PhraseApp

# Kubernetes on AWS

Tobias Schwab, Co-Founder of PhraseApp

# ops@phraserapp

- 👎 operating state
  - Backups
  - Migrations
  - Recovery
- 👍 self-healing, scalability, transparency
  - throw away infrastructure

# How we buy from AWS

1. How much?
    - operational time do we save?
    - how much is that worth?
2. How critical?
3. How convenient and flexible
4. How much lock-in

# Road to Kubernetes @ PhraseApp

- 2013: Docker in Production
- 2014: AMIs, AutoScalingGroups, Cloudformation
- 2015: Wunderproxy
- 2016: ECS vs. Kubernetes

# Kubernetes?

**Framework** to build
**distributed applications**

“The distributed Linux of the Cloud”

# Why Kubernetes?

- **API:** open, simple and awesome
- **runs (almost) everywhere**:
  - Localhost
  - Raspberry PI
  - bare metal
  - IAAS/PAAS
- **Ecosystem:** Plugins (e.g. Helm), Operators
- **CNCF:** Cloud Native Compute Foundation

Platinum Members

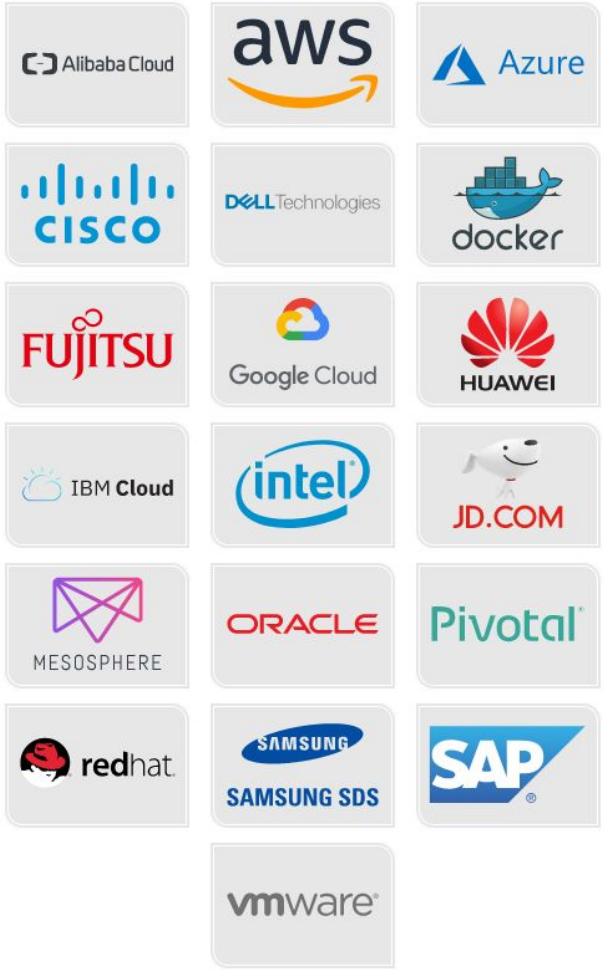

## Graduated

Kubernetes

Orchestration

Prometheus

Monitoring

## Incubating

OpenTracing

Distributed Tracing API

Fluentd

Logging

gRPC

Remote Procedure Call

containerd

Container Runtime

rkt

Container Runtime

CNI

Networking API

Envoy

Service Mesh

Jaeger

Distributed Tracing

Notary

Security

TUF

Software Update Spec

Vitess

Storage

CoreDNS

Service Discovery

NATS

Messaging

Linkerd

Service Mesh

Helm

Package Management

Rook

Storage

# Why not Google, Azure ...?

### Compute
EC2
Lightsail
Elastic Container Service
EKS
Lambda
Batch
Elastic Beanstalk

### Storage
S3
EFS
Glacier
Storage Gateway

### Database
RDS
DynamoDB
ElastiCache
Neptune
Amazon Redshift

### Migration
AWS Migration Hub
Application Discovery Service
Database Migration Service
Server Migration Service
Snowball

