# Democratizing MySQL: Cloud Managed to Kubernetes Managed

KubeCon 2019

00. Who we are

Case study

## Migrate from Cloud Managed SQL to K8S managed

**Sachin Manpathak**
Technical Lead @Platform9

## The story of building Presslabs Operator for MySQL

**Flavius Mecea**
Project Lead @Presslabs

01. Context
02. The Need
03. The Solution
04. Challenges
05. Future plans



# 01. Context



# 01.1 Who is Presslabs

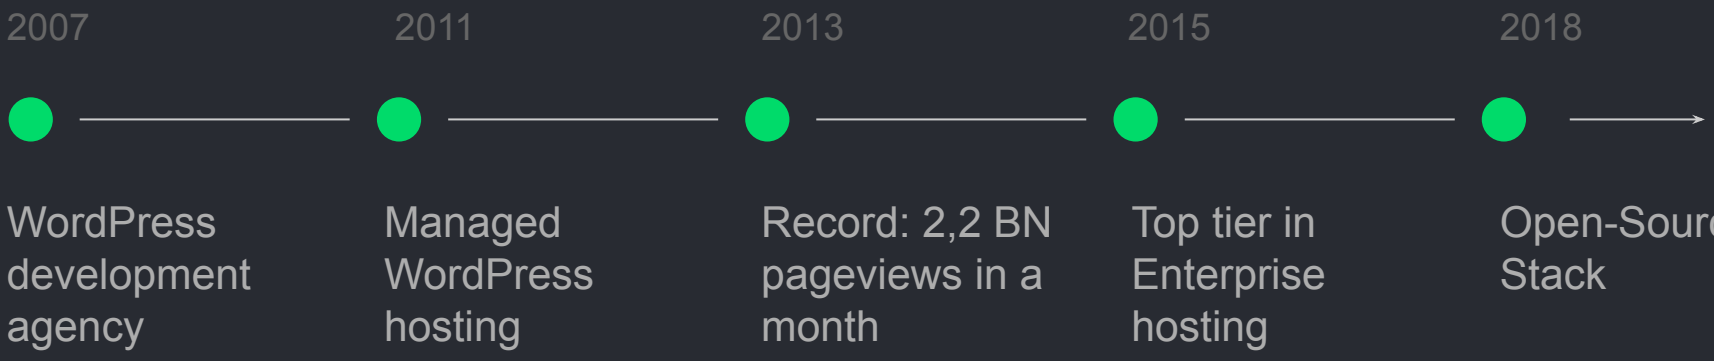

01.2 Presslabs mission

# Democratizing WordPress hosting infrastructure



# 01.3 Presslabs objectives

Open infrastructure using Kubernetes to run and operate WordPress

MySQL Operator for WordPress hosting

# 01.4 Why Kubernetes?

- Runs everywhere
- Open-source
- We had experience with containers before they were cool
- Our core services already run on Kubernetes since version 1.7
- Support for a lot of integrations

# 02. The Need



## 02. The Need

1 — Ease of operations

2 — Elasticity

3 — Service availability

4 — Data safety

5 — Observable



# 03. The Solution



## 03.1 MySQL Operator

A Kubernetes Operator for managing MySQL Clusters with **asynchronous** or **semi-synchronous** replication:

- ✓ Self-healing clusters
- ✓ Highly available reads
- ✓ Virtually highly available writes
- ✓ Replication lag detection and mitigation
- ✓ Resource abuse control
- ✓ Automated backups and restores



03.2

# Architecture overview

- **Control plane**
  - Operator
  - Orchestrator
- **Data plane**
  - MySQL deployment
- **Monitoring**
  - Prometheus

# 03.3 MySQL Cluster

- **Statefulsets**
- **Persistent volumes**
- **Services**
  - Master
  - Healthy nodes

