

# The Cloud is Inevitable

---

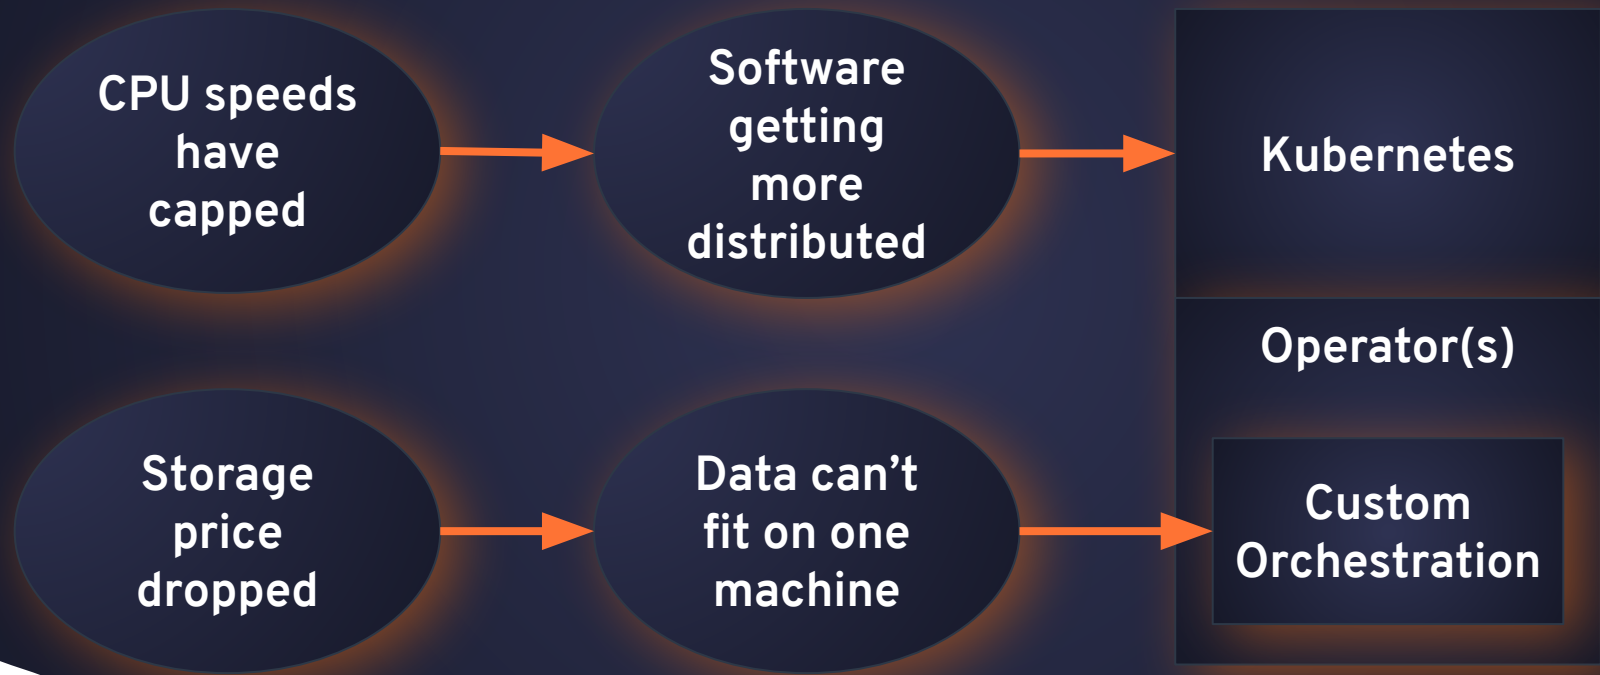
**Sugu Sougoumarane**

Co-Creator, Vitess 

CTO  planetscale

 ssougou

# Why are things **Changing**?



# Problem **Areas**


**Orchestration**

**Distribution**

**Transformation**


# Orchestration

# Scenario 1: Local Disk



- Data loss on Pod reschedule


# Scenario 2: Mounted Volume



- Unavailability on Pod reschedule
- Performance impact

# Scenario 3: Failover


Pod 0: Master



app server


- Cannot designate master
- Failover on Pod reschedule
- Who informs the app?

