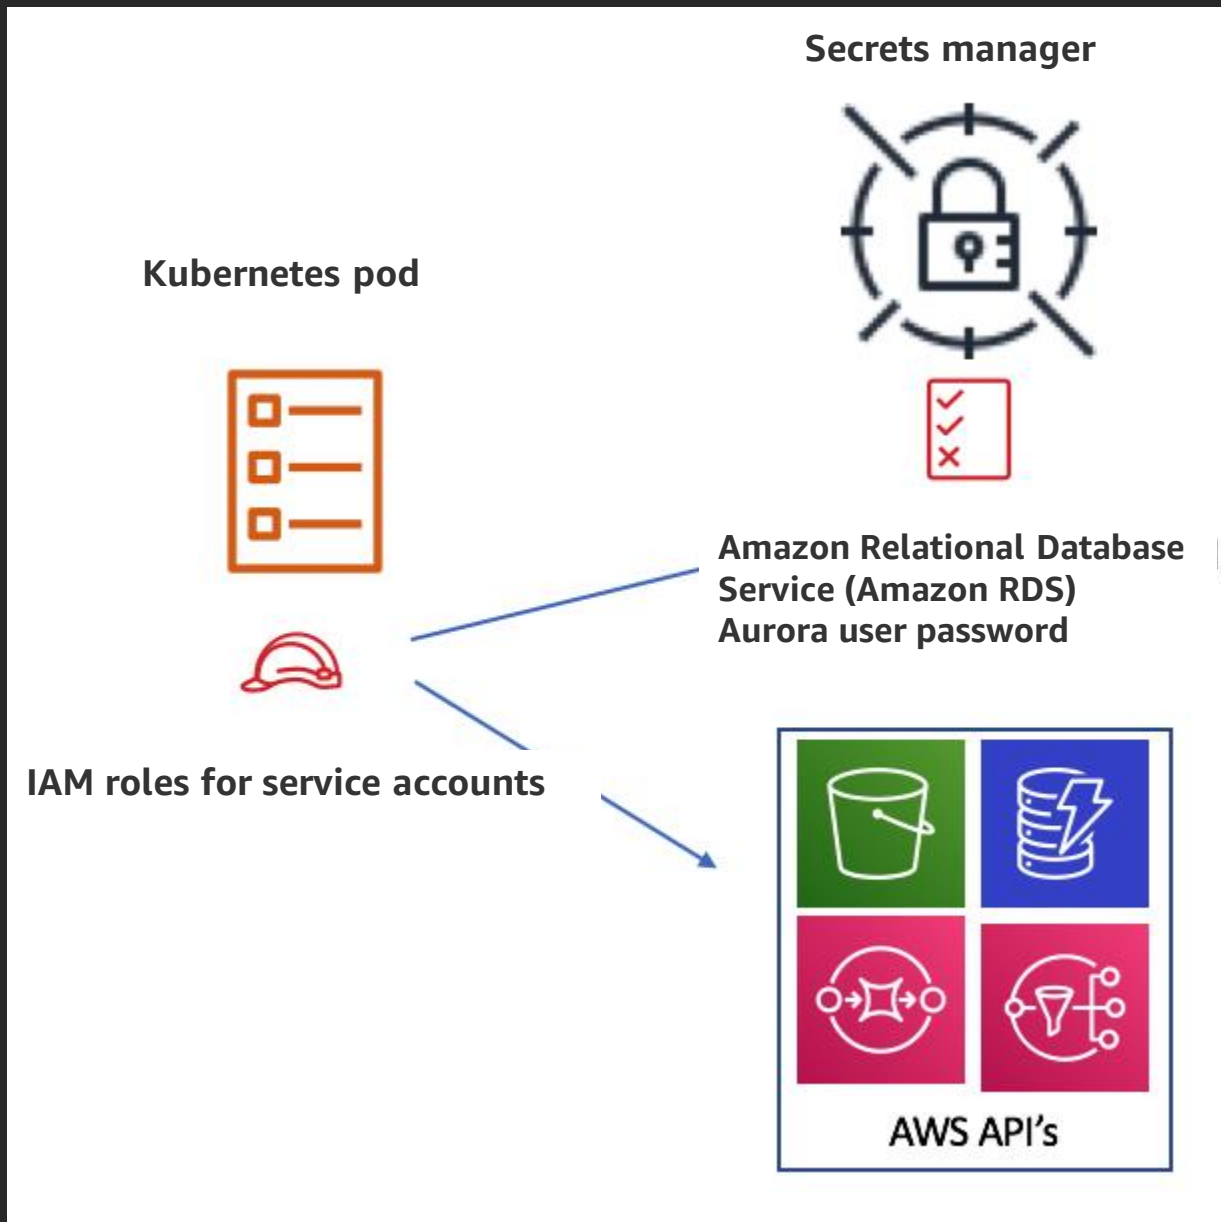
For each service:

- List access keys for the service

For each region the service is in:

- Get the currently active access keys from secret manager
- Use AWS SDK to create new credentials
- Store credentials in secrets manager
- Use the Kubernetes SDK to update cred secrets and trigger rolling restart
- Delete any inactive access keys in secrets manager

Cash platform: Security



Autoscaling

Cash platform environment: Kubernetes autoscaling

Horizontal pod autoscaler

+

Placeholder pods

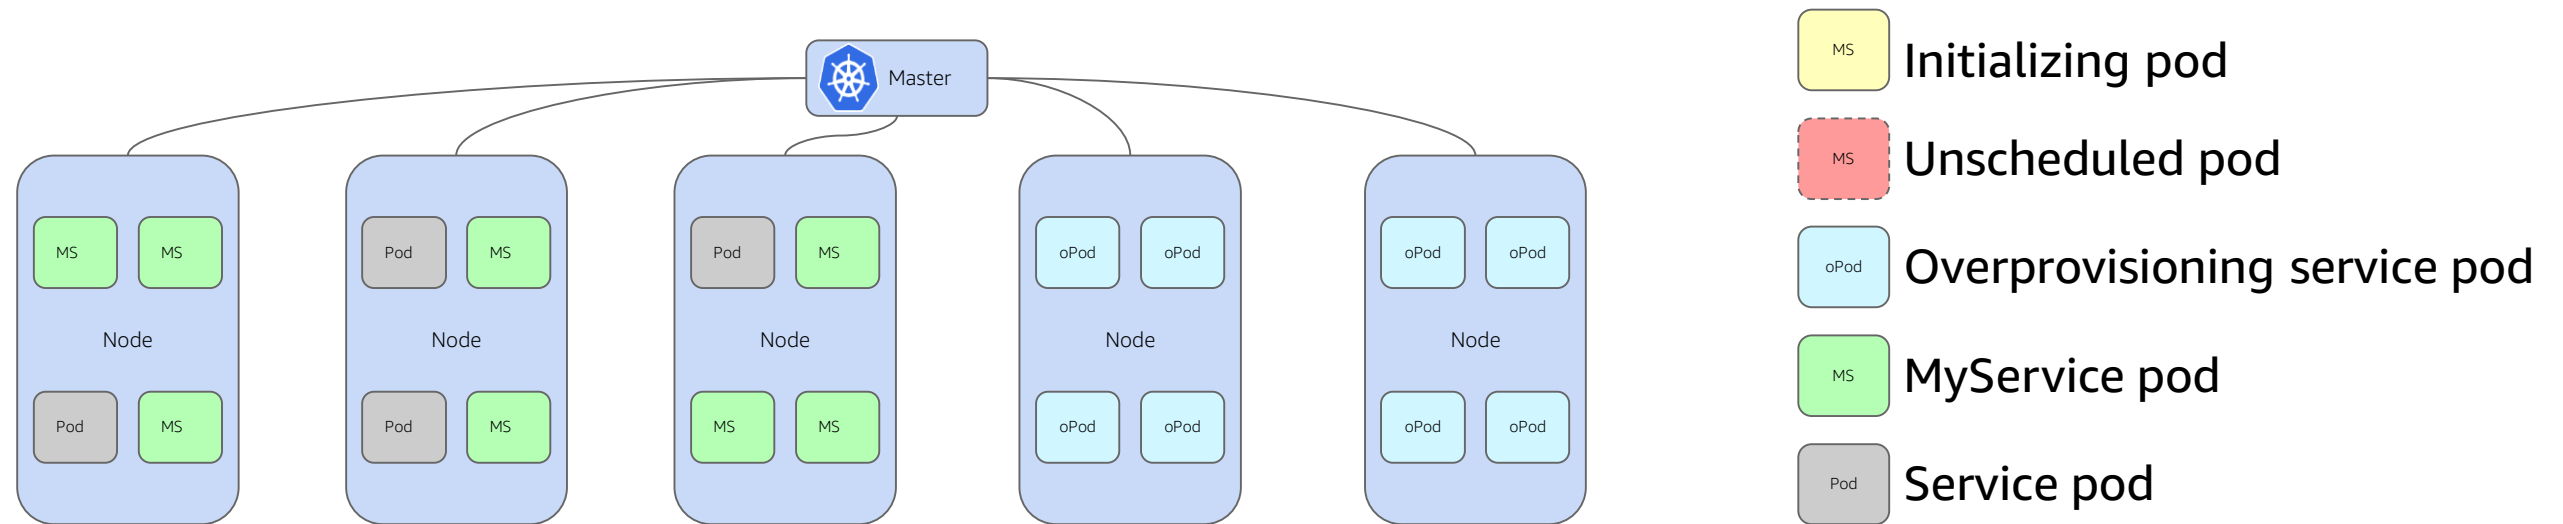
+

Cluster autoscaler

=

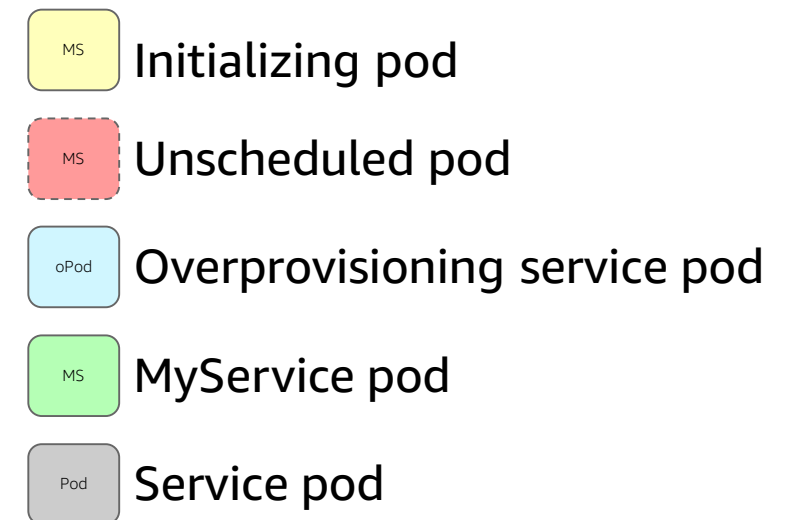
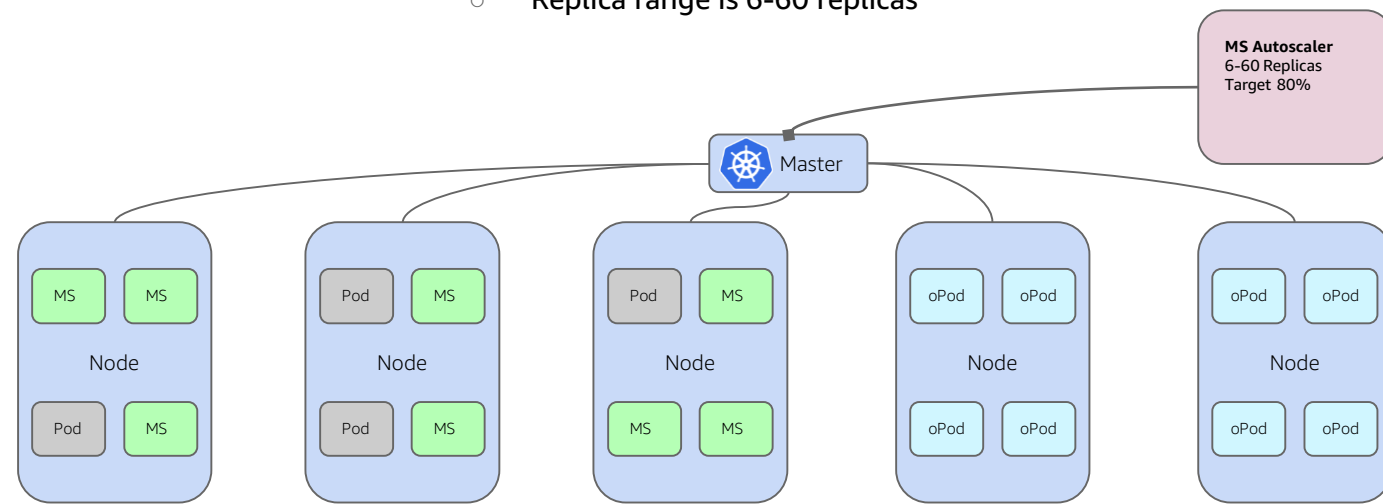
<3

