Building Cloud Native GraphQL Applications on Kubernetes with YugabyteDB and Hasura

Nikhil Chandrappa Ecosystem Engineering

yugabyteDB

Gavin Johnson

Senior Product Marketing Manager,
Yugabyte Cloud, Yugabyte

Product Marketer

RudderStack
New Relic
AT&T

Tech Consultant

Deloitte (MBA Intern -> Senior Consultant -> Manager)

Sys Admin

Developer

Nikhil Chandrappa

Lead Software Engineer,
Ecosystem Engineering, Yugabyte

Data Engineer
Pivotal (Acq. VMware)

Systems Engineer
Infosys

Kubernetes is massively popular in Fortune 500s

- Walmart - Edge Computing
  KubeCon 2019 https://www.youtube.com/watch?v=sfPFrvDvdlk

- Target - Data @ Edge
  https://tech.target.com/2018/08/08/running-cassandra-in-kubernetes-across-1800-stores.html

- eBay - Platform Modernization
  https://www.ebayinc.com/stories/news/ebay-builds-own-servers-intends-to-open-source/

Why stateful services on K8s?

Containerized data workloads running on Kubernetes offer several advantages over traditional VM / bare metal based data workloads including but not limited to:

- Better cluster resource utilization 
- Portability between cloud and on-premises
- Frictionless multi-tenancy with versioning
- Simple and selective instant upgrades
- Robust automation framework can be embedded inside **CRDs** (Custom Resource Definitions) or commonly referred as ‘K8S Operator’. 

Getting started with GraphQL

- Provides GraphQL abstractions over your Database
  a. Query, mutate data via GraphQL constructs
  b. Build your schema and evolve your domain models
- Evolve your API without versioning
- Provides out of the box pagination and filtering
- Collate disparate Data sources
- Eventing Support

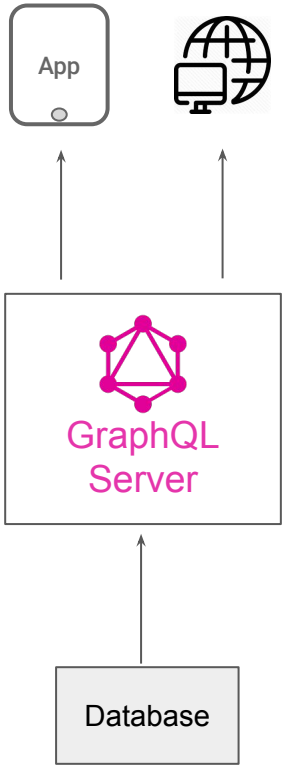

Cloud native relational database for cloud native applications

yugabyteDB

Distributed SQL database for transactional applications.
100% open source. Runs on any cloud.

PostgreSQL Compatibility

Resilience and High Availability

Horizontal Scalability

Geographic Distribution

ACID Transactions

Security

YugabyteDB and Hasura on K8S

Design decisions for Cloud Native Databases

Cloud Native

- Scale on-demand
- API driven admin operations
- Prometheus based metrics

Resiliency

- Against cloud infrastructure failures
- Automated and predictable data recovery