Pod 1: Replica



# Scenario 4: Discovery

Pod 0: Master




app server




- How do you initialize the replica?
- Who informs the app?

Pod 1: Replica



Pod 2: New Replica





# Vitess Features

---

Discovery

---

Failovers

---

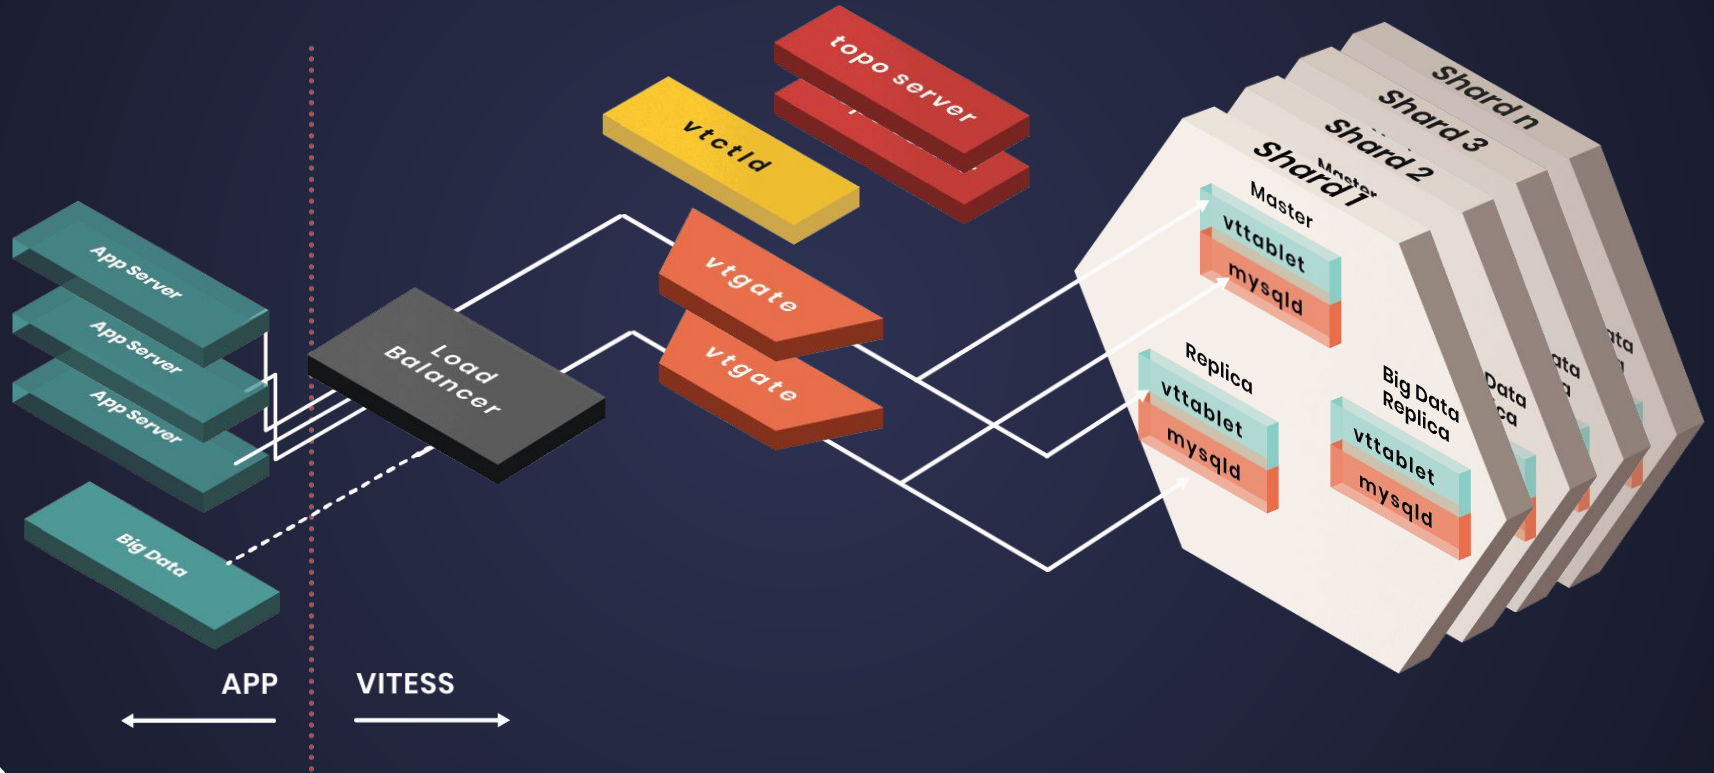
Backups

---

Schema

# Distribution

# Vitess Architecture



# Transformation

## VReplication

# Use Cases

Materialized Views

Migrate Data

Upgrade MySQL Versions

Reshard

Stream Everything out

Change Sharding Key

Anonymize PII

# Slack Migration

**1000+**  
Shards

**100+**  
TB

**100+**  
KQPS

# Estimate

4 Engineers  
3-5 Years