Cash platform environment: Kubernetes autoscaling

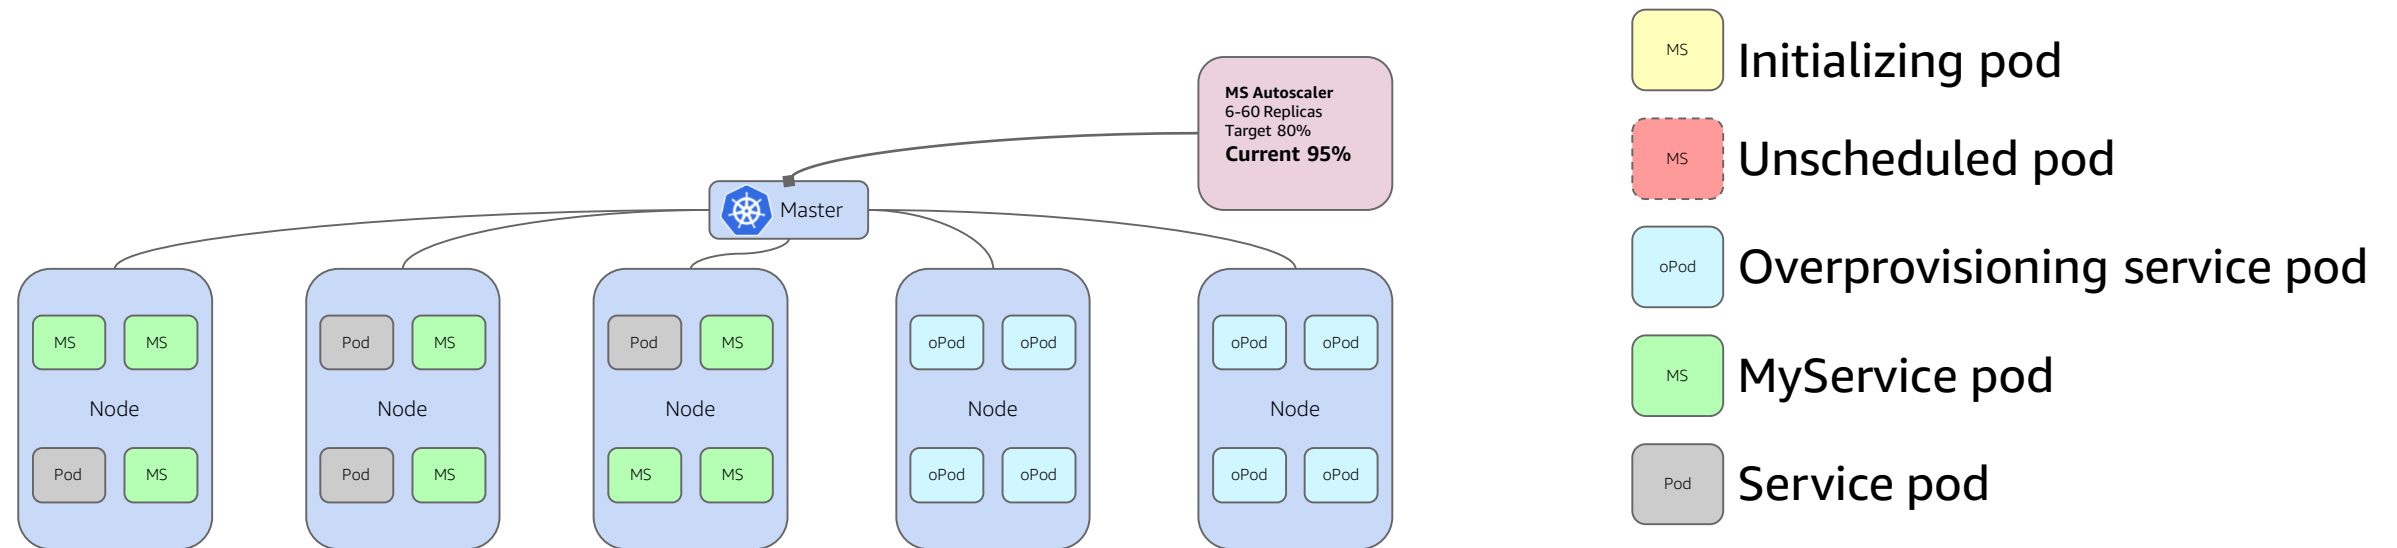


Cash platform environment: Kubernetes autoscaling

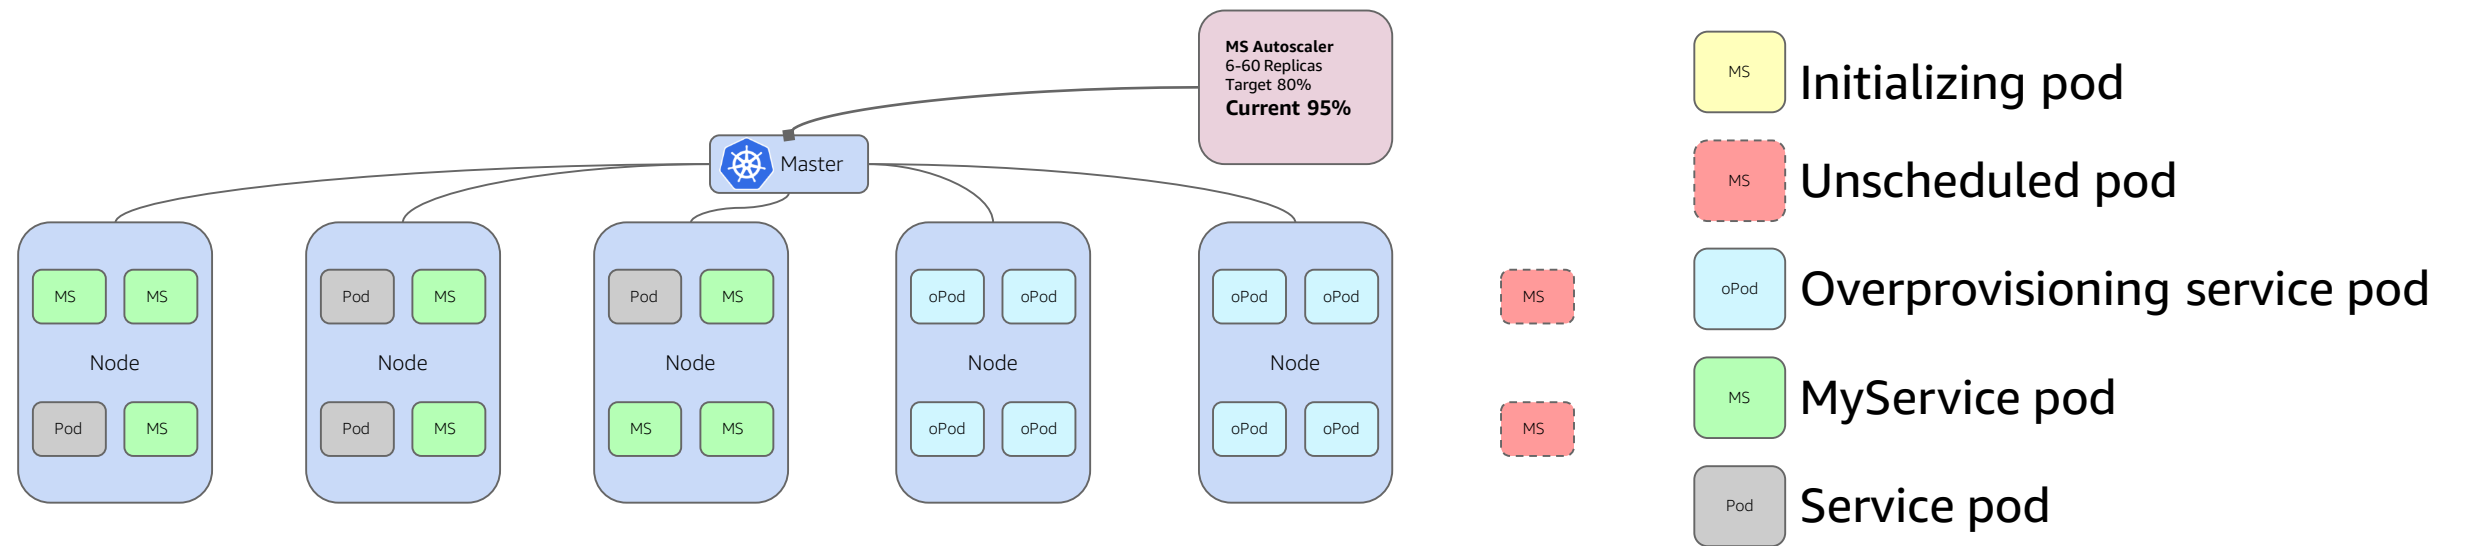
- MyService (MS) has 8 replicas and a Horizontal Pod Autoscaler
 - CPU threshold is 80%
 - Replica range is 6-60 replicas



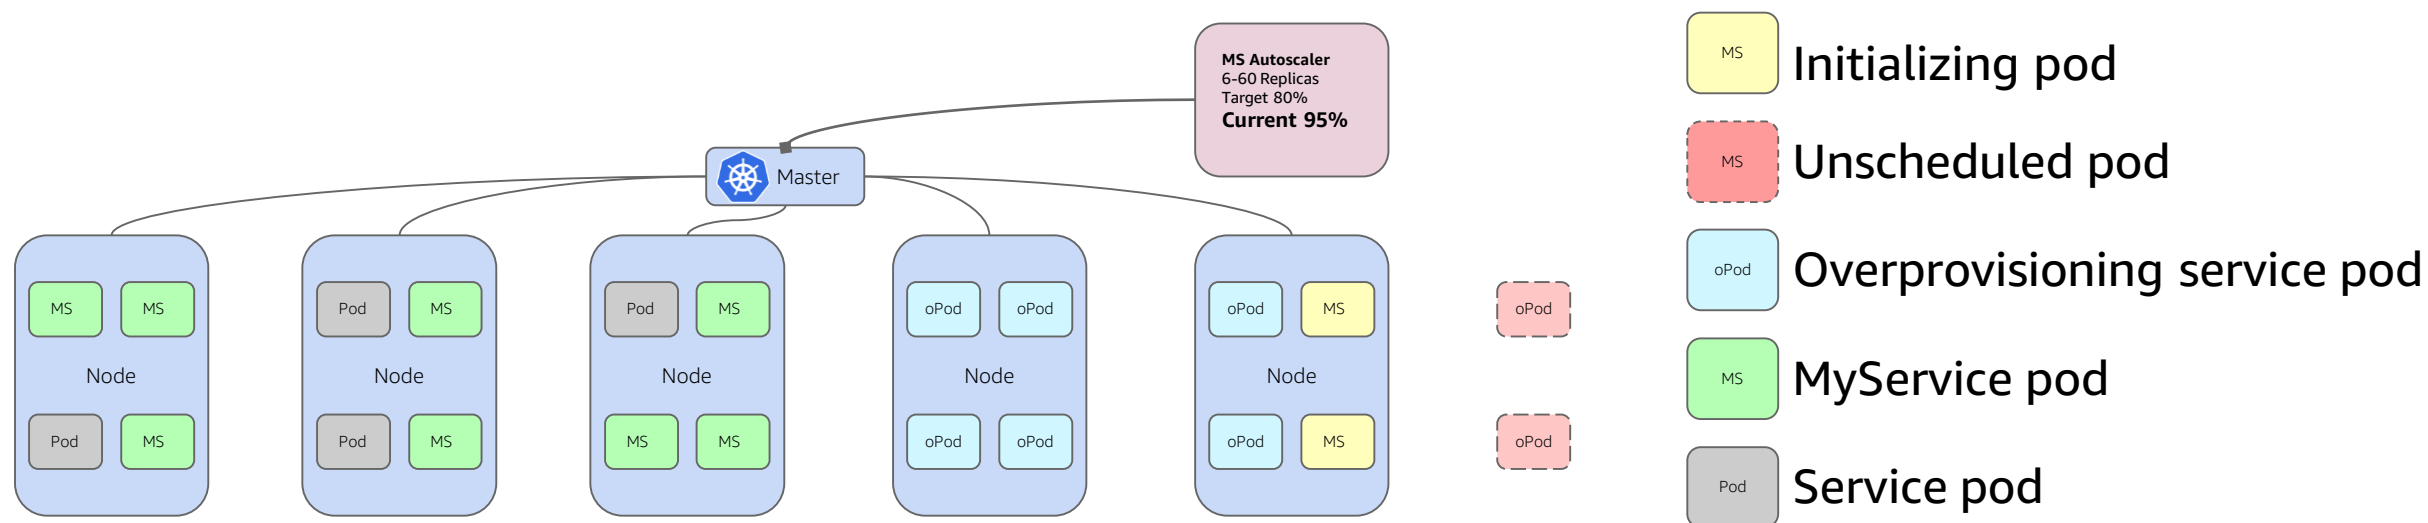
Cash platform environment: Kubernetes autoscaling