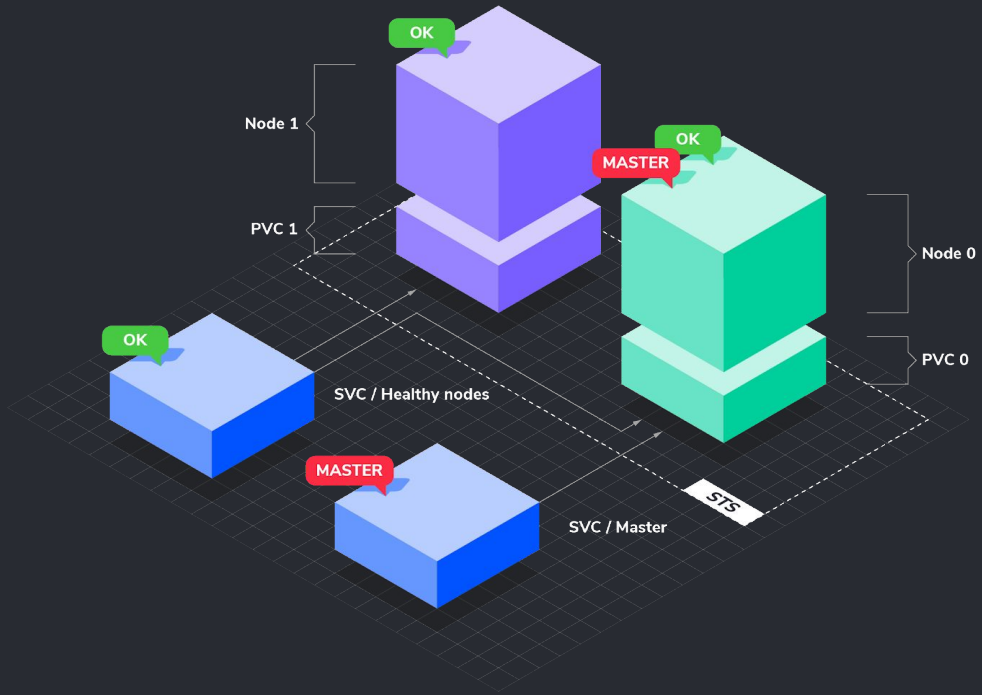

03.4

# Data plane

- **Init**:
  - MySQL configuration
- **Main**:
  - Percona Server for MySQL
- **Sidecar**:
  - Monitoring
  - Automated backups and restores
  - Resource abuse control