# VReplication

1 Engineer  
1 Year

Zero Downtime  
Transparent to App

# Yelp PII Filtering

**800+**

Tables

**120**

With PII

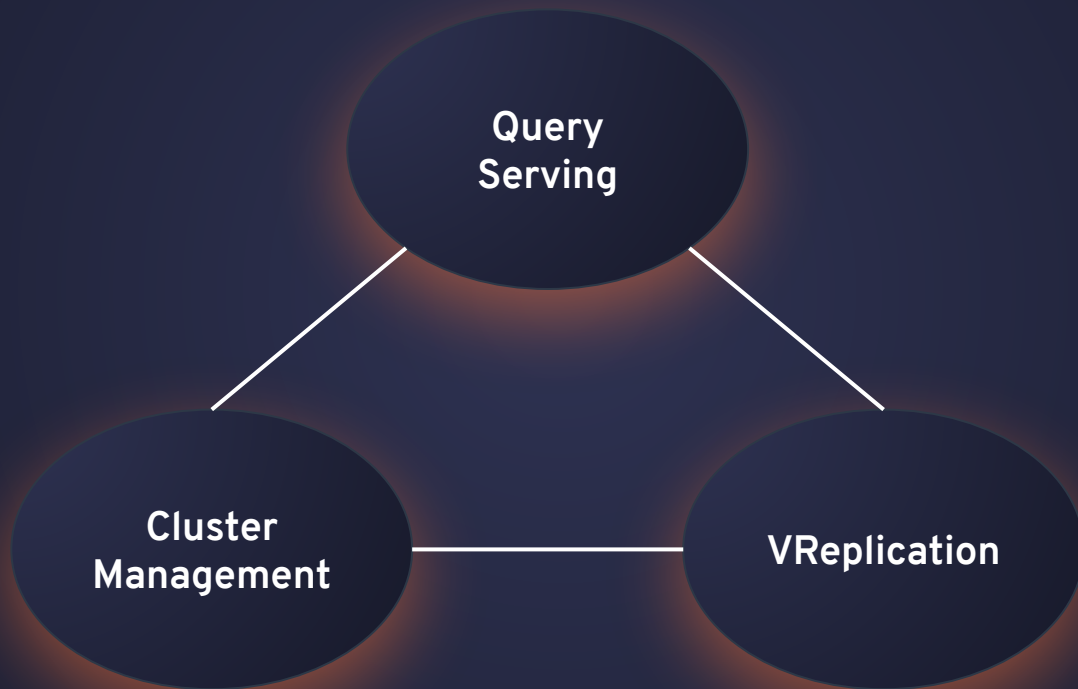
**700+**

Rules



*“Refactor Data like you Refactor Code”*

# Vitess Pillars



# Vitess serves **millions of QPS** in production

 YouTube

 New Relic®

 slack

*Flipkart*

 HubSpot

*peak*

 Pinterest

 Square

 BetterCloud

 weave

**GitHub**

**JD.** 京东  
.COM

 nozzle

**Quiz of Kings**

 stitchlabs

# Q&A

Website [vitess.io](https://vitess.io)

Shlomi's Talk [sched.co/ePp6](https://sched.co/ePp6)

Slack: [vitess.slack.com](https://vitess.slack.com)