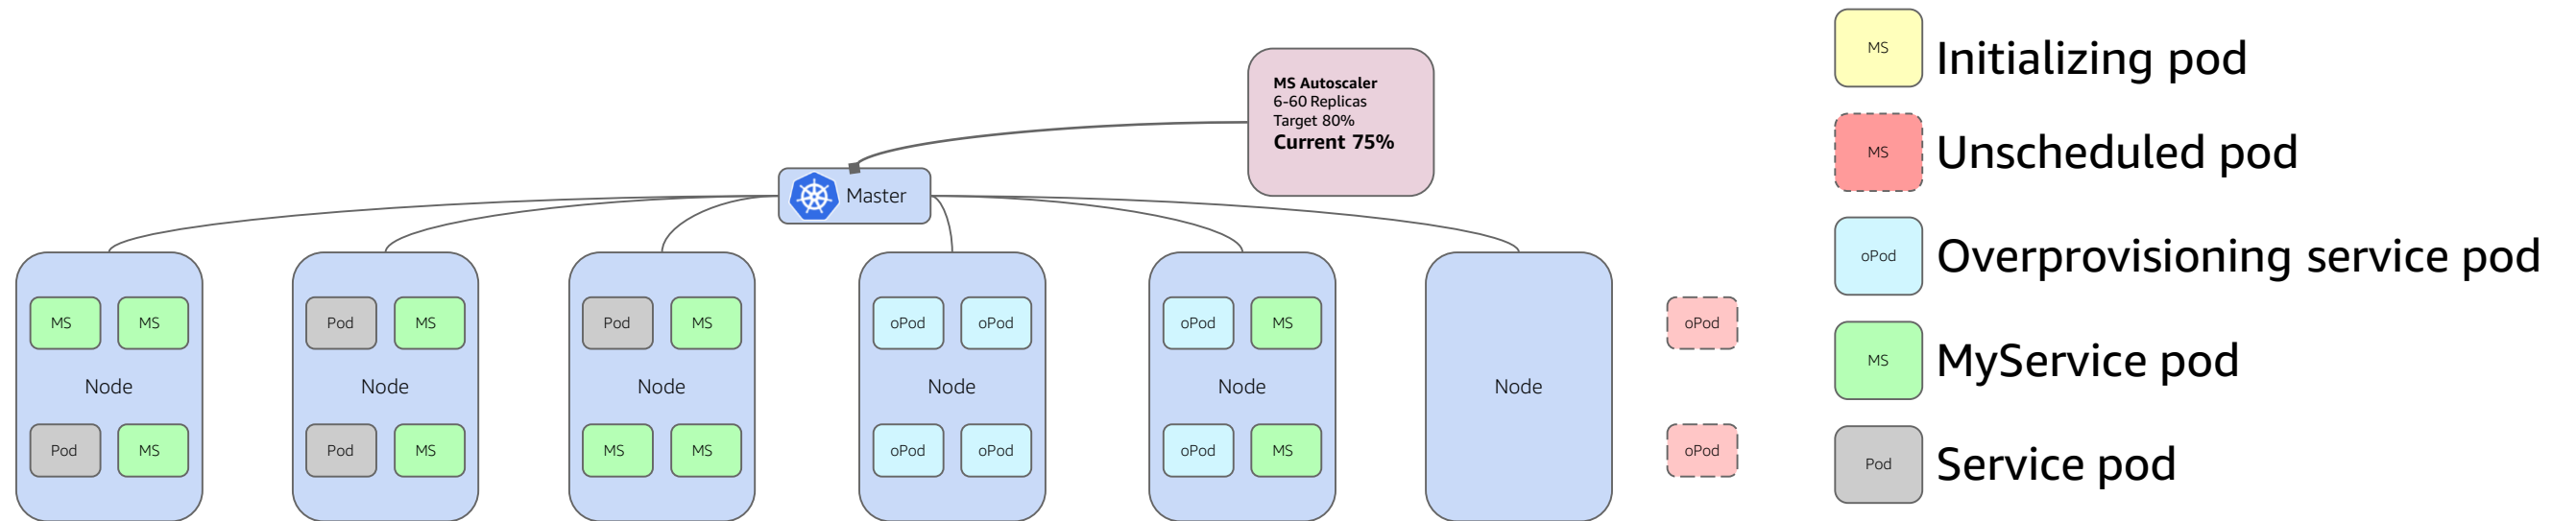
Cash platform environment: Kubernetes autoscaling



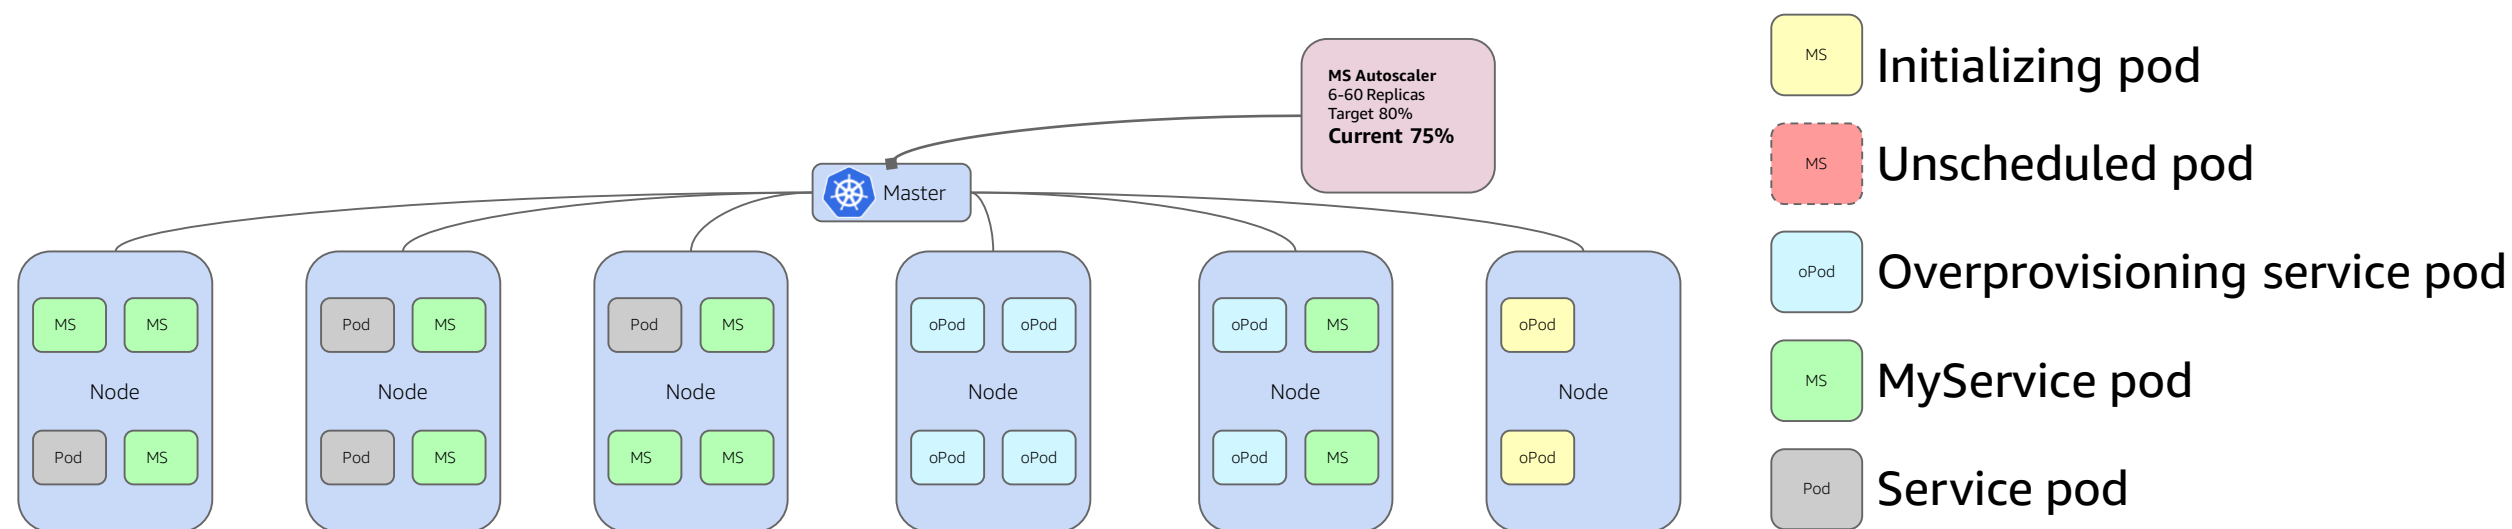
Cash platform environment: Kubernetes autoscaling