### Networking & Content Delivery
VPC
CloudFront
Route 53
API Gateway
Direct Connect

### Developer Tools
CodeStar
CodeCommit
CodeBuild
CodeDeploy
CodePipeline
Cloud9
X-Ray

### Management Tools
CloudWatch
AWS Auto Scaling
CloudFormation
CloudTrail
Config
OpsWorks
Service Catalog
Systems Manager
Trusted Advisor
Managed Services

### Media Services
Elastic Transcoder
Kinesis Video Streams
MediaConvert
MediaLive
MediaPackage
MediaStore
MediaTailor

### Machine Learning
Amazon SageMaker
Amazon Comprehend
AWS DeepLens
Amazon Lex
Machine Learning
Amazon Polly
Rekognition
Amazon Transcribe
Amazon Translate

### Analytics
Athena
EMR
CloudSearch
Elasticsearch Service
Kinesis
QuickSight
Data Pipeline
AWS Glue

### Security, Identity & Compliance
IAM
Cognito
Secrets Manager
GuardDuty
Inspector
Amazon Macie
AWS Single Sign-On
Certificate Manager
CloudHSM
Directory Service
WAF & Shield
Artifact

### Mobile Services
Mobile Hub
AWS AppSync
Device Farm
Mobile Analytics

### AR & VR
Amazon Sumerian

### Application Integration
Step Functions
Amazon MQ
Simple Notification Service
Simple Queue Service
SWF

### Customer Engagement
Amazon Connect
Pinpoint
Simple Email Service

### Business Productivity
Alexa for Business
Amazon Chime
WorkDocs
WorkMail

### Desktop & App Streaming
WorkSpaces
AppStream 2.0

### Internet Of Things
IoT Core
IoT 1-Click
IoT Device Management
IoT Analytics
Greengrass
Amazon FreeRTOS
IoT Device Defender

### Game Development
Amazon GameLift


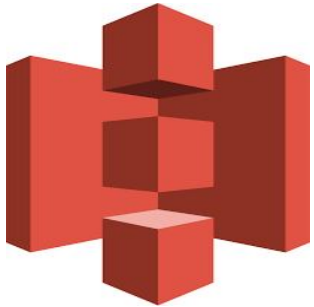



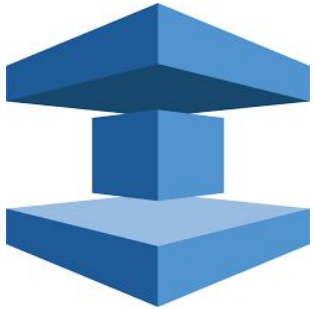

# Kubernetes

- **Node:** worker machine for Pods
- **Pod:** set of containers on the same Node
- **Label:** key value pair for all resources
- **Service**: logical set of Pods with the same Labels
- **ReplicaSet**: set of Pods with same configuration
- **Deployment**: well defined Pod Updates (via ReplicaSets)

# Kubernetes: Master

- **api-server**: front-end for control plane
- **etcd**: backing distributed storage
- **scheduler**: assign newly created pods to nodes
- **replication controller**: maintain correct number of pods in a Replica Set

# Kubernetes: Node

- **kubelet**: manages Pods scheduled on Node
- **kube-proxy**: network rules and connection forwarding