Day 2 Operations

- Zero downtime upgrades
- Automated backup/recovery

YugabyteDB deployed as StatefulSets

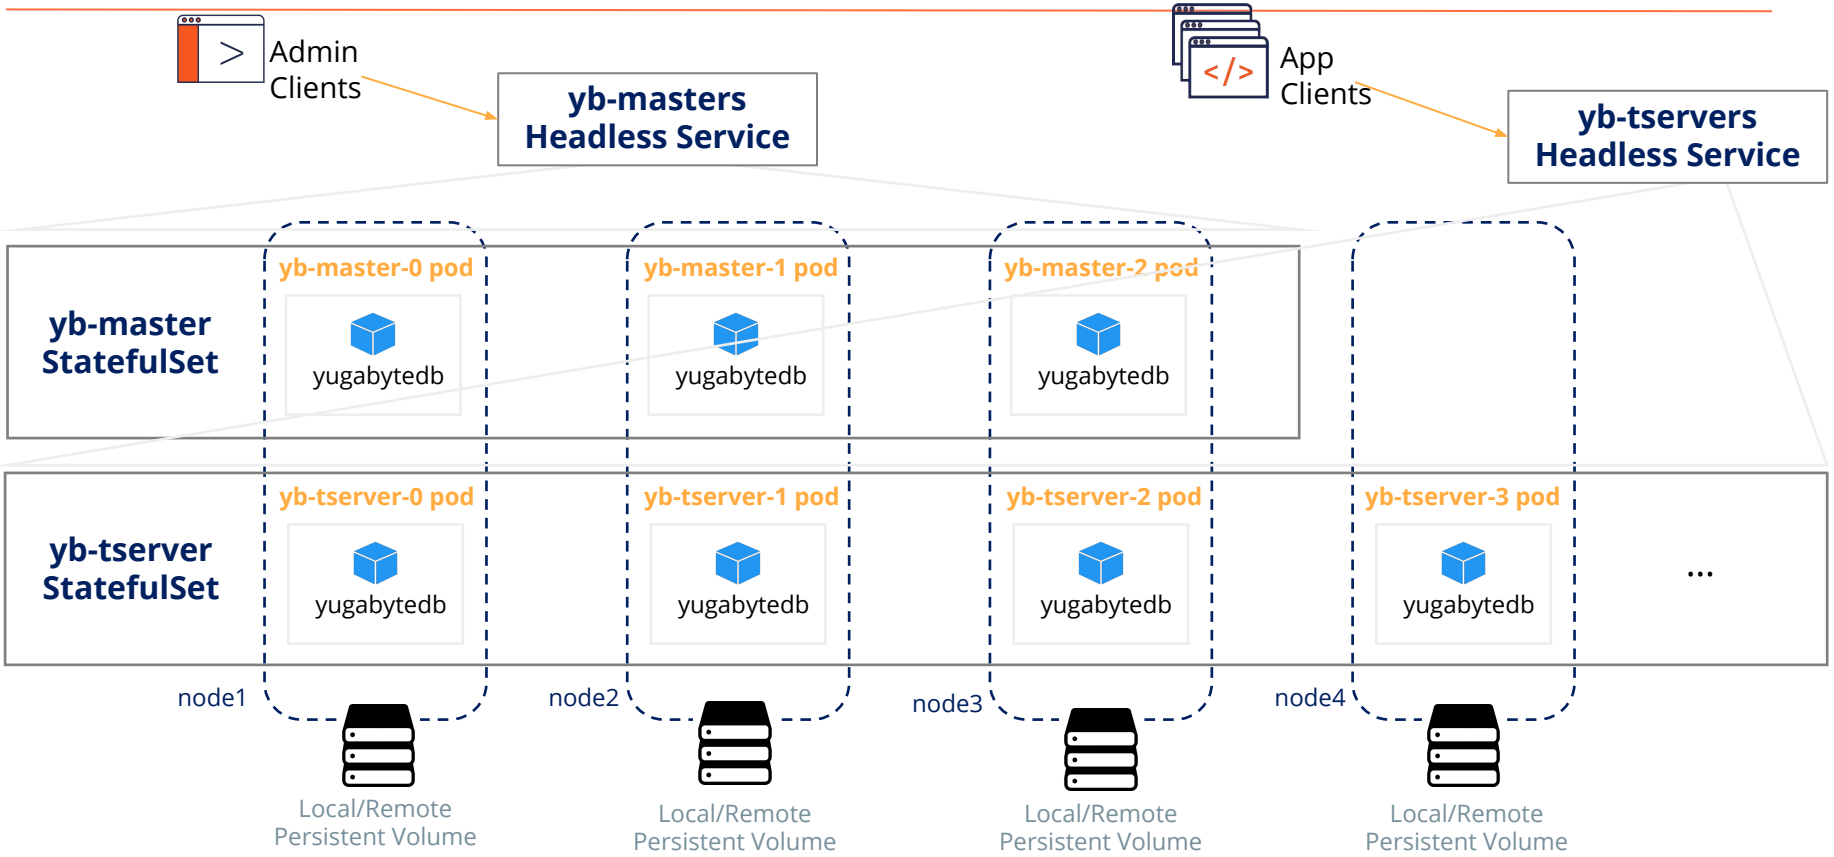

Automating day 2 operations

HANDLING FAILURES

Pod failure handled by K8S automatically

Node failure has to be handled manually by adding a new slave node to K8S cluster

