

**“Kubernetes  
*is not* for  
databases”**

That phase  
is **over.**



# About me



- EMEA Sales Engineer Lead at EDB
- Based in Amsterdam
- Database since 1990
- Oracle Database version 4
- Former Oracle (23 years), DB, Middleware, Exadata, OCI



kevinckli

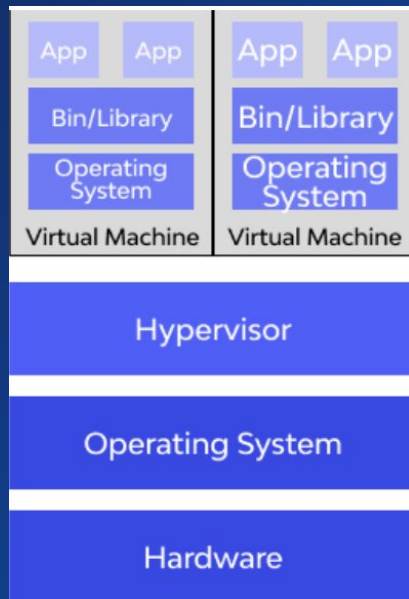


@kevinckli

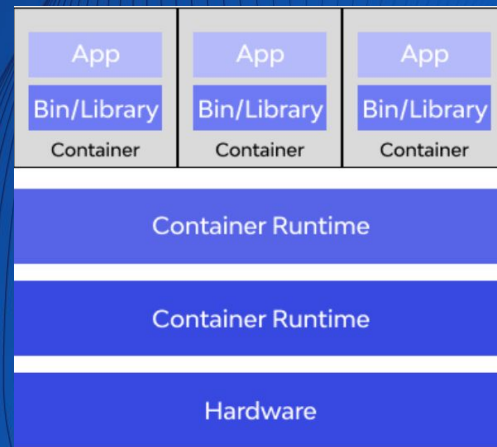
# What is Kubernetes?



Bare Metal



Virtualization



Kubernetes (Cloud Native)

# Two points of clarification

## Point 1: Cloud Native $\neq$ Cloud

**Cloud Native does not mean it is deployed on the public cloud. It can be, but it doesn't have to be.**

Kubernetes, containers, and DevOps are optimized for cloud-native architectures using microservices.

But they can be used across on-premises, private cloud, and public cloud environments.

## Point 2: Containers $\neq$ Kubernetes

**Kubernetes (k8s) and containers are not interchangeable terms. They are not the same thing.**


A container packages up code and all its dependencies so the application runs quickly and reliably from one computing environment to another.

Containers are managed and orchestrated by Kubernetes, from a single control plane.

# An overview of Automated Postgres on Kubernetes

# PostgreSQL: World's most popular Open Source database

- Aka Postgres
- Open source Community
- 28+ years of innovation
- “Rock solid”
- **One primary, multiple replicas**
- **More customers:** Than any dedicated PostgreSQL company
- **More experts:** Leading PostgreSQL contributors
- **More innovation:** Positioned to lead in enterprise PostgreSQL and hybrid cloud



# Introducing Kubernetes

Kubernetes is an open-source platform designed to automate deploying, scaling and managing containerized applications

Originally built for stateless applications, Kubernetes can be taught to manage stateful applications as well

Operators" are extensions to Kubernetes that can teach Kubernetes "new tricks"

A Postgres operator could help DBAs and developers automate day 1 and day 2 work and focus on value drivers







**Dave Pitts** (He/Him) **Author**

17h (edited) ...

Database Engineer, Adyen

This was a **\*\*really fun event\*\*** and there were lots of good followup questions/discussions on a broad variety of database topics:

- \* Getting started with Python and SQLite
- \* Running databases on k8s - a very hot topic in the DB Engineering world
- \* Newer NoSQL vs traditional Relational DB engines - we went into the pro's and con's of both approaches.
- \* Building Prometheus dashboards: I have only ever built a single dashboard (but I do love it some very interesting maths). Prometheus is another very exciting open-source project.
- \* We also has some discussions about some banks still running on mainframes. I enjoyed working with COBOL, DB2, JCL and even some Assembler in the 1990s, but by 1999 I was focusing on Java, WebApps, Unix and eventually decided that databases were going to be "my thing".