# Kubernetes: User

- **kubectl**: CLI tool to talk to the API

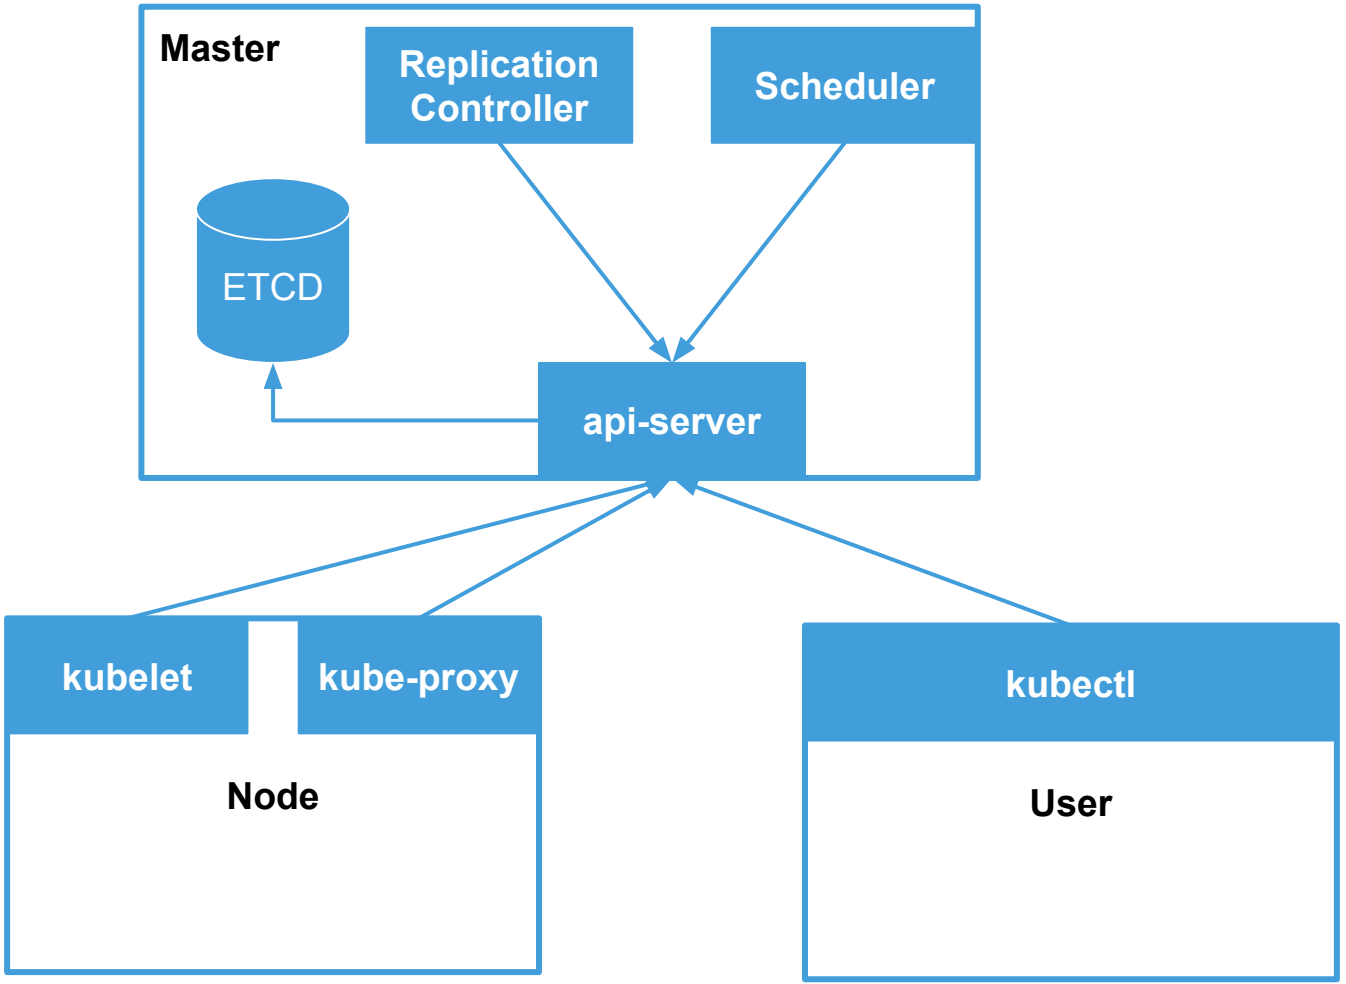
Master
Replication Controller
Scheduler
ETCD
api-server
kubelet
kube-proxy
Node
kubectl
User

# The art of running Kubernetes

- **etcd** most critical component of cluster
- **certificate management**
- **container network**
- scheduler, replication controller, api-server, kubelet all **stateless**

# Our path to Kubernetes on AWS

- Evaluation on GCP
- DIY/“how hard can it be” on AWS
- kubeadm with **UserData** and **AutoScalingGroup** based NodeGroups