Local storage failure has to be handled manually by mounting new local volume to K8S

ROLLING UPGRADES

Supports two *upgradeStrategies*: *onDelete* (default) and *rollingUpgrade*

Pick rolling upgrade strategy for DBs that support zero downtime upgrades such as YugaByteDB

New instance of the pod spawned with same network id and storage

BACKUP & RESTORE

Backups and restores are a database level construct

YugaByteDB can perform distributed snapshot and copy to a target for a backup

Restore the backup into an existing cluster or a new cluster with a different number of TServers

Hasura connectivity with YugabyteDB on Kubernetes

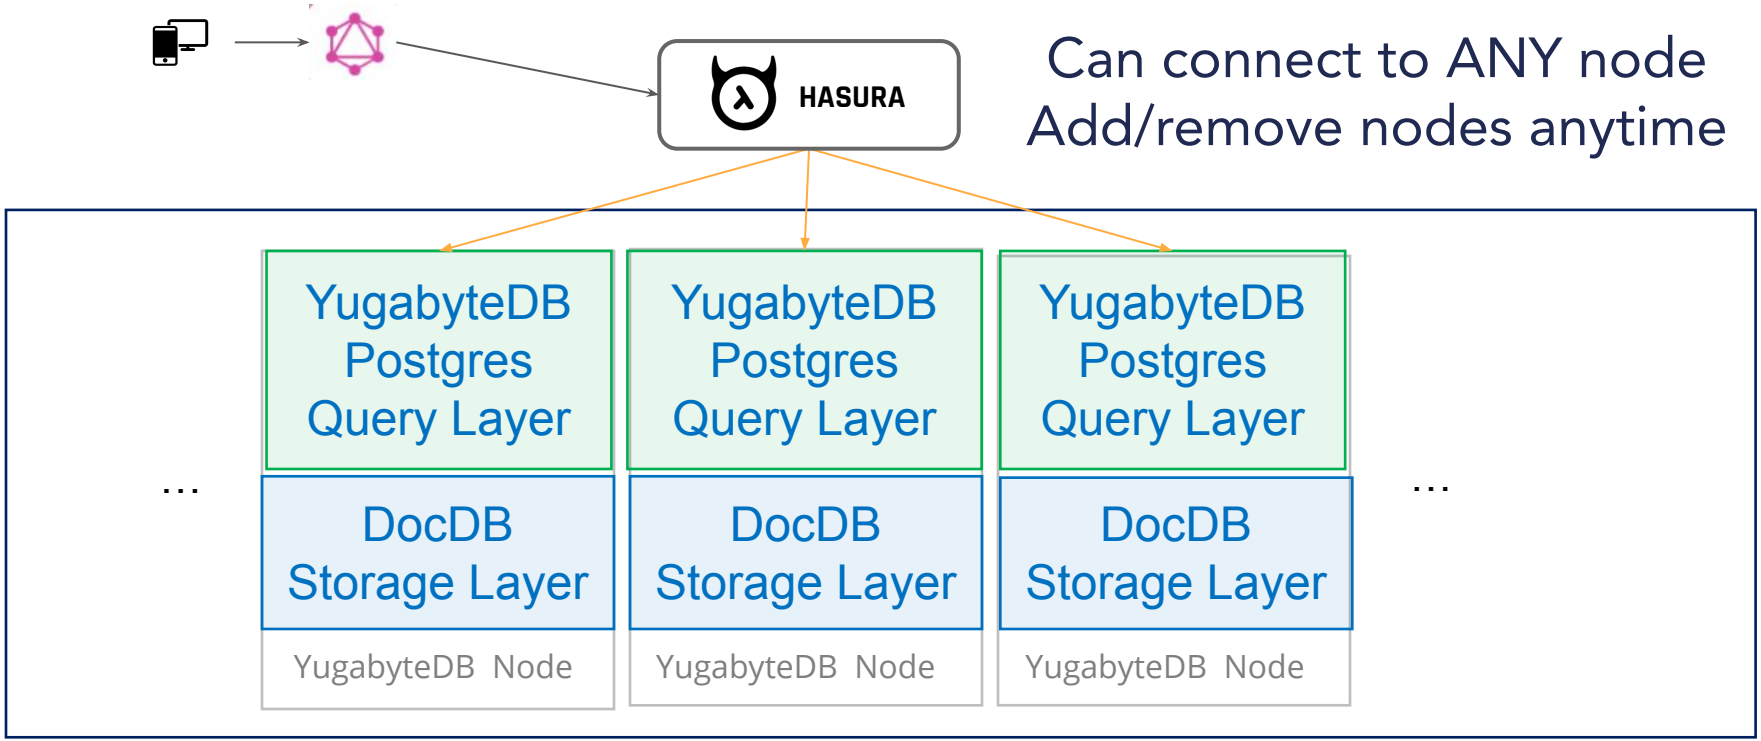

Building Real Time Poll App Using Hasura and Yugabyte on Kubernetes

Scaling out GraphQL services

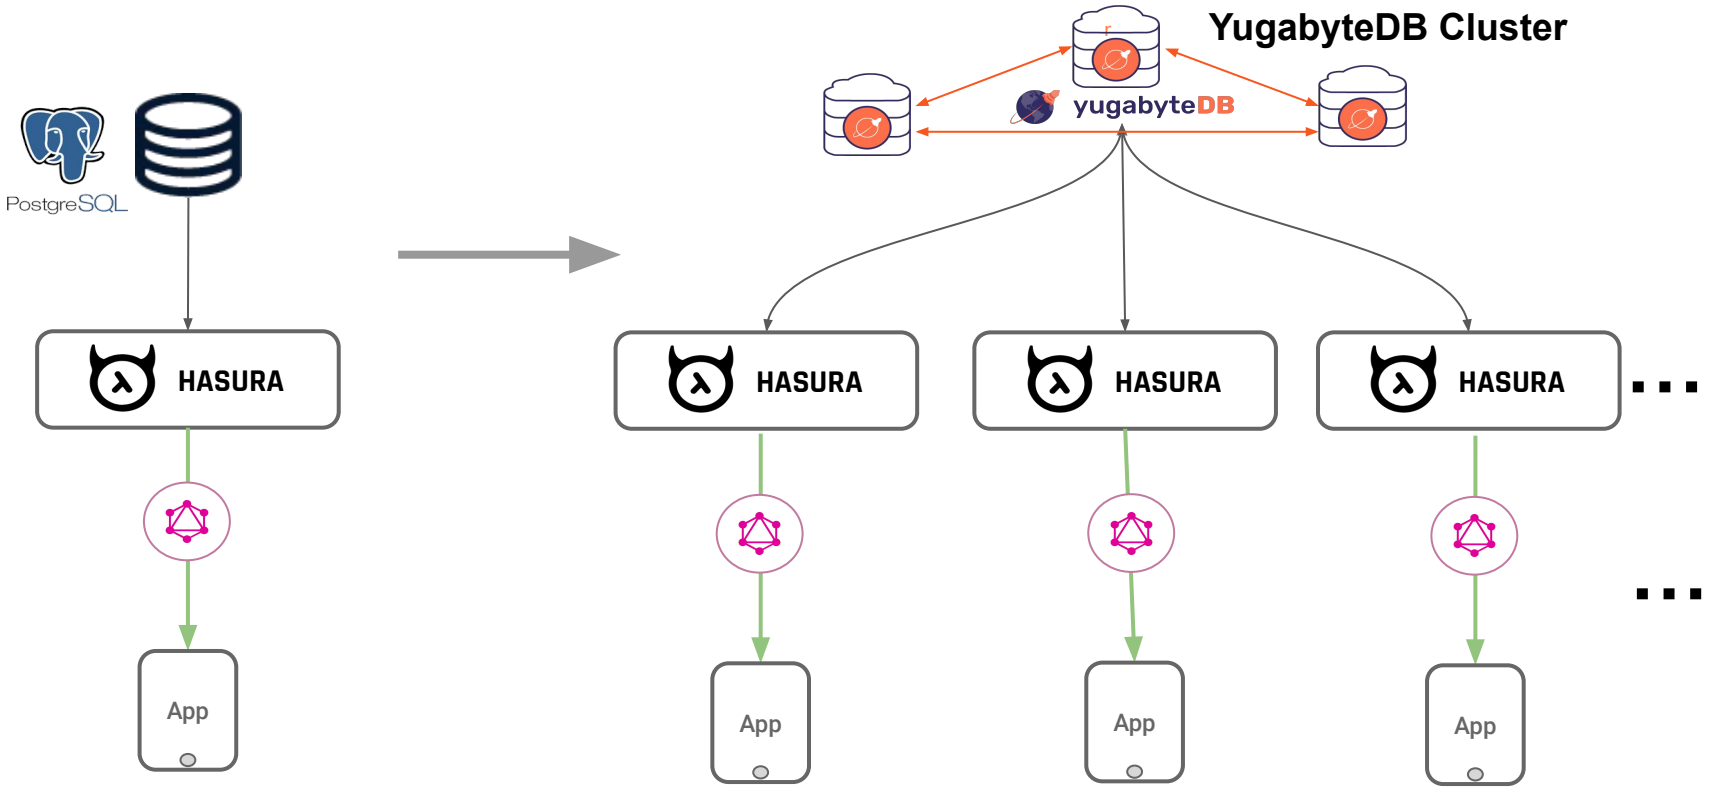

Try YugabyteDB

Distributed SQL database built for cloud native applications

yugabyteDB

Distributed SQL database for transactional applications.
100% open source. Runs on any cloud.

PostgreSQL Compatibility	Resilience and High Availability
Horizontal Scalability	Geographic Distribution
ACID Transactions	Security

Admiral HRT Kroger CIPHERTRACE narvar Plume TURVO VOICELAND xignite