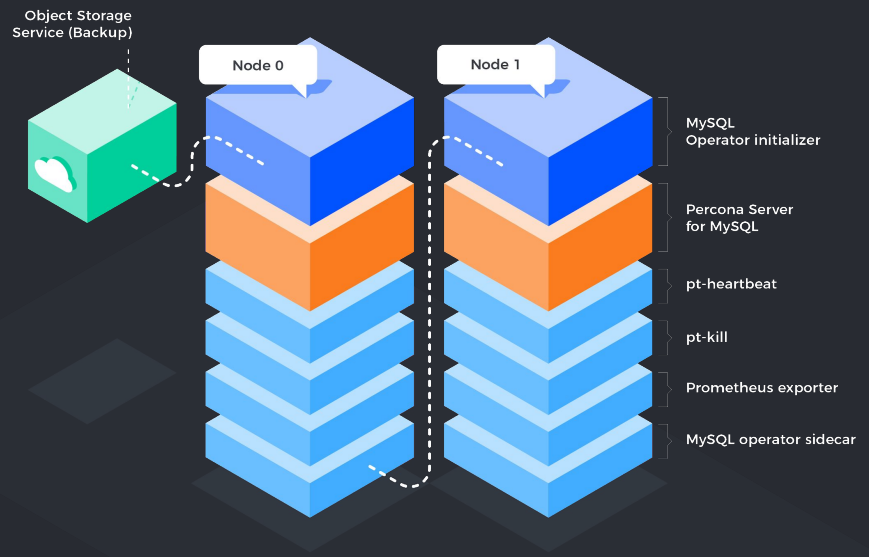

## 04. Challenges

1

Orchestrator sync and clean-up

2

PVC clean-up

3

CRD upgrade/ validation

4

MySQL Upgrade



# 04.1 Orchestrator sync and clean-up

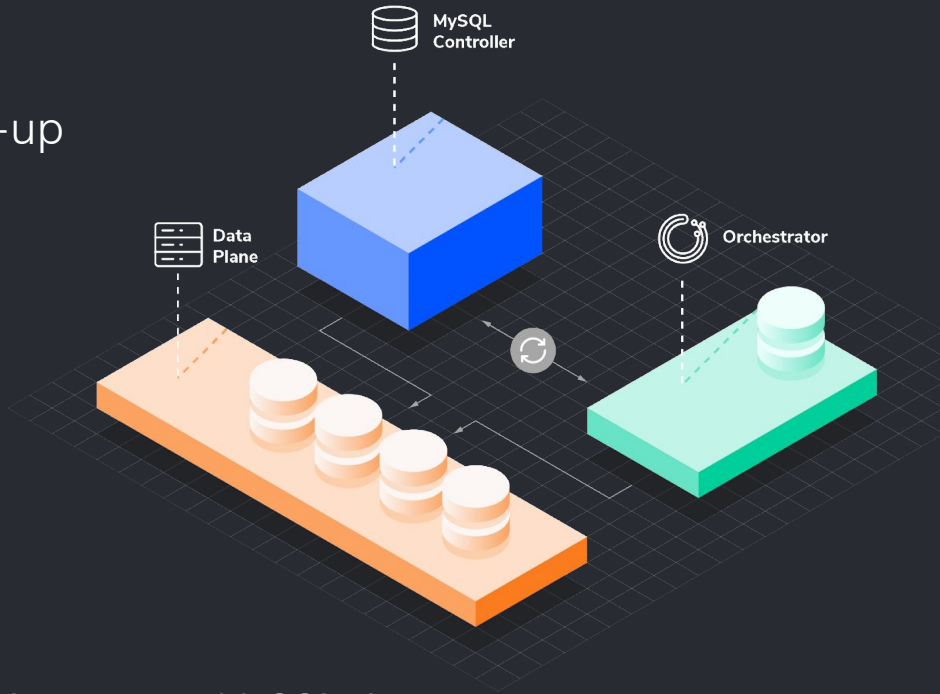


- Orchestrator is a MySQL high availability and replication management tool
- State reconciliation between Orchestrator and Kubernetes
- Information flow:

  Kubernetes → Operator → Orchestrator → MySQL cluster

# 04.2 PVC cleanup

- MySQL nodes keeps data into Persistent Volumes
- Scale down does not delete PVCs
- Scale up may be an issue because of obsolete data
- Special case for Node 0