# The day our cluster died 😩 😫 😣

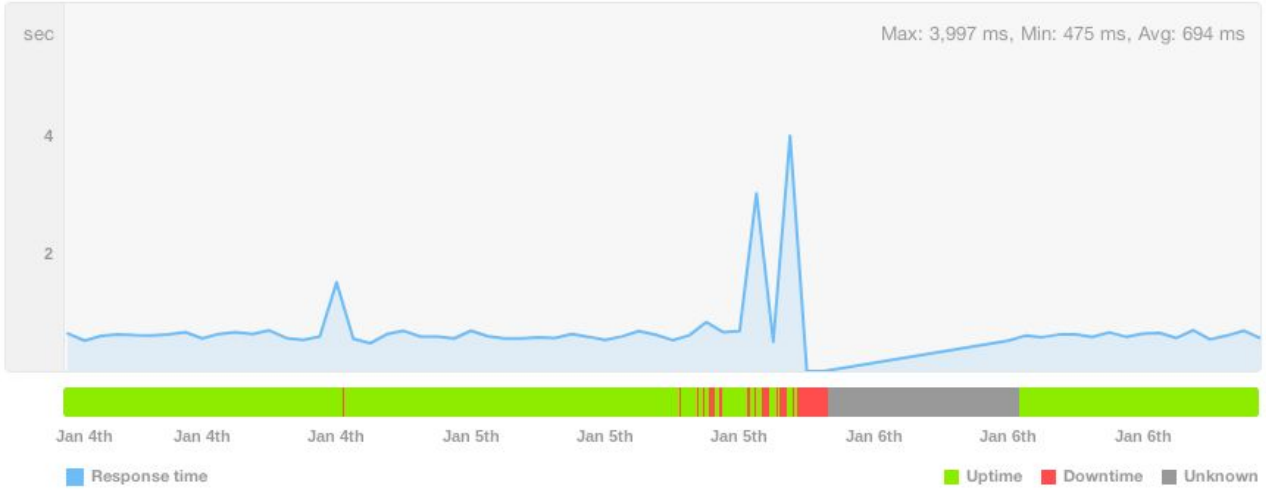


DOWNTIME

4 hours

(14 outages)

UPTIME

93.54%

# DEMO

# Kops: Kubernetes Operations

- “kubectl for clusters”
- supports AWS, GCE (beta) and VMware vSphere (alpha)
- CLI with kubernetes controllers
- cluster config stored in S3
- high availability (multi-AZ) for all components (masters and nodes)
- fully automated management of
  - certificates
  - container network (VPC RouteTable by default)
  - updates (master and nodes)

# Learnings and Issues

- always start with Multi-AZ masters!
- automatic patching of kubeconfig contexts
- IAM do not always follow “principle of least privilege”
  - ELB access on masters
  - ECR access on nodes
- shared secrets by default

# EKS: Elastic Container Service for Kubernetes

- Available in us-east-1, us-west-2 and eu-west-1 (since **2018-09-05**)
- fully Managed **Kubernetes Control Plane**
- Infrastructure and Instance Group management via **CloudFormation**
- Role based access control via **IAM**
- much tighter security by default

# EKS: Elastic Container Service for Kubernetes

- only available in **selected regions**
- only available on **AWS**
- **automatic master updates**
- **manual** setup and update process of **instance groups**
- costs: ~ 144 USD/Cluster/Month

# Kubernetes and AWS

- expose services via **ELB**
- Container images in **ECR**
- log to **Kinesis Firehose** via **fluentd**
- Log management with **Kibana/ElasticSearchService**
- persistent Pod state on **EBS** via **PersistentVolumeClaims**

# What’s next?

- Cloud Controller Manager?
- Better tooling for EKS (e.g. eksctl)?
- Fargate for EKS!

# Resources

- https://kubernetes.io/
- https://github.com/kelseyhightower/kubernetes-the-hard-way
- https://github.com/kubernetes/kops
- https://aws.amazon.com/eks/
- https://helm.sh/

PhraseApp

**Always looking for talent!**

**@tobstarr**
**tobias@phraseapp.com**
**phraseapp.com**