Postgres reimagined for a cloud native world

	PostgreSQL	Google Spanner	YugabyteDB
SQL Ecosystem	✓ Massively adopted	✗ New SQL flavor	✓ **Reuse PostgreSQL**
RDBMS Features	✓ Advanced *Complex*	✗ Basic *cloud-native*	✓ **Advanced** *Complex and cloud-native*
Highly Available	✗	✓	✓
Horizontal Scale	✗	✓	✓
Distributed Txns	✗	✓	✓
Data Replication	Async	Sync	Sync + Async

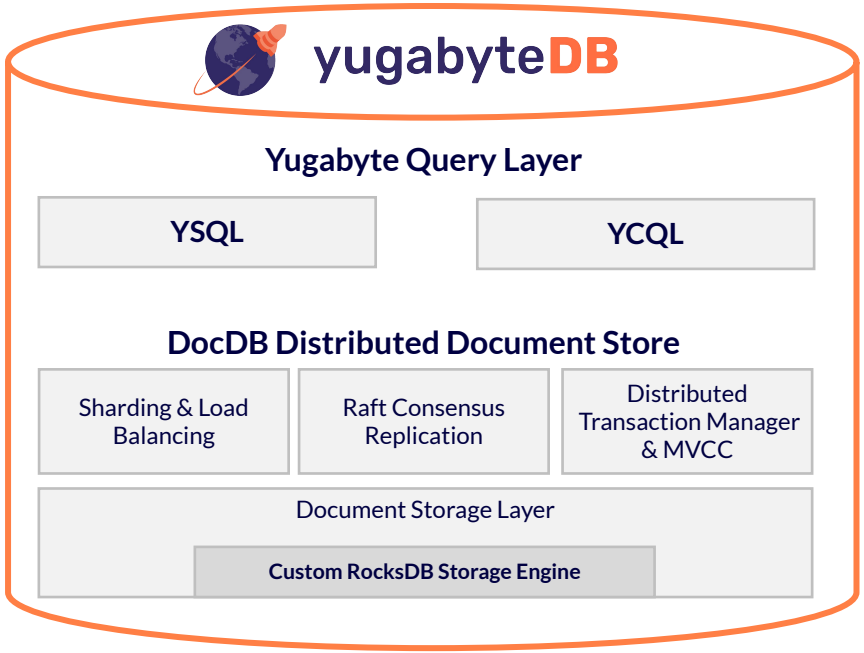

Target use cases

Systems of Record and Engagement

Resilient, business critical data

- Identity management
- User/account profile
- eCommerce apps - checkout, shopping cart
- Real time payment systems