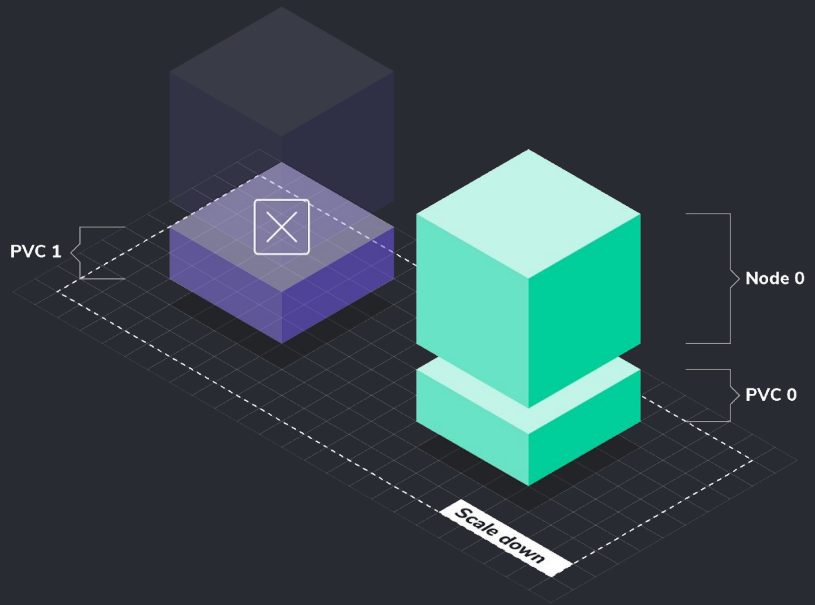

# 04.3 CRD upgrade / validation

- ➔ CRD versioning ( alpha from 1.13)
- ➔ Helm hook

```
helm.sh/crd-install
```

- ➔ No CRD validation



# 04.4 MySQL Upgrade

- Default policy:

```
Rolling Updates
```

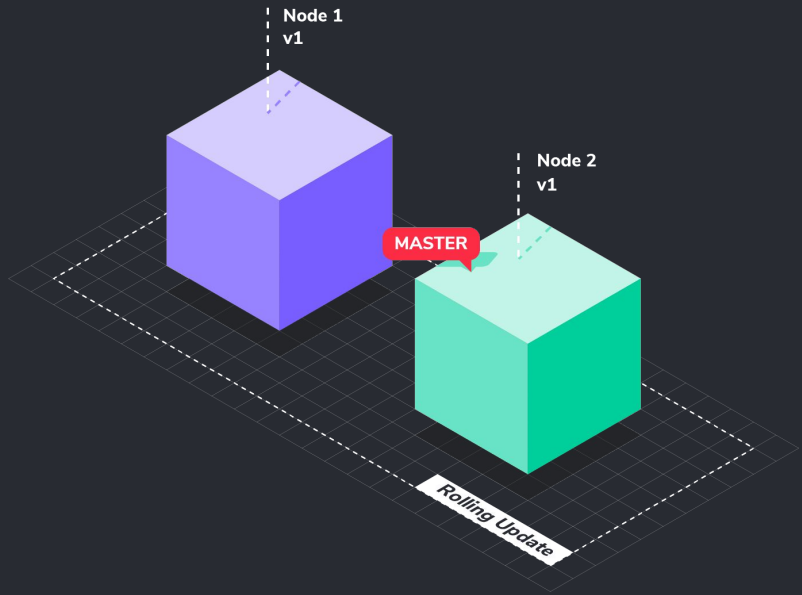

# 04.4 MySQL Upgrade

- Default policy:

```
Rolling Updates
```

- Not gentle for MySQL

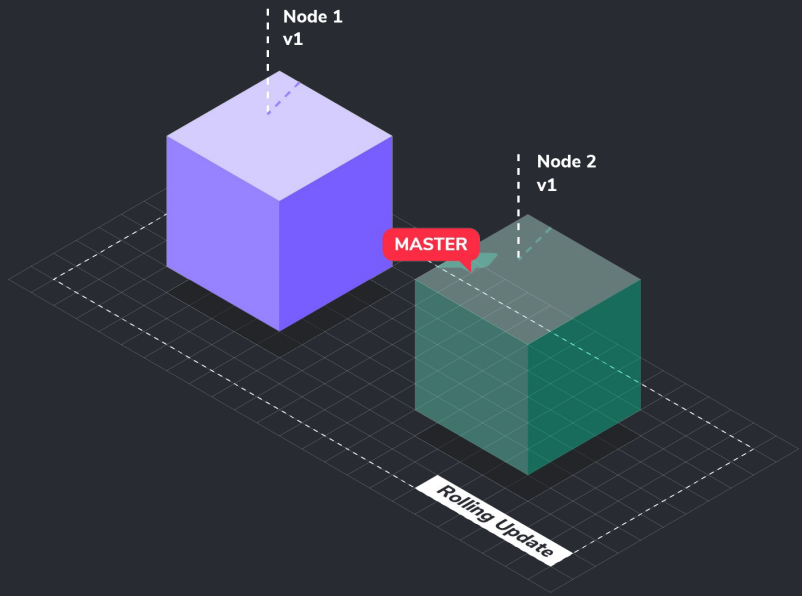

# 04.4 MySQL Upgrade

- Default policy:

```
Rolling Updates
```

- Not gentle for MySQL
- Master should be the last one standing

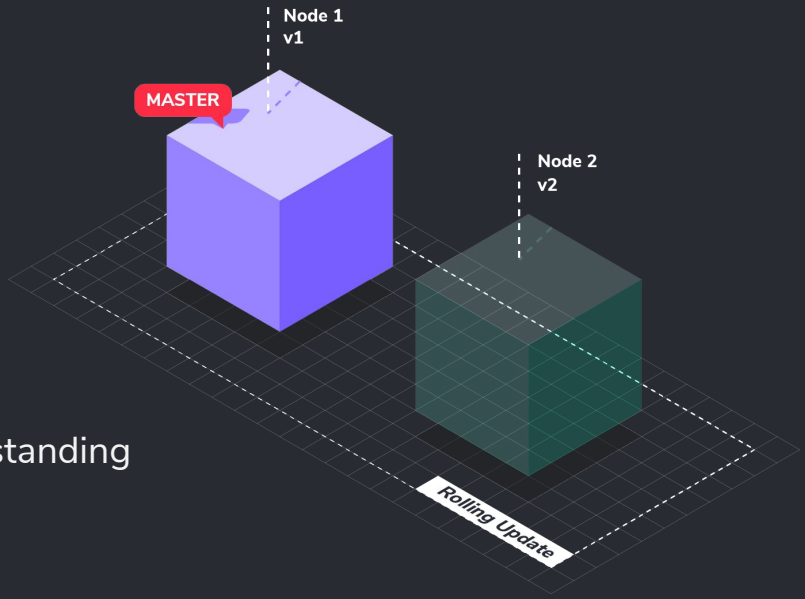

# 04.4 MySQL Upgrade

- ➔ Default policy:

  `Rolling Updates`

- ➔ Not gentle for MySQL
- ➔ Master should be the last one standing
- ➔ Recommended policy: `On Delete`
- ✘ Pod finalizer
- ✘ Container lifecycle hooks

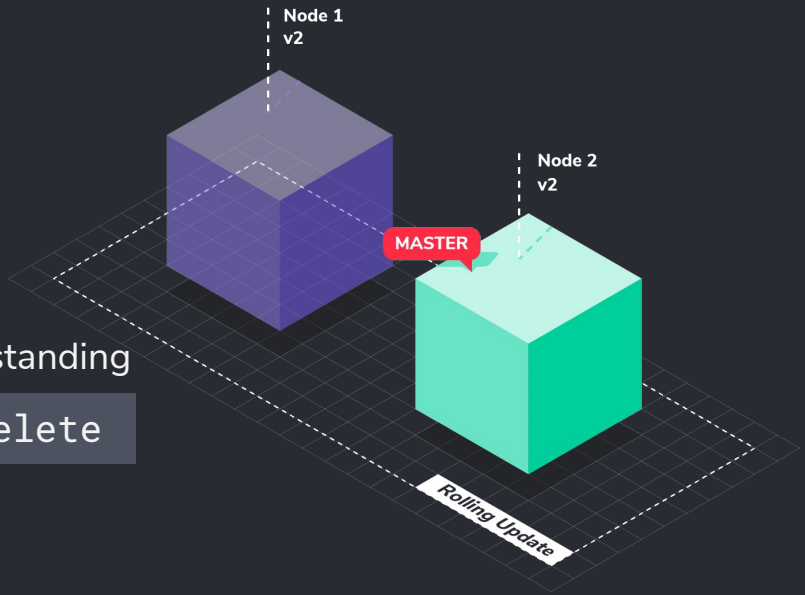

# 05. Current status



# 05.1 Future Plans

- ➔ Integration with Google Cloud Marketplace, OperatorHub.io, AWS Marketplace
- ➔ CRD Validation and webhooks
- ➔ Proxy SQL integration
- ➔ Backup policies



05.2

# Project status

github.com/presslabs/mysql-operator

#mysql-operator

https://kubernetes.slack.com/

## Community

★ Star 158

Fork 41

## Contributors

Presslabs
Smart Managed WordPress Hosting
Romania https://www.presslabs.c...

Heureka
Czech republic http://www.heureka.cz

Platform9 Systems Inc
http://www.platform9.com

Kinvolk
We build technologies for modern cloud-native infrastructure
Berlin, Germany https://kinvolk.io hello@kinvolk.io

AGRI TERRA

PLATFORM9

# Who we are

| | |
|---|---|
| **Mission** | Accelerate enterprise hybrid clouds with a hybrid cloud platform that just works |
| **Key Metrics** | **300 cloud regions** managed globally<br>**500,000 cores** of compute under management |
| **Investors** | Menlo, Canvas Ventures, Redpoint Ventures, Hewlett Packard Enterprise |
| **Key Customers** | Autodesk, Snapfish, Cadence, Aruba, EBSCO, Rodan+Fields, Providence Health Care, Veritas |



PLATFORM9

# Problem of Scale

- Customer base growth == Substantial increase in public cloud costs
- At ~300 cloud regions, just RDS bill amounted to 10s of thousands



PLATFORM9

# Infra evolution

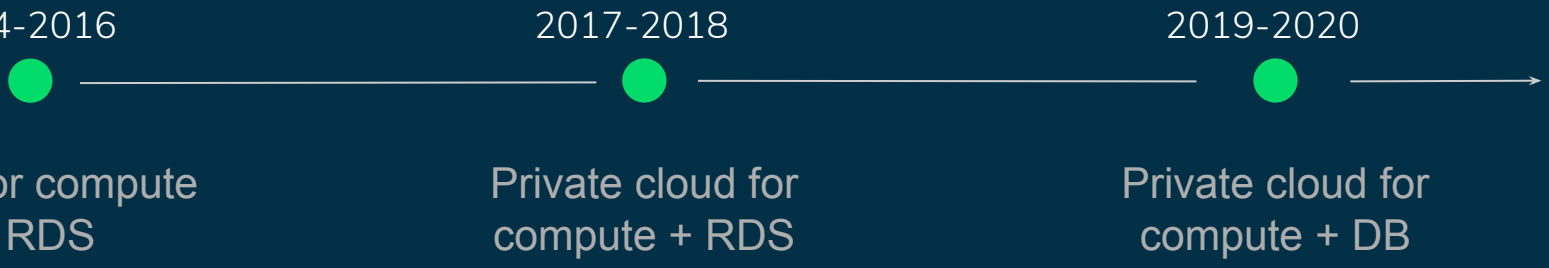




PLATFORM9

# Hurdles to DBaaS

- Automation set up to use of public self service API
- Reliance on RDS snapshots, performance charts and alerting
- No MySQL expertise
- Required Comparable performance
- Needed drop-in replacement for MySQL to minimize impact