Cash platform environment: Kubernetes autoscaling



Cash platform environment: Kubernetes autoscaling



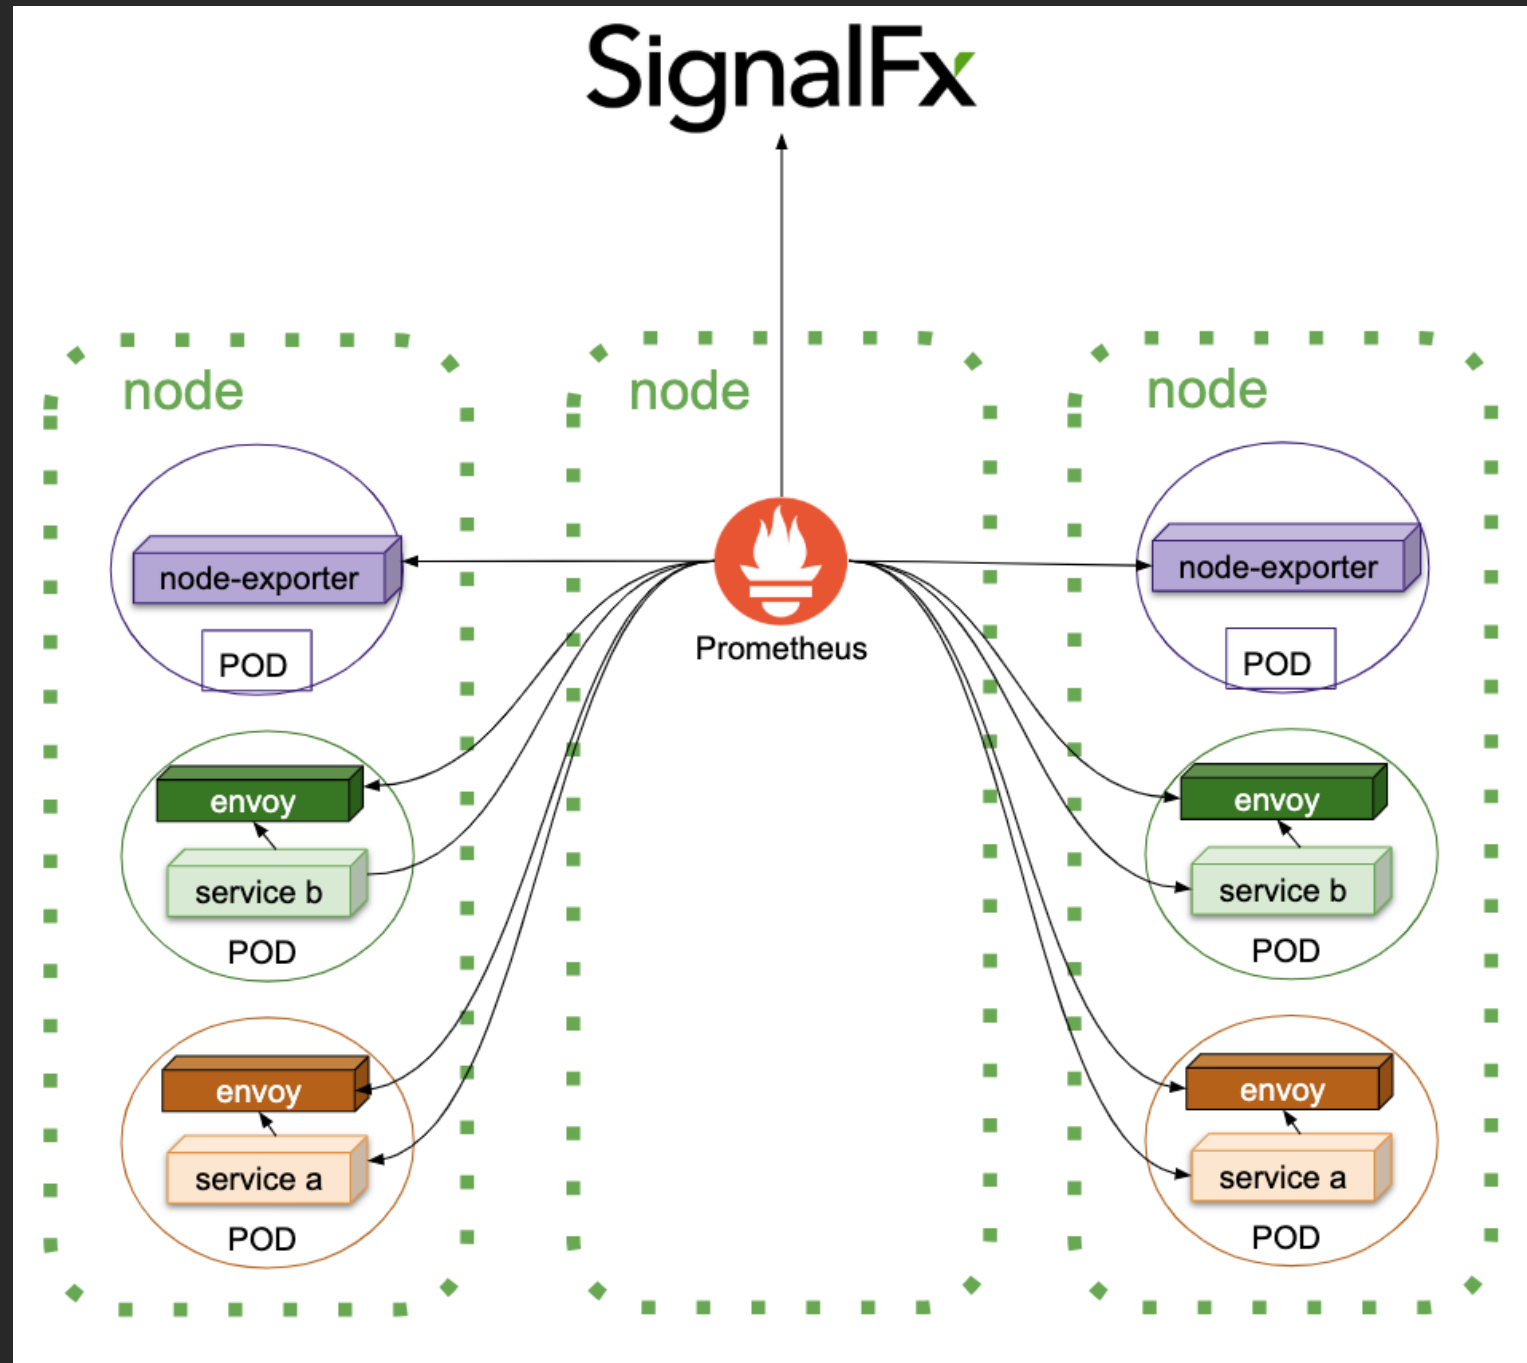
Metrics & alerting

Cash platform: Metrics & alerting



pagerduty

Cash platform: Metrics & alerting



Cash platform: Metrics & alerting

namespaceSelector:

any: true

selector:

matchExpressions:

- key: cash_app

operator: Exists

Cash platform: Metrics & alerting

kind: DaemonSet

template:

spec:

containers:

- image: quay.io/prometheus/node-exporter:v0.17.0

name: node-exporter

Cash platform: Metrics & alerting

kind: Deployment

template:

spec:

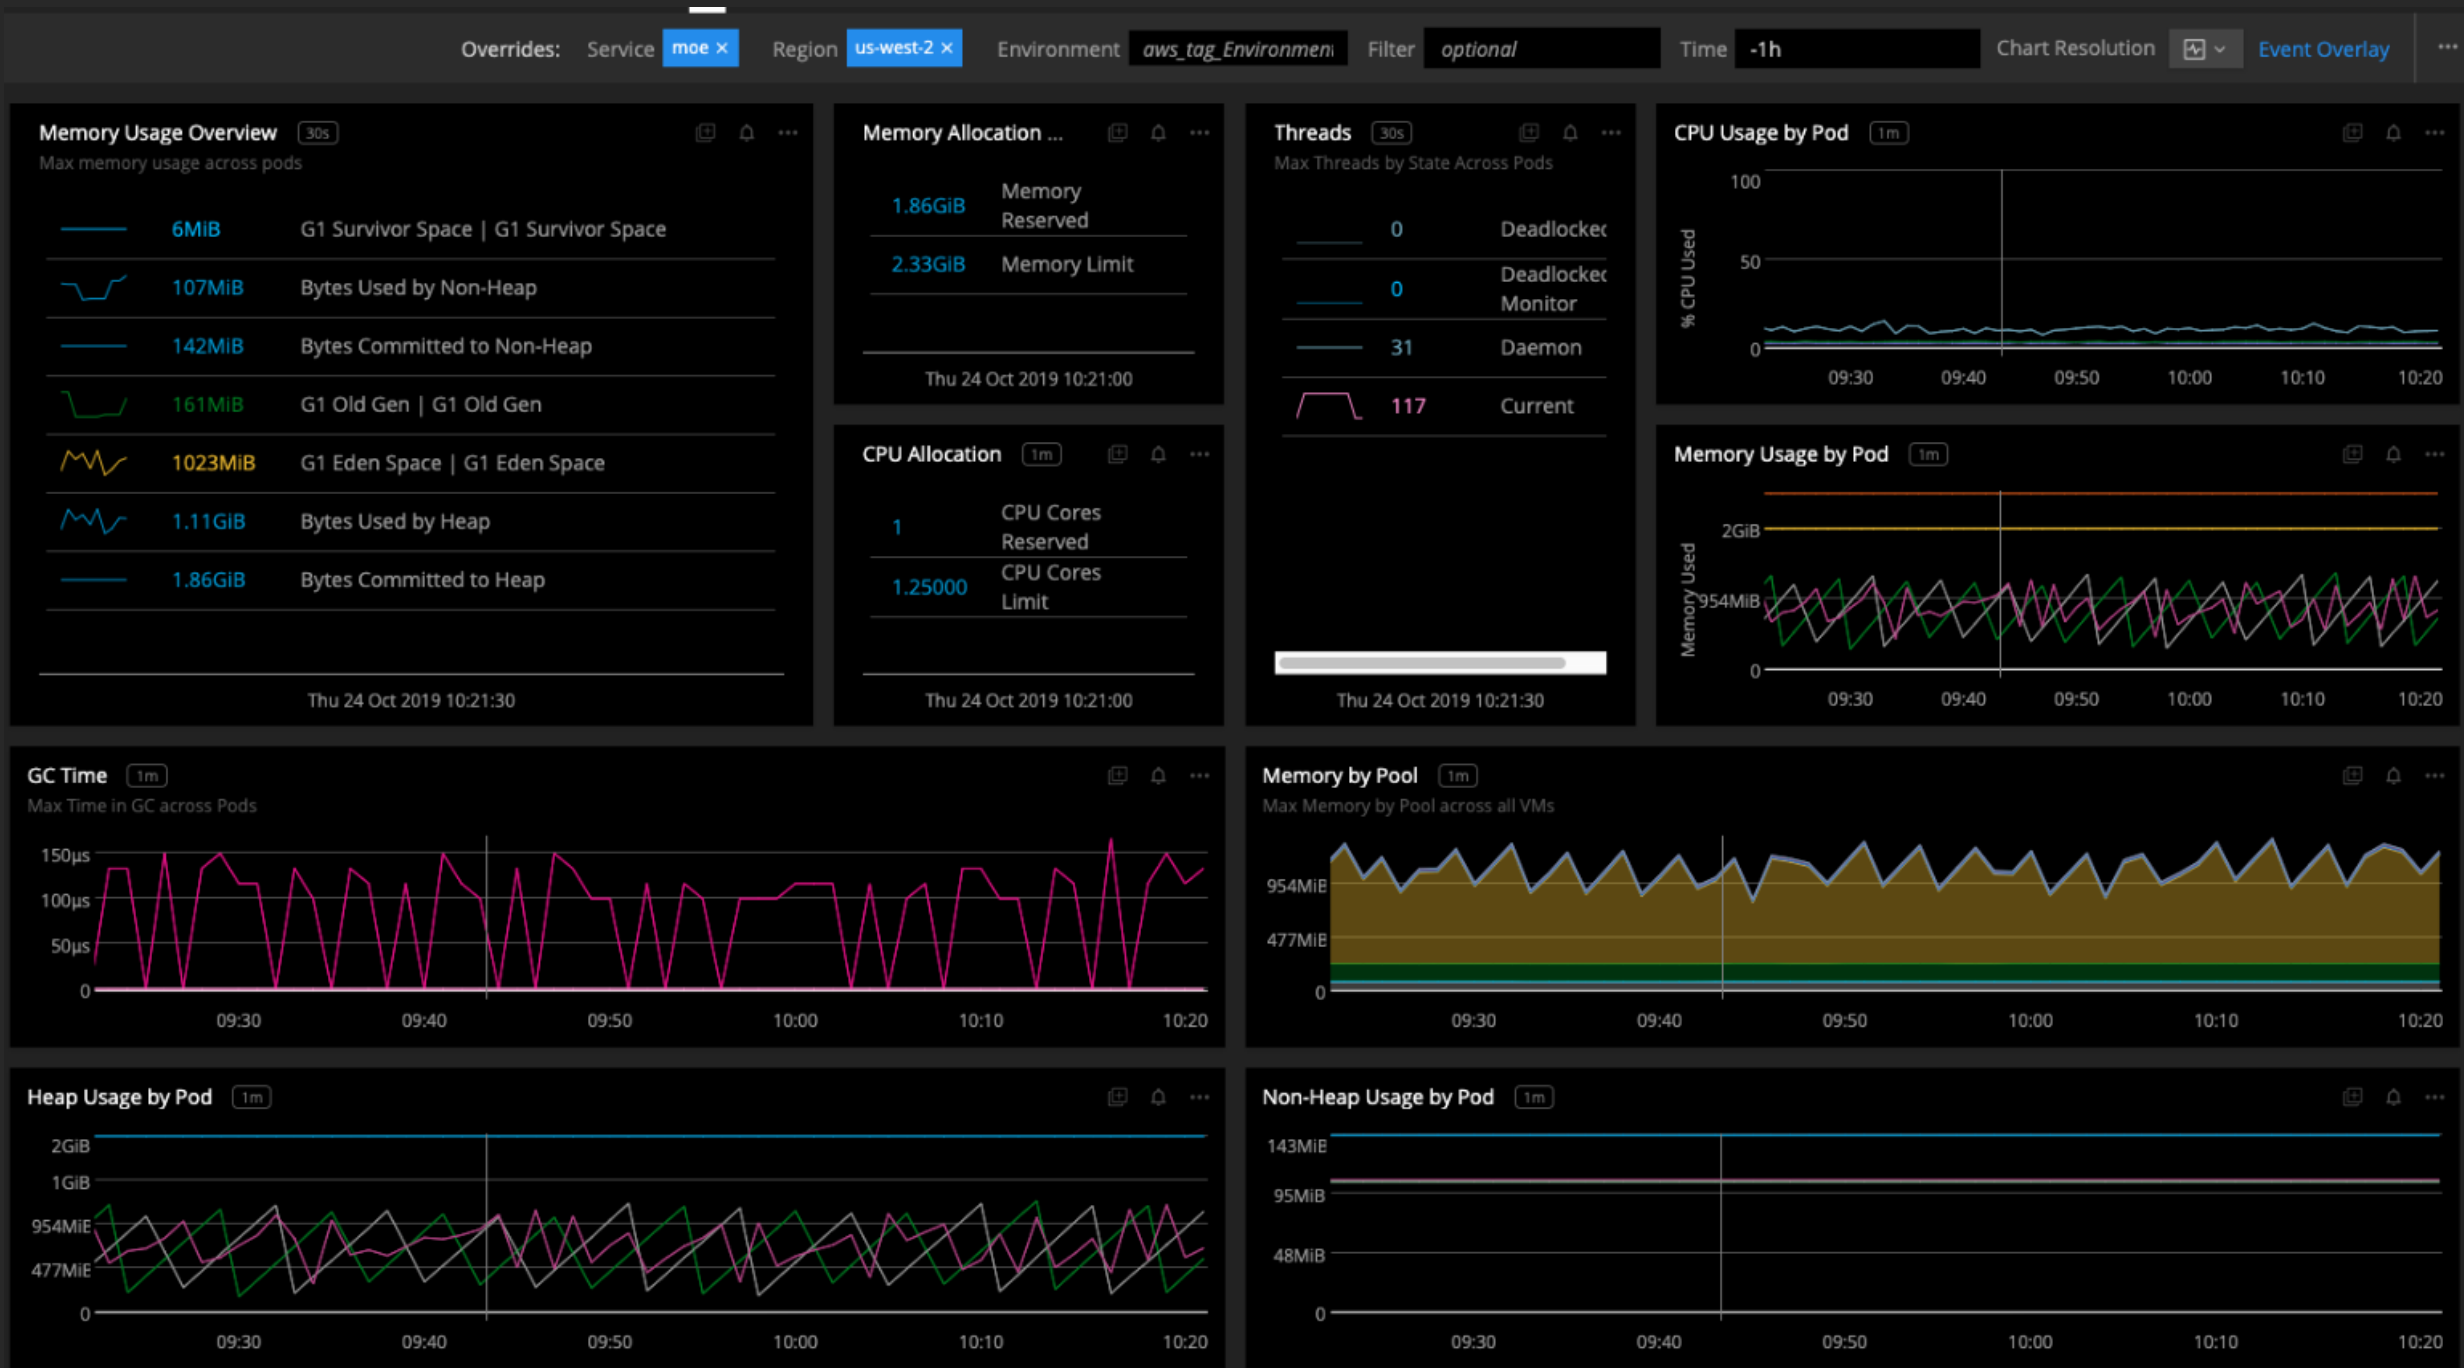
containers:

- image: quay.io/signalfx/gateway

Cash platform: Metrics & alerting

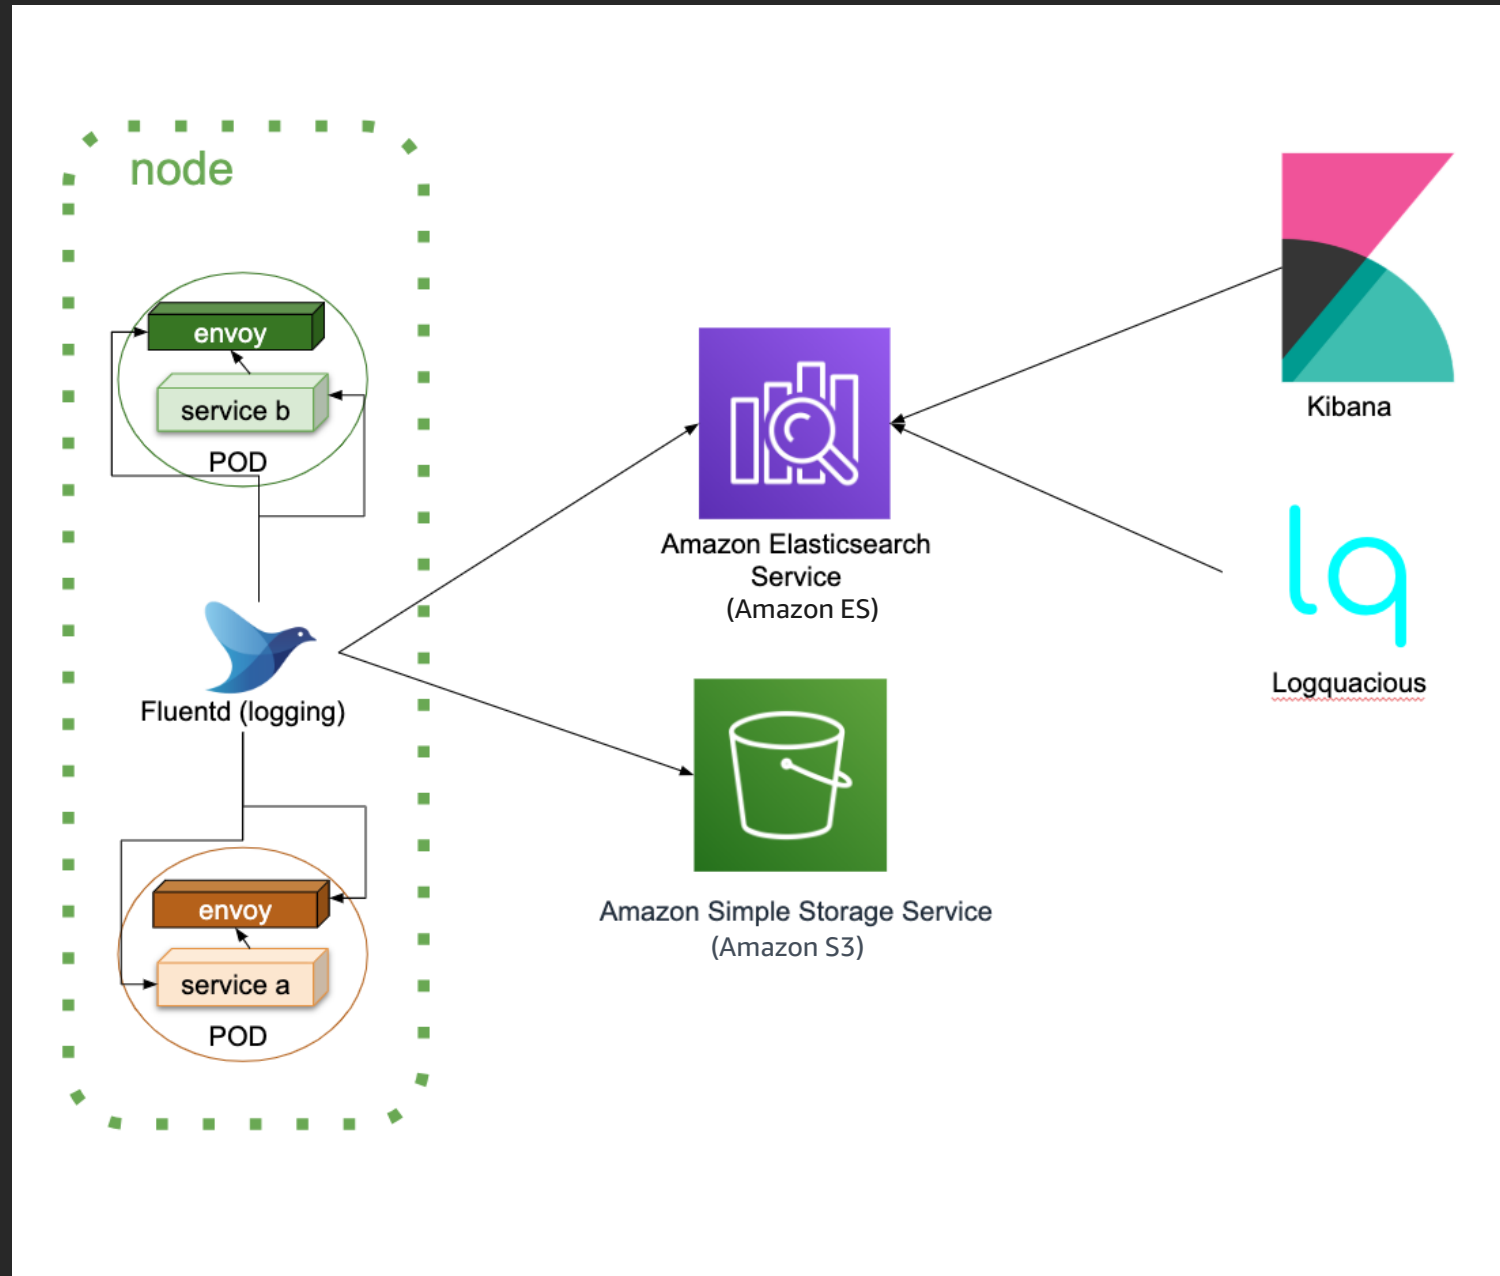
```
resource "signalfx_dashboard_group" "moe" {  
    name      = "Moe (Terraformed)"  
    description = "Dashboard for Moe"  
}  
  
module "larry_nyuk_nyuk_5xx_detector" {  
    source = "../../../../../modules/detectors/rpc/outgoing/5xx_ratio"  
  
    team    = var.team  
  
    service = local.service  
  
    action  = "larry.nyuknyuk"  
}
```