Looking forward to the [#outintech](#) [#madrid](#) Xmas event (Dec 5th) 🍷🎅



You can  
same a  
unders  
Cloud

[Traduci il](#)



[Sc](#)  
[@kels](#)  
[a Kub](#)  
[start l](#)

[Mostr](#)

5:21 PM

# Why Kubernetes?

Kubernetes already includes many system services needed for managing software

Load Balancing,  
and SDN



Health  
checks



Storage  
management



Automated  
Scheduling



Scalability:  
scale-up/down



Rolling  
Deployments



# Why run Postgres on Kubernetes?

## (Business) benefits of using Postgres on Kubernetes



### Go to market sooner

Self-service database creation and maintenance, instant delivery of new clusters.

Keep development fast!



### Keep staff focused on business objectives

Restoring a backup? Configuring some GUC parameter? Updating a cluster? Easily done, easily replicated, even with DBA support


Lower the load on DBAs!



### Keep the business running

When a node fails, CNPG will self-heal with little to no service interruption, and zero required effort from the DBA team. Scaling is as easy as pie!

Keep customers happy!



*“How do I synchronize the state of a Postgres database in Kubernetes?”*



## Storage level replication



~~Generic way to run stateful workloads~~

Good for applications without advanced replication mechanisms

## Postgres level replication




Founded on the Write Ahead Log


HA clusters via RO replicas

Sync replication controlled at the transaction level

Shared nothing architectures



*“How do I manage the full lifecycle  
of a Postgres database in  
Kubernetes?”*




CloudNativePG

# Run PostgreSQL. The Kubernetes way.

CloudNativePG is the Kubernetes operator that covers the full lifecycle of a highly available PostgreSQL database cluster with a primary/standby architecture, using native streaming replication.

[View on GitHub](#)



## Autopilot

It automates the steps that a human operator would do to deploy and to

## Data persistence

It doesn't rely on statefulsets and uses its own way to manage persistent

## Designed for Kubernetes

It's entirely declarative, and directly integrates with the Kubernetes API





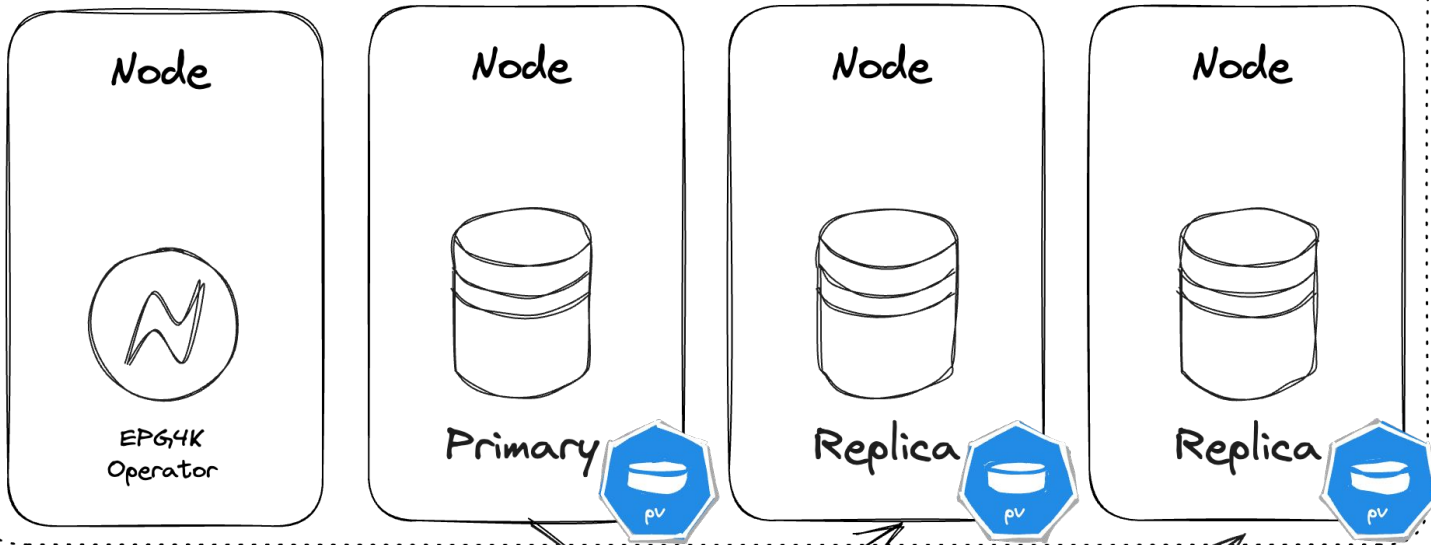
## How does that work?!