PLATFORM9

# Searching for DBaaS

| Requirement | Mysql-Operator Featureset |
| --- | --- |
| Simple, self service, open API | **K8s CRD implementation** |
| Drop in replacement: MySQL | **Percona: 100% compatible** |
| Automated backups, API driven recovery | **Scheduled Backups to S3** |
| High Availability & Failover | **Replica support with automated failover with Orchestrator** |
| Open Source | **Yes** |
| Built-in monitoring | **Yes: Prometheus metrics** |



PLATFORM9

# On-prem Architecture

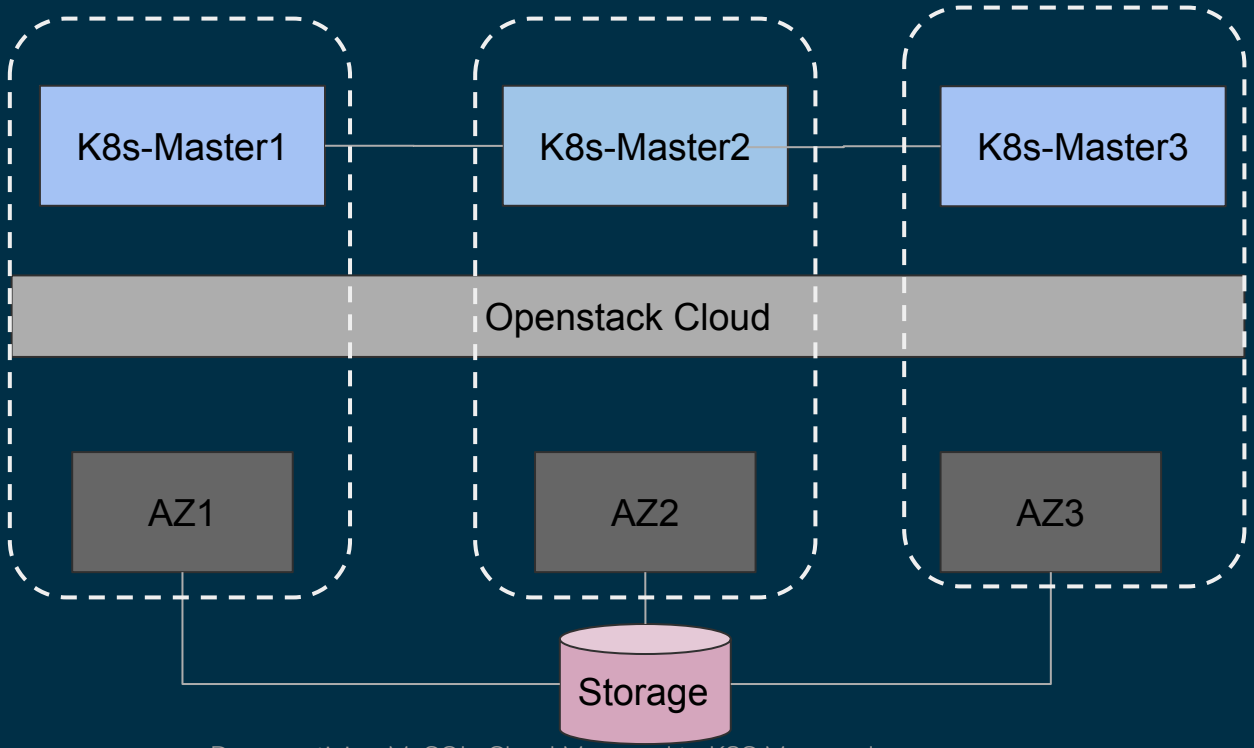

PLATFORM9

# The Rollout

**Current State**

- ~10 accounts using MySQL managed on K8s by the operator
- 3 managed multi-master K8s clusters: Dev, Stage and Prod
- Automated failover with 3-AZ deployment

**Plan**:

- 100% deployments managed with MySQL operator
- Standardize on Operator Paradigm: Prometheus Monitoring, Log collection, etc.



# Thank You!