Cash platform: Metrics & alerting

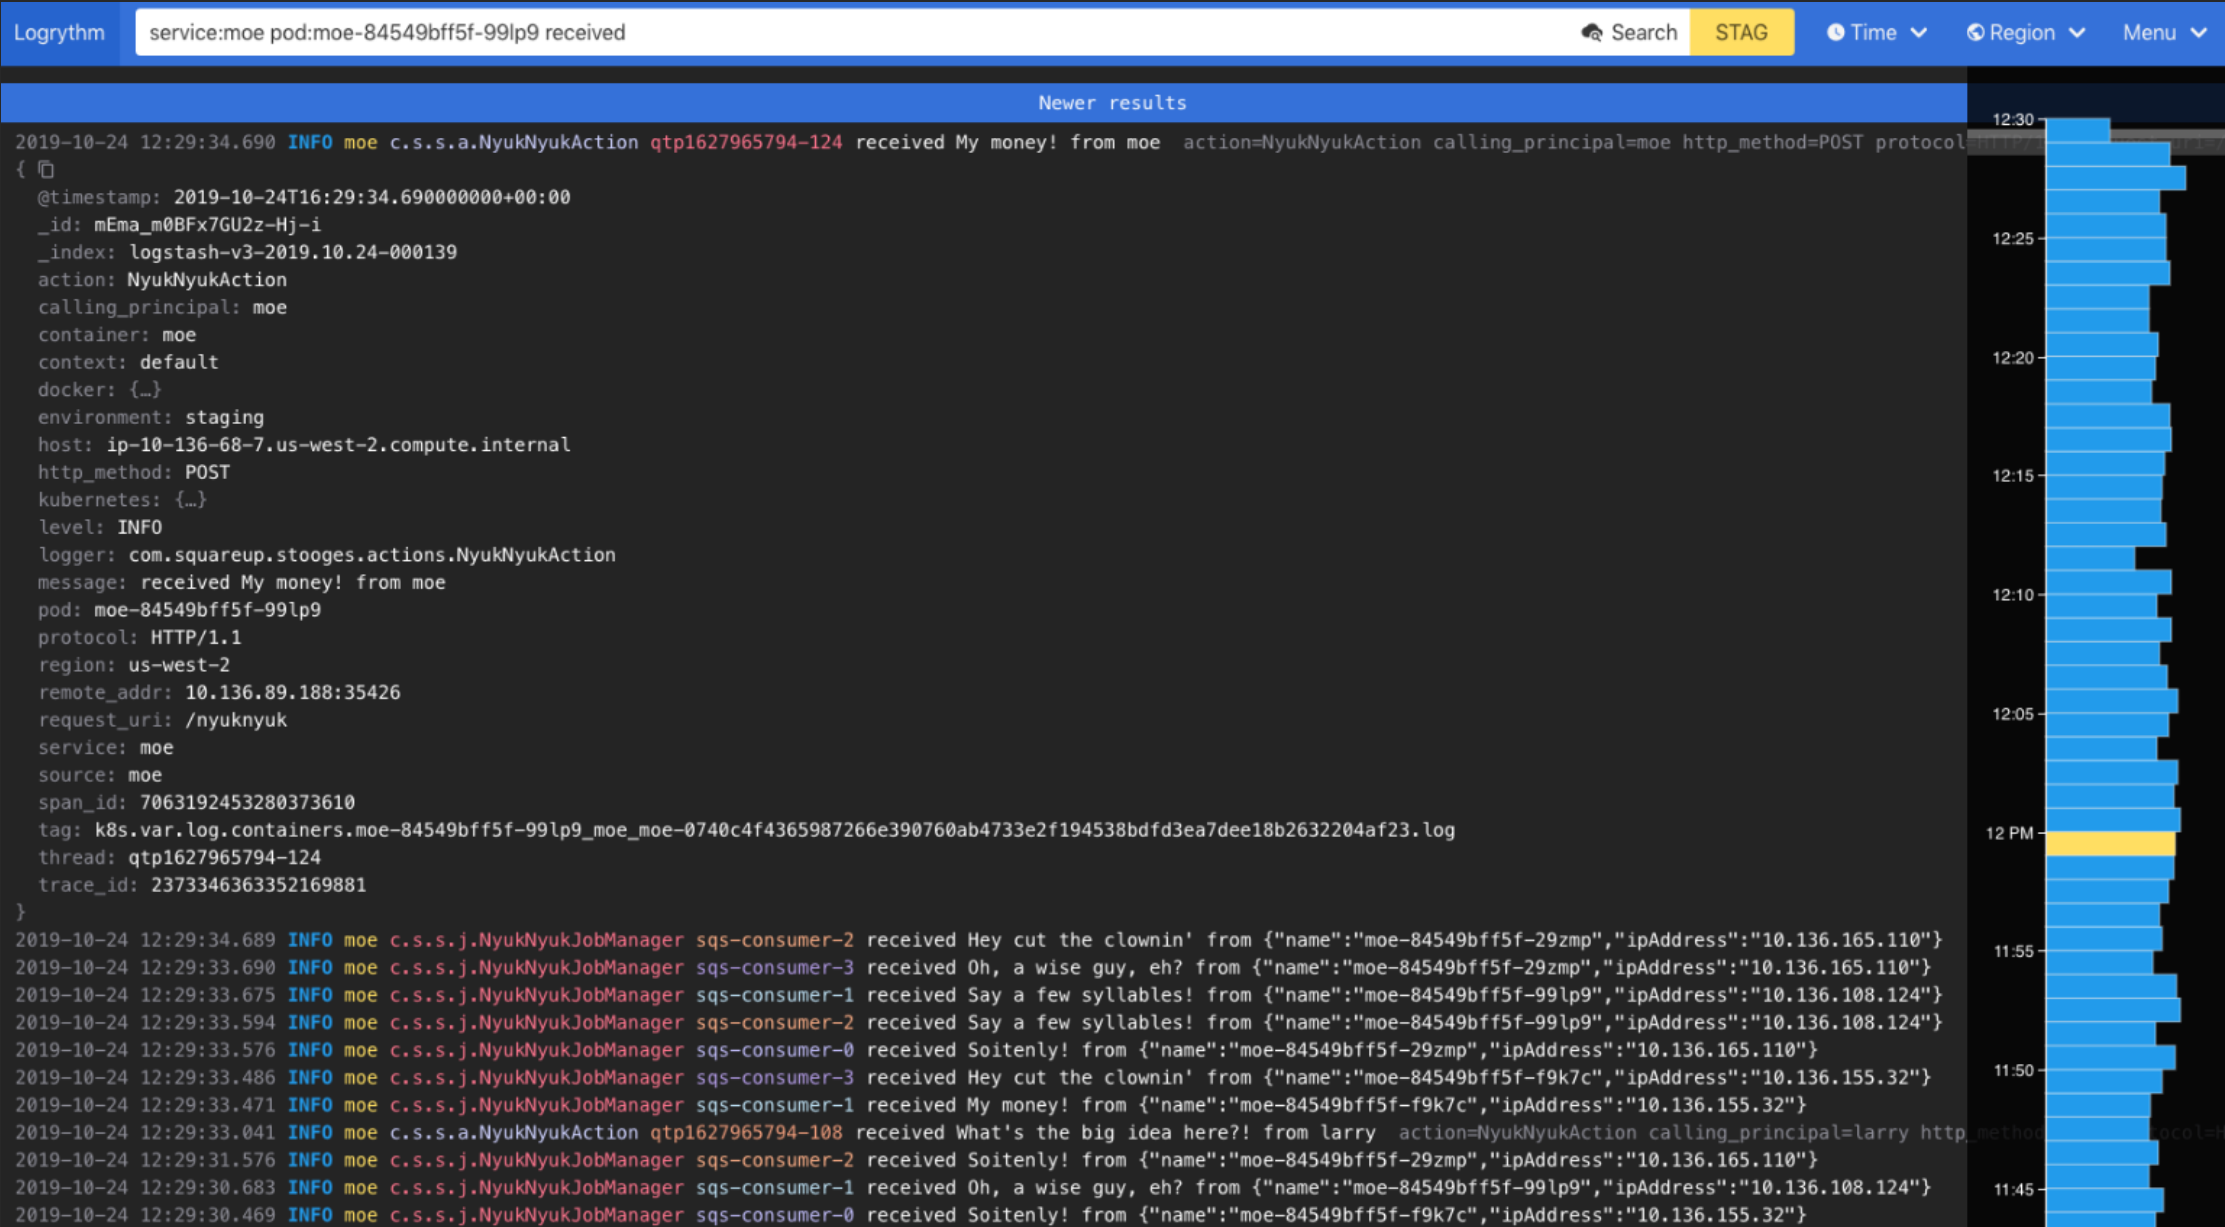


Logging & tracing


Cash platform environment: Logging



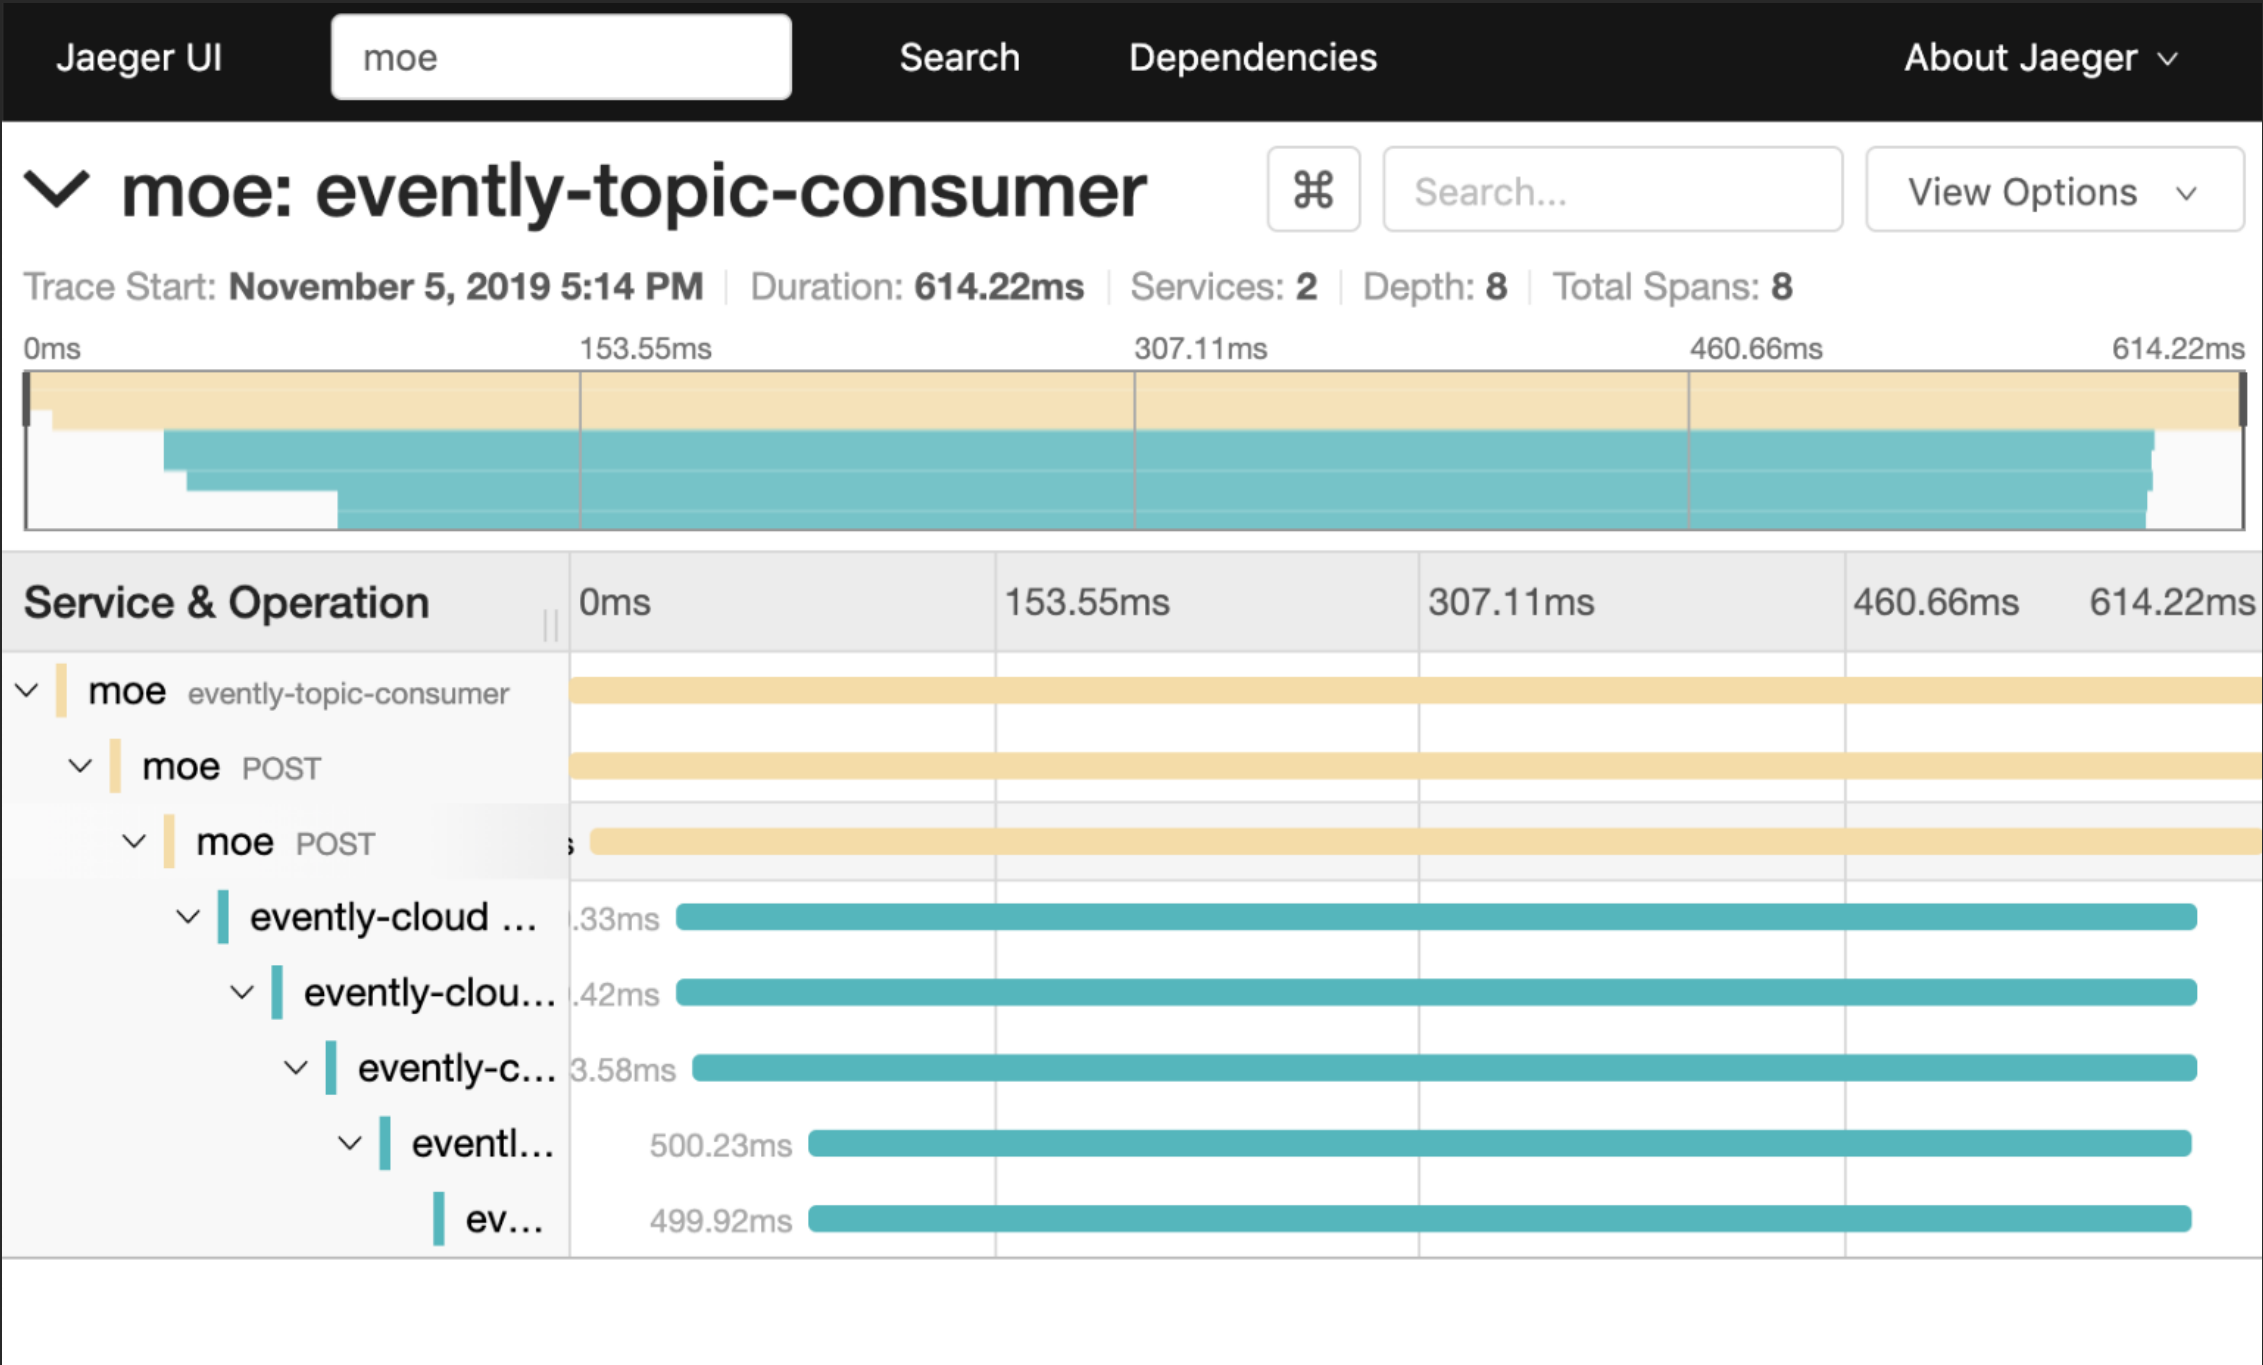
Cash platform environment: Metrics & alerting