- The CNPG operator "codifies" EDB's expert knowledge on how to deploy, manage, and use Postgres databases.
- The operator then gives you an opinionated way of leveraging that codified knowledge by providing it with "cluster declarations" that the operator will act upon.
- Simplest example of that is "there is a Postgres 15.3 cluster, consisting of 3 instances, with 10Gi storage per instance"



```
1  apiVersion: postgresql.cnpg.io/v1
2  kind: Cluster
3  metadata:
4    name: my-cluster
5    namespace: edb-postgres-demo
6  spec:
7    instances: 3
8    primaryUpdateStrategy: unsupervised
9    primaryUpdateMethod: switchover
10   imageName: quay.io/enterprisedb/postgresql:15.3
11   storage:
12     size: 10Gi
```

# Kubernetes cluster



# EDB Postgres for Kubernetes



## EDB Technology

Stack of enterprise grade products designed, developed, and supported by EDB

Exploiting Postgres workloads in Kubernetes



## Kubernetes Experience

EDB Postgres for Kubernetes operator - based on the open source CloudNativePG operator

Adds EDB Postgres Advanced Server workloads with a primary/standby architecture



## Red Hat Certified

Red Hat Certified Kubernetes Operator

Available on Red Hat OpenShift




# Kubernetes



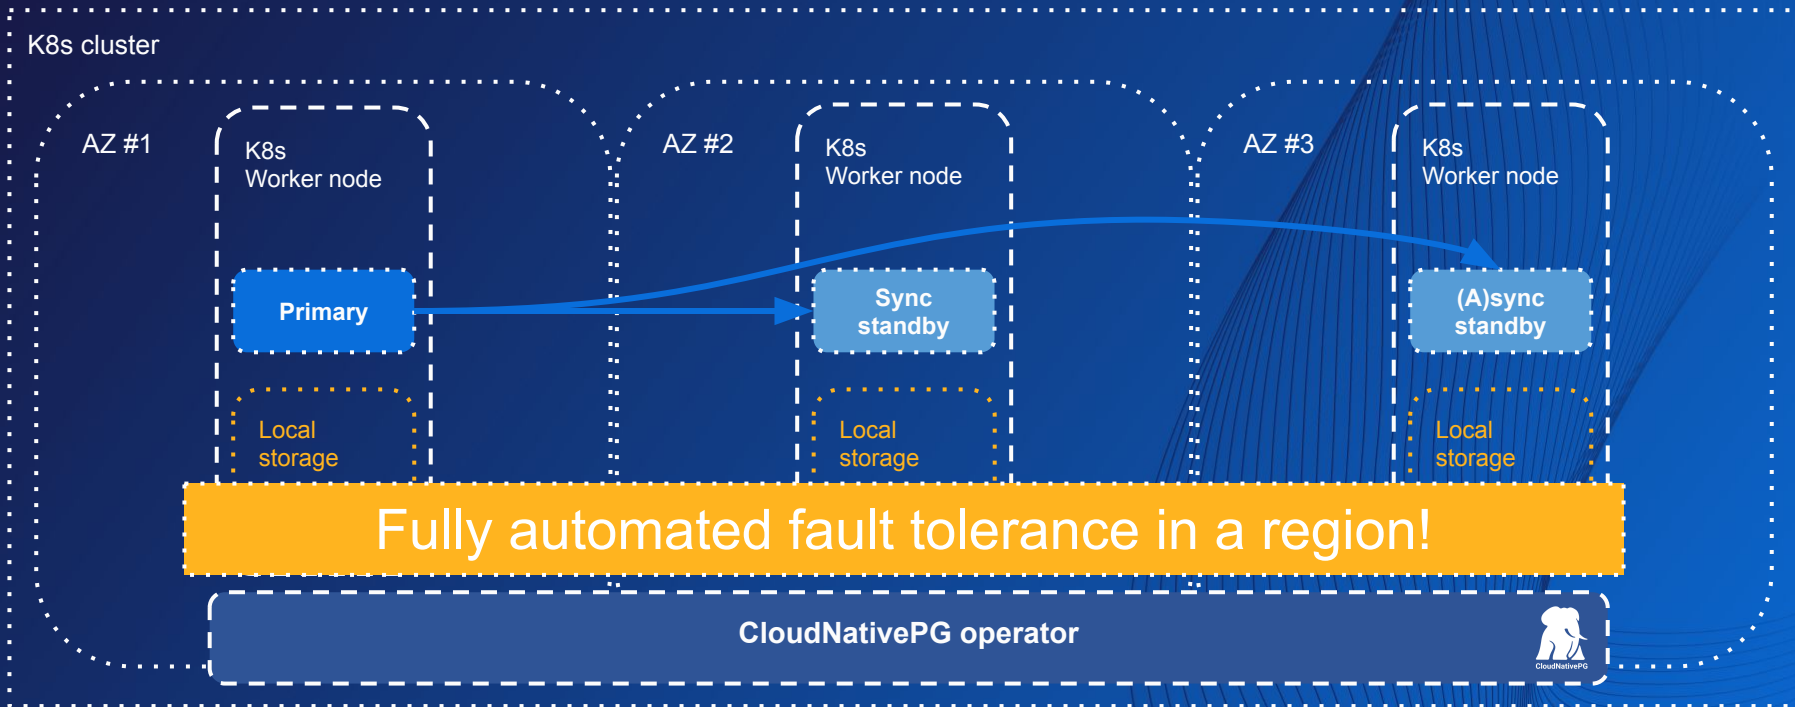
# kubernetes

- **Kubernetes is the means, not the goal**
  - Kubernetes is not the panacea of all problems
  - Find out why Kubernetes is used in your organization
- **Kubernetes is an open source Standard Distributed OS**
  - Private (including on-premise), Public, Hybrid and Multi Cloud
  - Bare metal and VM Kubernetes
  - Self-managed or managed by a third-party (like a CSP)
- **Vibrant inclusive community around the CNCF**
  - KubeCon is the largest open source community conference in the world
- **You can run databases like Postgres in Kubernetes (in production)**
  - You need Kubernetes skills, Postgres skills and a reliable operator
  - Storage and Architecture are critical, like in VMs