Event Data and IoT

Handling massive scale

- Vehicle telemetry
- Stock bids and asks
- Shipment information
- Credit card transactions

Geo-Distributed Workloads

Sync, async, and geo replication

- Vehicle telemetry
- Stock bids and asks
- Shipment information
- Credit card transactions

Yugabyte offers flexible consumption models

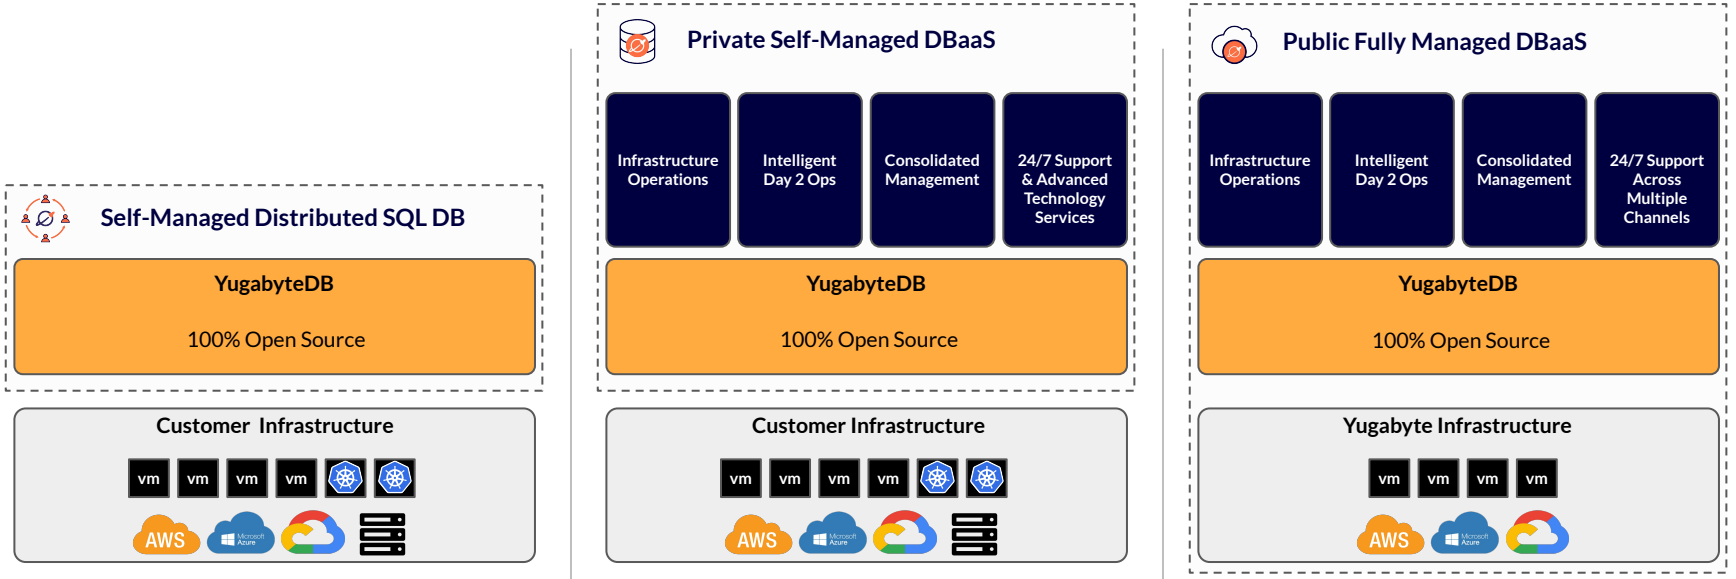

YugabyteDB

- Self managed
- Manual infrastructure operations
- Manual database day 2 operations
- Community support

Yugabyte Platform

- Self managed (or with Yugabyte Services)
- Automated infrastructure operations
- Automated database day 2 operations
- 24/7 support available
- Advanced technology services

Yugabyte Cloud

- Fully managed by Yugabyte
- Consolidated management console
- Automated infrastructure operations
- Automated database day 2 operations
- 24/7 support

Thank You

Join us on Slack:
www.yugabyte.com/slack

Star us on GitHub:
github.com/yugabyte/yugabyte-db

yugabyteDB