Cash platform: Tracing




Cash platform: Tracing



Self-service tooling

Provision a service

```
sqm service create --service-name=jay-reinvent --  
env=staging --registry-owner-role=cash-product-platform-  
eng --skip-sentry-setup
```



Provision a database

```
sqm database create --service-name=jay-reinvent --  
env=staging --cluster-id=jay-reinvent-001 --database-  
name=jayreinvent --replicas=1
```



Engineering empowerment

Engineering velocity

Minimize time from idea to production

105 provisioned services since mid-2018

Focus on application code

Opinionated application framework

Strong tooling

Provision higher-level concepts like a database, queue, and service

Architecture

Publish domain events

Publish an event containing what was transformed

Subscribe to domain events

Transform payload to domain relevant context, persist

RPCs

Our mobile apps are synchronous

Migration

Breaking up the monolith

Two extremely large services

- App traffic

- Ledgering/card transactions

- Hinders developer velocity

Cloud

- Platform is ready

- Frameworks are ready

Build it

- Breaking up the monolith

- Avoiding more code in the monolith

Breaking up the monolith

Identify targets

Determine what domains belong to our monolith

Moving data

Backfills

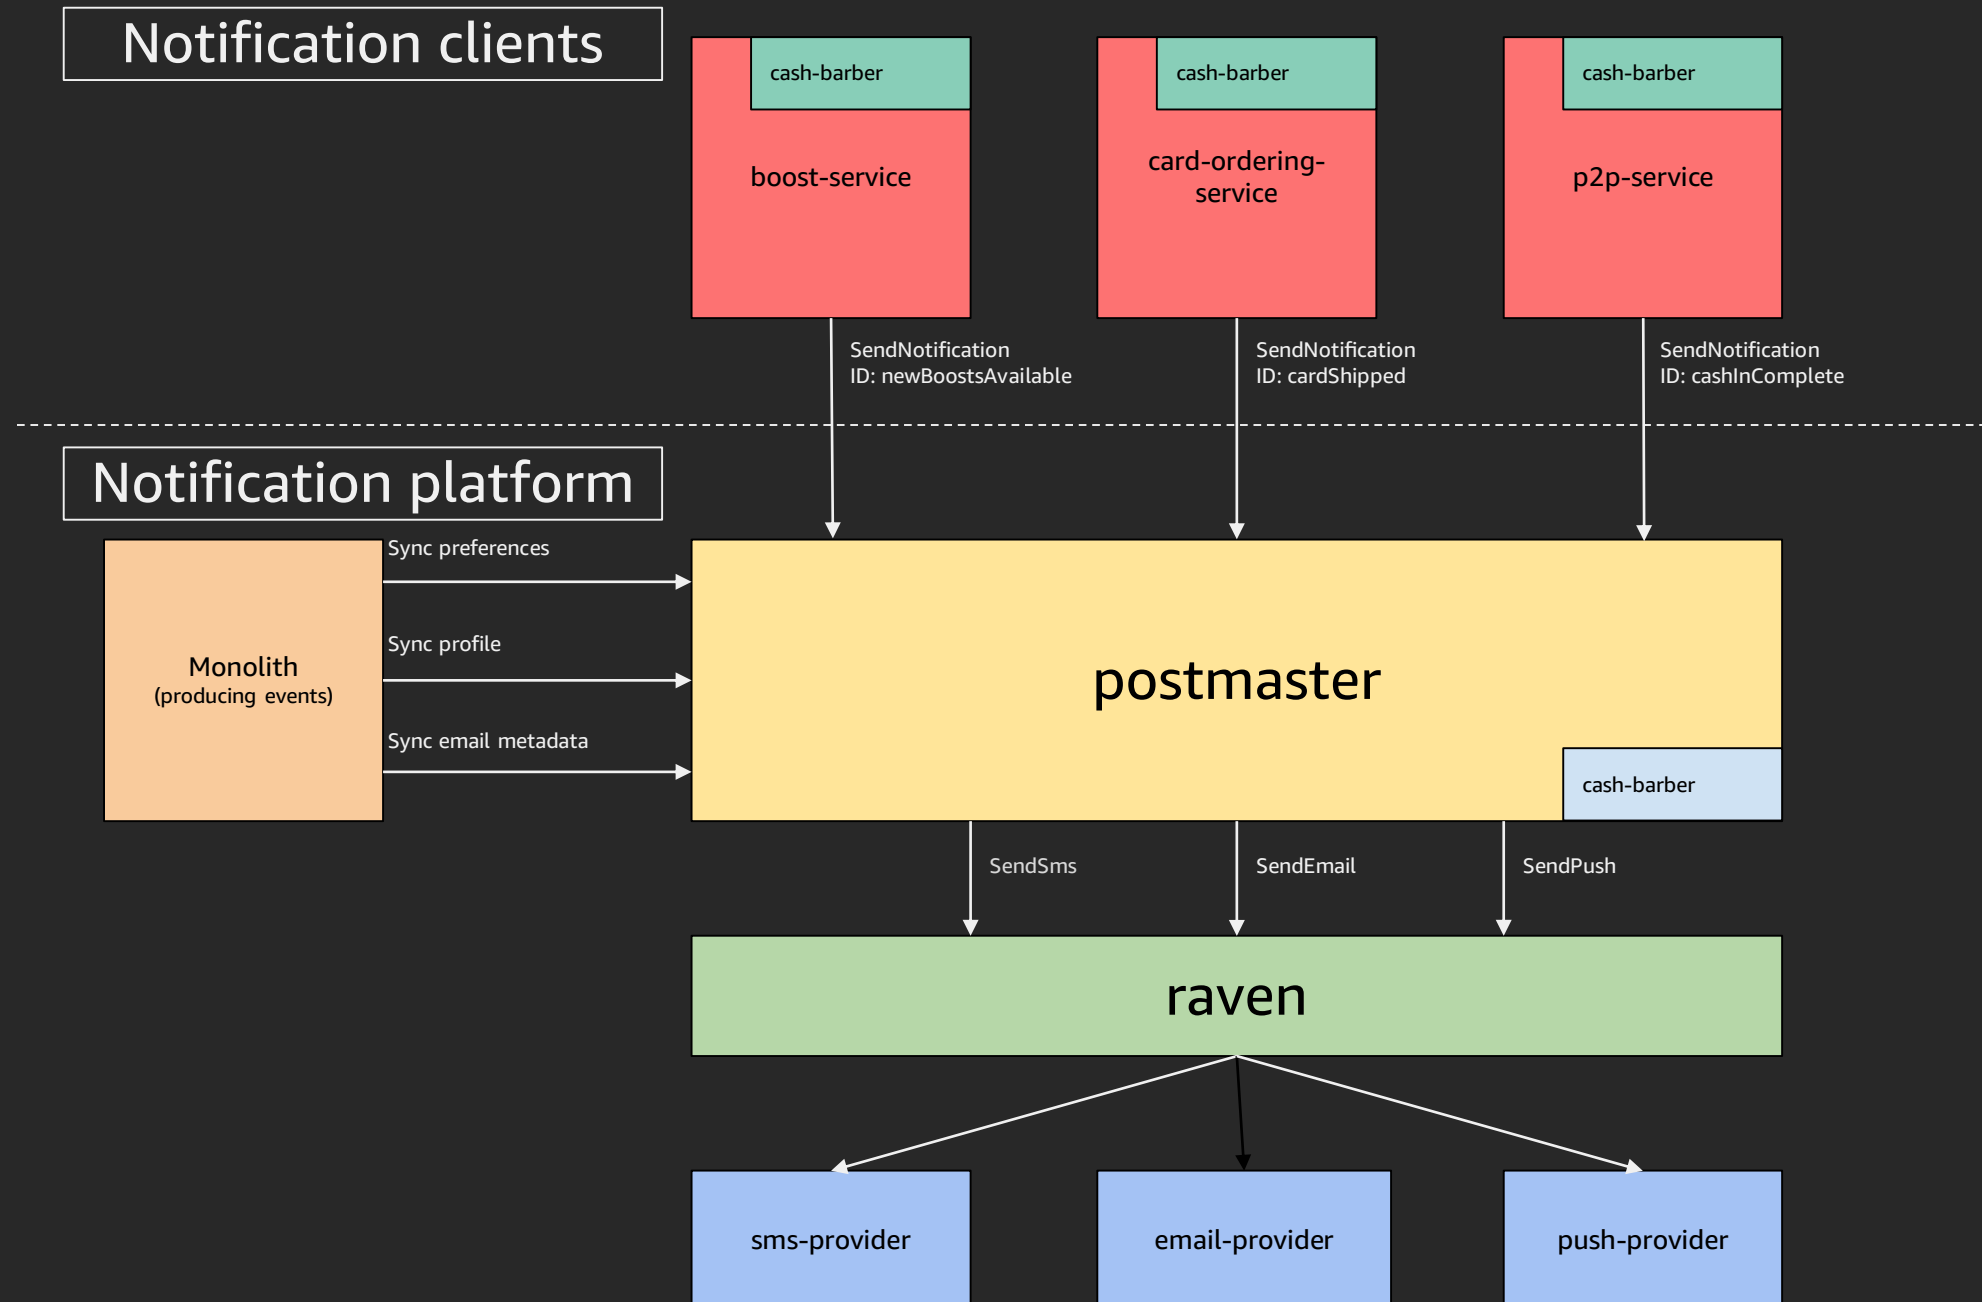
Dual reading/writing

Two copies, one source of truth

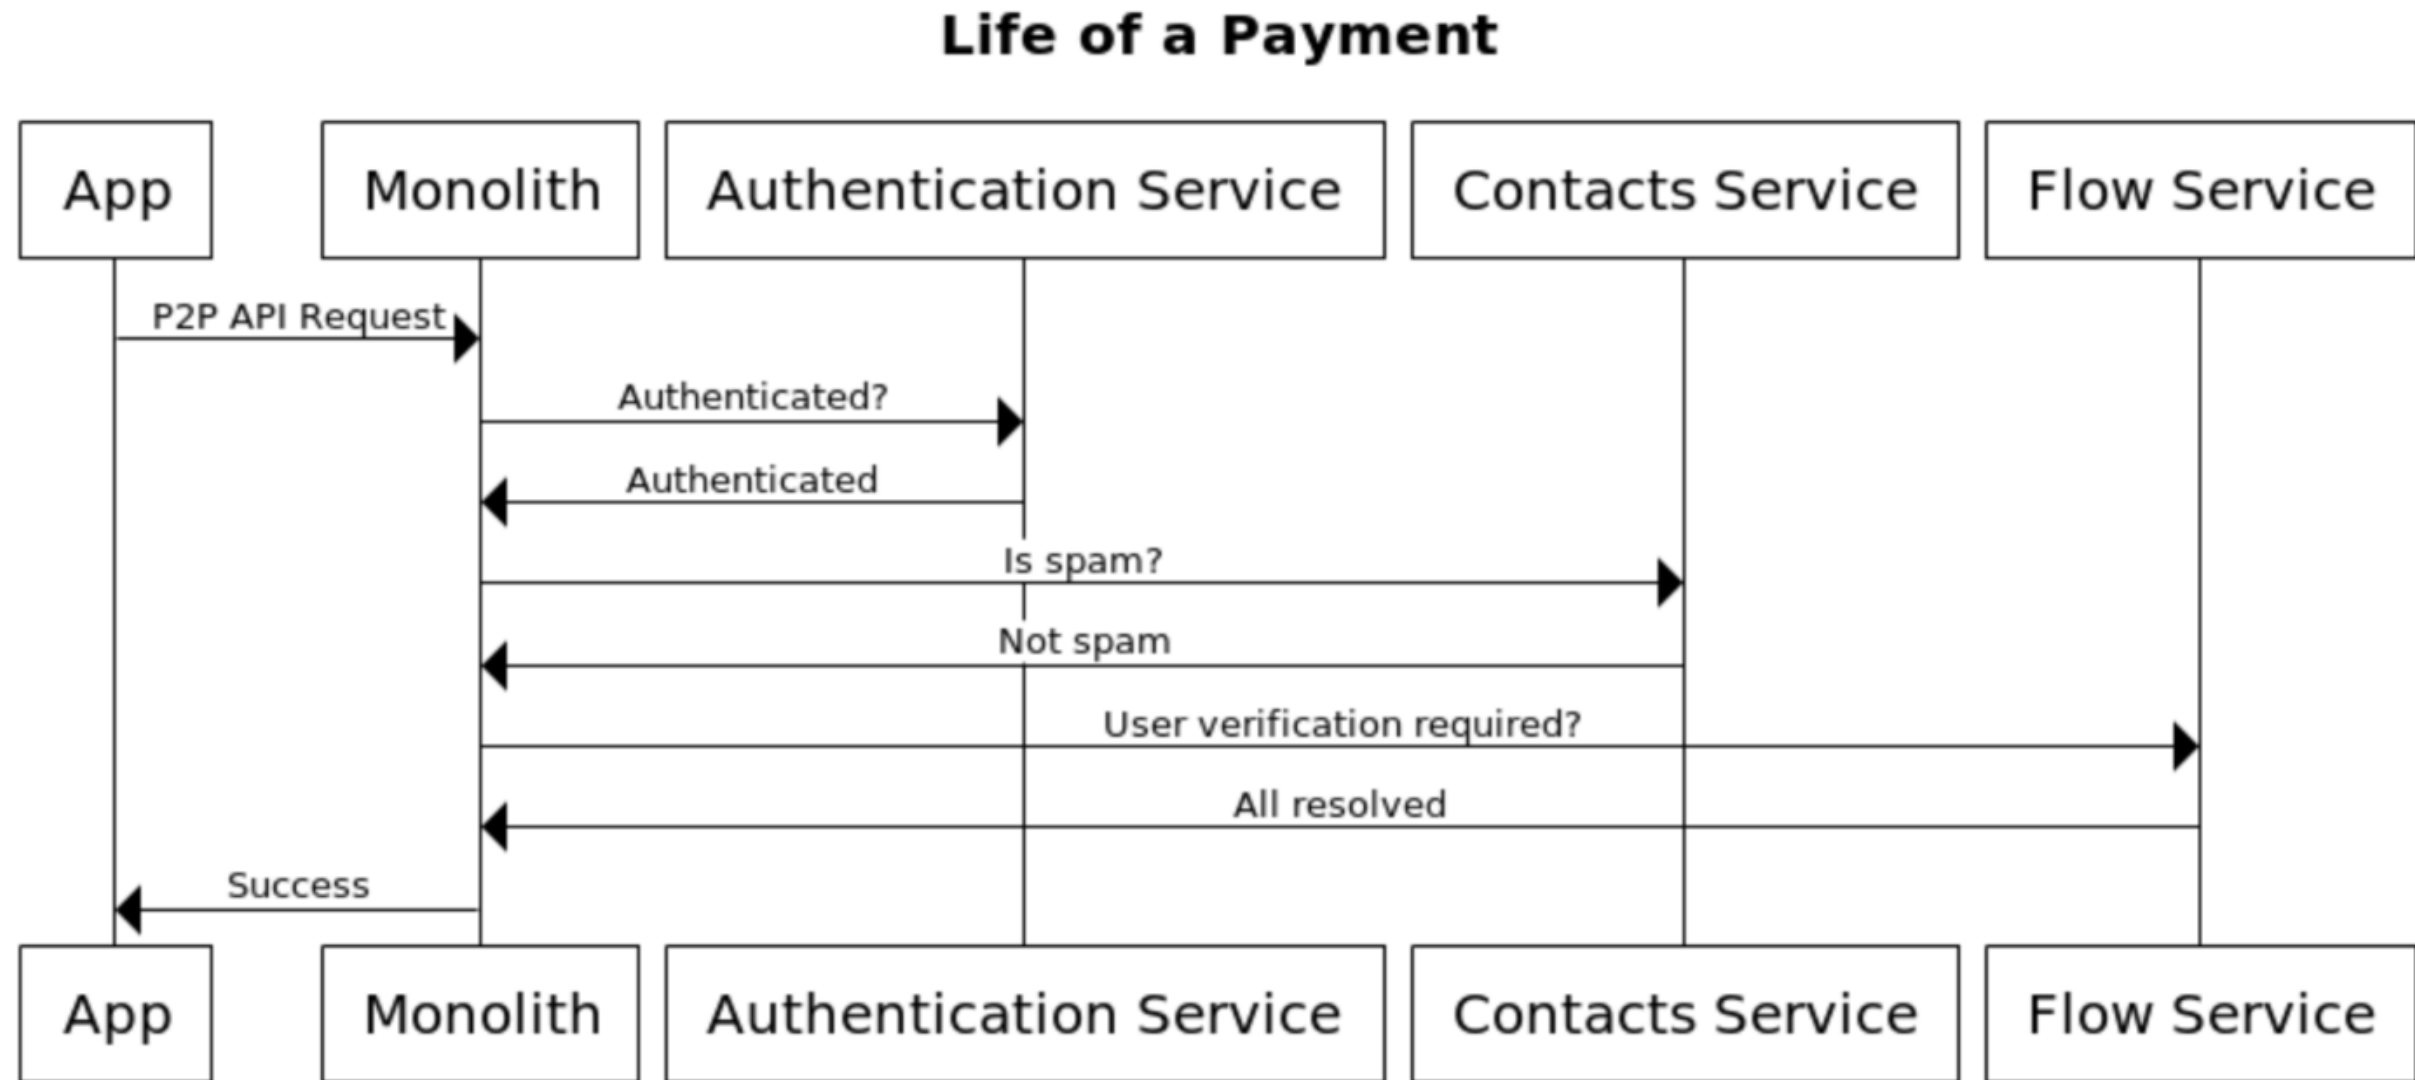
One diff

Migration examples

Life of a notification



Life of a payment



One last thing

Open sourcing the Cash App platform

github.com/cashapp

Cash App team is committed to open sourcing as much of our cloud platform as possible



Open sourcing the Cash App platform: misk

github.com/cashapp/misk

Kotlin-based Kubernetes aware service container

Misk

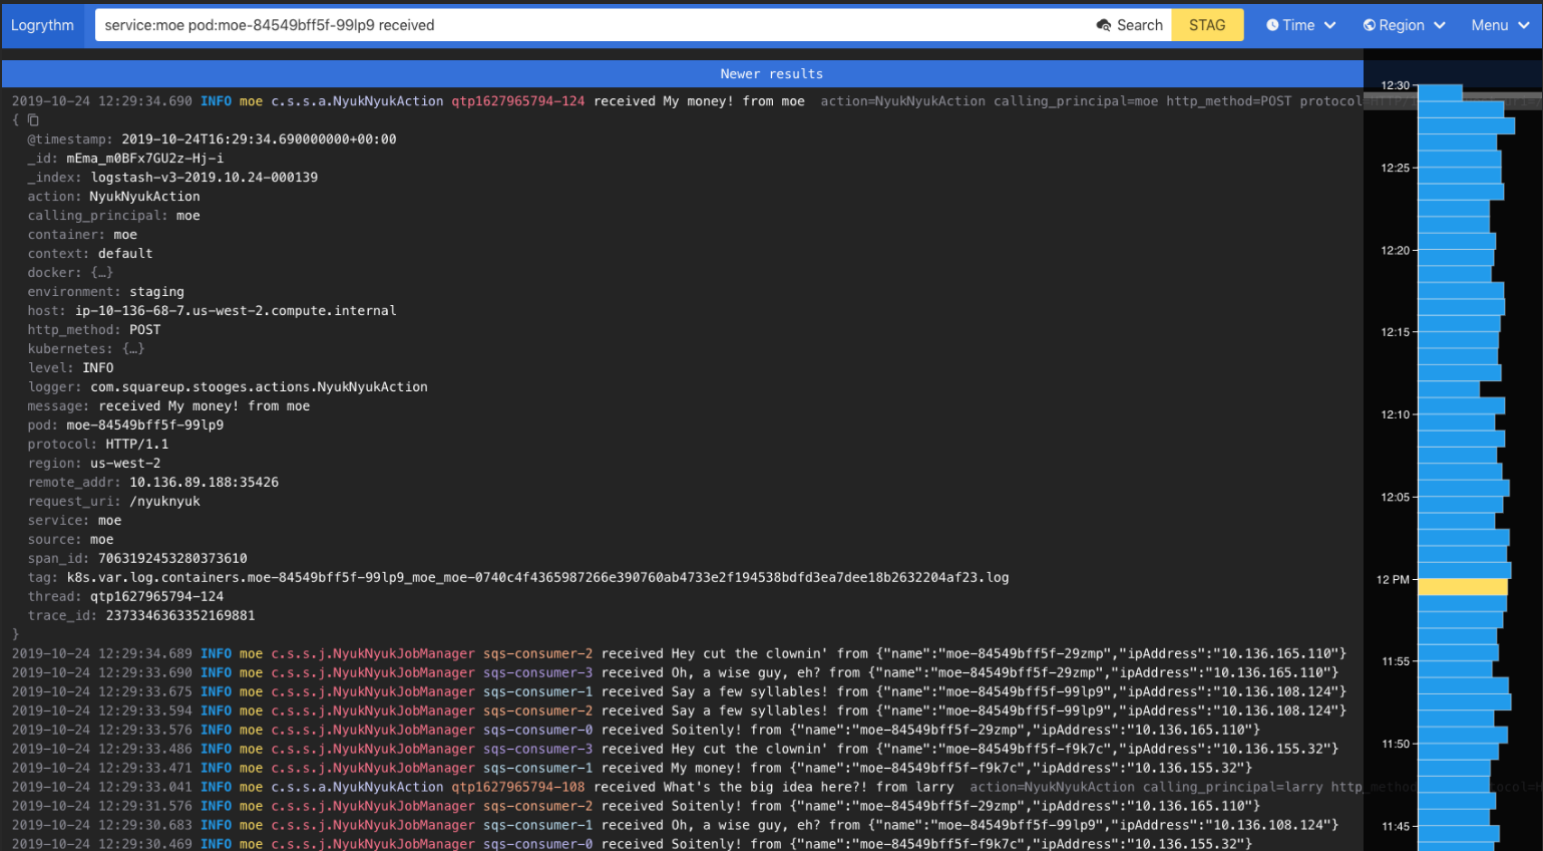
Open sourcing the Cash App platform: backfila