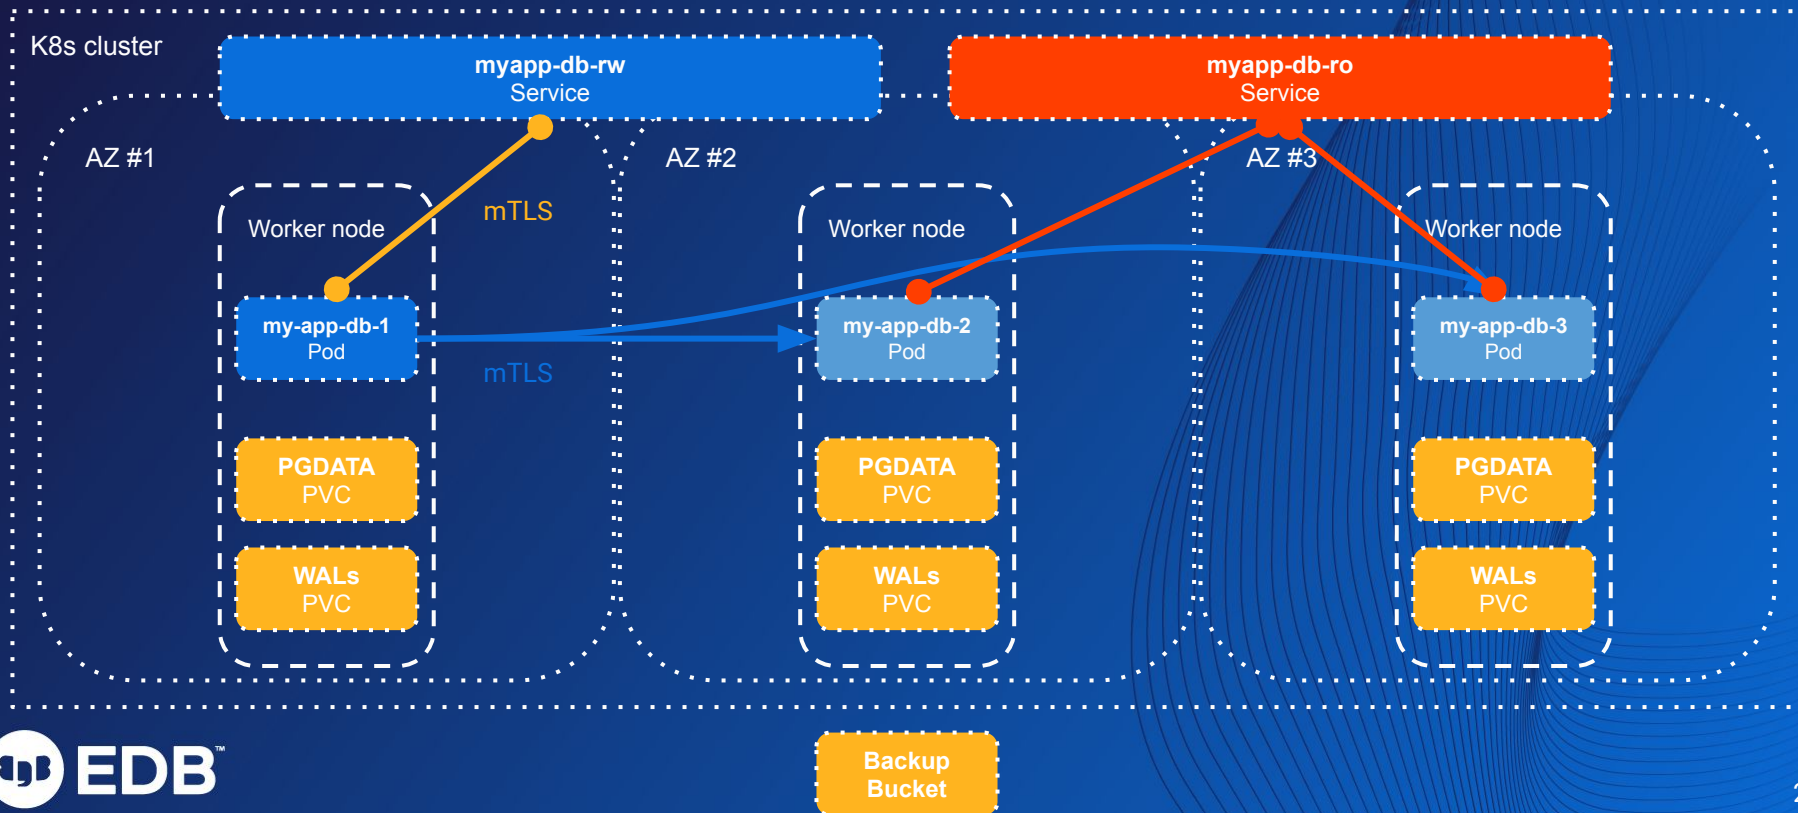
# A common HA scenario in traditional environments



# How a HA Postgres cluster might look in Kubernetes



# This is what happens under the hood





# Example of Kafka monitoring




# Next stop: Very Large Databases (VLDB)

- **Support for Kubernetes 'VolumeSnapshot' API**
  - Transparent incremental & differential copy for backup & recovery
  - PVC cloning
- **In-place major upgrades with pg\_upgrade**
- **Declarative tablespaces**
  - Including temporary ones

# EDB Postgres Distributed for Kubernetes

Always On Multi-Location: Active/Active, Active/Passive, Active/DR



## More from EDB

- **Community 360 Plan for the open source stack**
  - Kubernetes, Postgres and CloudNativePG
- **EDB Postgres for Kubernetes**
  - Certified with Red Hat OpenShift
  - Advanced Security features & Database migration (Oracle Compatibility)
- **EDB Postgres Distributed for Kubernetes**
- **EDB BigAnimal (Managed Service in public cloud - Azure, AWS, Google)**
- **CloudNativePG on GKE and GKE Autopilot Marketplace**

# How CloudNativePG changes the Postgres DBA role

- **Most infrastructure related problems are automated**
- **You as a DBA are crucial in the organization: scale yourself!**
  - Protect Postgres, from day 0
  - Focus on day 2 operations



# Conclusion

- **Key takeaways:**
  - Postgres is the world's most Popular, Admired & Desired\* database by professional developers
  - State synchronization via Postgres native replication
  - Running Postgres (or any database) inside Kubernetes requires an operator
- **Benefits:**
  - No vendor lock-in
  - Private, Public, Hybrid and Multi-Cloud
  - Own your data!
  - Low RTO and RPO business continuity goals within a region, out of the box
  - Add Postgres to your GitOps pipelines for microservice databases

**“Kubernetes  
*is also for*  
databases”**



CloudNativePG has been originally created by



Reach out to us for support on the entire stack!