github.com/cashapp/backfila

A data migrator / batched work coordinator that runs over http

Open sourcing the Cash App platform: lq

Logquacious (lq) is a fast and simple Elasticsearch-backed log viewer



github.com/cashapp/logquacious

Open sourcing the Cash App platform: kruit

Kubernetes Realtime UI Toolkit: Build real-time user interfaces for Kubernetes easily



github.com/cashapp/kruit

Open sourcing the Cash App platform: csop

Cash Service Operator: Makes microservices operation simple and secure

us-west-2 (100%)		<div>Autoscaling</div> <div>Restart</div> <div>Deploy</div> <div>Scale</div> <div>Stop</div>			
Autoscaling	disabled				
Image	moe @ 3f59947a envoy @ b12a8ef3				
Deployed at	Wed 6-Nov-19 4:39 AM				
State	MinimumReplicasAvailable - Deployment has minimum availability.				
Last Change	restart on Wed 6-Nov-19 4:39 AM by rotate-aws-credentials-job				
Replicas Available	3 out of 3				
Pods	<div>moe-5874975969-5whz9</div> <div>moe-5874975969-q5szh</div> <div>moe-5874975969-zfrqc</div>				

Deploy moe in us-west-2

SHA to Deploy

fa0920efaf9cda2037c64b08762de3cd91ac7ed5

Deploy Diff

Deploying moe from SHA `3f59947a7f024b160dab1456fb2dec447b13548b` to SHA `fa0920efaf9cda2037c64b08762de3cd91ac7ed5`

Changes

`100b6ba5460` Adopt the polyrepo plugin for artifact publishing

Also upgrade Gradle to 5.6.4 @jwilson (2019-11-03 03:04:43 +0000 UTC)

`8f6a9f496a1` Fix AdminDashboardAccess import @adrw (2019-10-18 18:15:13 +0000 UTC)

`87ecf2b17bd` Bump Misk and Skim with AdminDashboard changes @adrw (2019-10-18 18:02:59 +0000 UTC)

`1cbf62f84e4` Explicitly depend on misk-aws @rhall (2019-10-12 00:37:56 +0000 UTC)

Dependency Updates

`fa0920efaf9` 2019-11-05 05:14:28 +0000 UTC

`51063f951e3` 2019-11-04 20:15:24 +0000 UTC

Close Deploy

github.com/cashapp/csop

Thank you!



Please complete the session
survey in the mobile app.