

EDB

Postgres for the AI Generation

How to run PostgreSQL in Kubernetes

Sergio Romera,

Senior Manager, EMEA Sales Engineer (South)


26th Sept 2024



Sergio Romera



- **Senior Manager, Sales Engineer @ EDB**
- Based in France
- Database fanatic since 1997
- Developer, DBA, Architect, Sales Engineer
- Companies: BNPParibas, Quest Software, Oracle
- Postgres Certified Professional
- Kubernetes Cloud Native Associate certified
- Oracle Certified Associate
- AWS and Azure certified



EDB
Postgres® AI



EDB POSTGRES AI PLATFORM

UNIFIED WORKLOAD MANAGEMENT

TRANSACTIONAL

ANALYTICAL

ARTIFICIAL INTELLIGENCE

SINGLE PANE OF GLASS ADMINISTRATION

HYBRID DATA ESTATE

INTELLIGENT OBSERVABILITY

ENTERPRISE SECURITY

HYBRID AND MULTI-CLOUD DEPLOYMENT

PUBLIC CLOUD
(MANAGED)

PRIVATE CLOUD
(SOFTWARE)

ON PREMISES
(APPLIANCE)

EXTENSIBILITY

CSP INTEGRATIONS

DEVOPS TOOLING

KUBERNETES TOOLING

GENAI & LLM INTEGRATIONS

LAKEHOUSE INTEGRATIONS

PLATFORM TOOLS AND SERVICES

MIGRATION
PORTAL

CONTINUOUS HIGH
AVAILABILITY

BACKUP AND
RECOVERY

Delivered with world-class
strategic partners:

carahsoft



Red Hat



NUTANIX



Why Kubernetes

One of the benefits of Kubernetes is that it makes building and running complex applications much simpler.

Services, Load Balancing and Networking

Health Checking

Storage Management

Automated Scheduling

Scalability: Scale out/Scale down

Rolling Deployments



Kubernetes timeline

- 2014, June: Google open sources Kubernetes
- 2015, July: Version 1.0 is released
- 2015, July: Google and Linux Foundation start the CNCF
- 2016, November: The operator pattern is introduced in a blog post
- 2018, August: The Community takes the lead
- 2019, April: Version 1.14 introduces **Local Persistent Volumes**
- 2019, August: EDB team starts the Kubernetes initiative
- 2020, June: we publish this blog about benchmarking local PVs on bare metal
- 2020, June: Data on Kubernetes Community founded
- 2021, February: EDB Cloud Native Postgres (CNP) 1.0 released
- 2022, May: **EDB donates CNP** and open sources it under CloudNativePG



Enabling the same PostgreSQL everywhere

From self-managed to fully managed DBaaS in the Cloud

- Same applications
- Faster innovation
- Performance and scalability
- Stability, security and control
- Seamless integration



Private



Hybrid



Multi-cloud



Public



Bare Metal



Virtual Machines



Containers



A kubernetes operator for Postgres



Kubernetes adoption is rising and it is already the de facto **standard orchestration tool**



PostgreSQL clusters “**management the kubernetes way**” enables many cloud native usage patterns, e.g. spinning up, disposable clusters during tests, one cluster per microservice and one database per cluster



CNP tries to encode years of experience managing PostgreSQL clusters into **an Operator which should automate all the known tasks a user could be willing to do**

Our PostgreSQL operator must simulate the work of a DBA



Win Technology



 **Red Hat**
OpenShift



CloudNativePG



EDB
Postgres® AI



CloudNativePG/EDB Postgres for Kubernetes

CloudNativePG



- Kubernetes operator for PostgreSQL
- “Level 5”, Production ready
- Day 1 & 2 operations of a Postgres database
- Open source (May 2022)
 - Originally created by EDB
 - Apache License 2.0
 - Vendor neutral openly governed
 - 2500+ stars on GitHub
- Extends the K8s controller
 - Status of the `Cluster`
 - “no Patroni, No statefulsets”
- Immutable application containers
- Fully declarative

EDB Postgres for Kubernetes

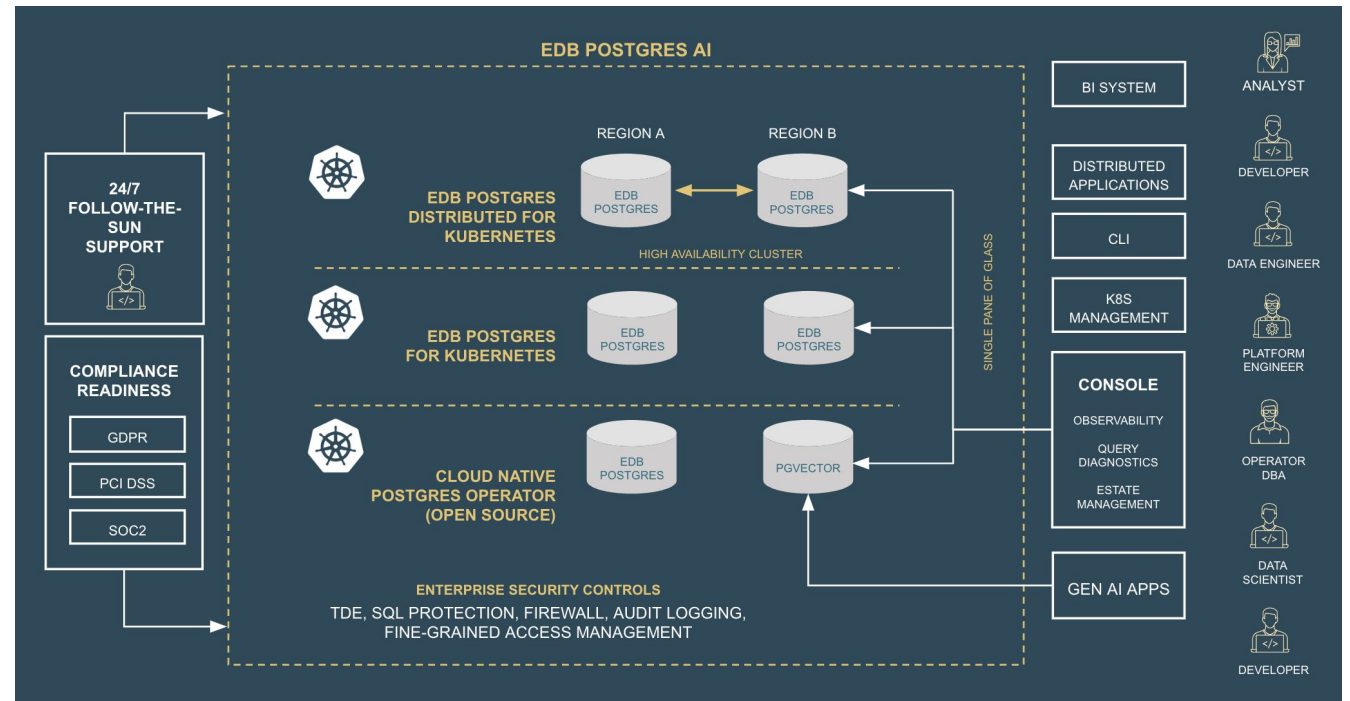
- Fork of CloudNativePG
 - +
- Provides Long Term Support
- Access to EDB Postgres Extended (TDE)
- Access to EDB Postgres Advanced (TDE + Oracle Compatibility layer)
- Red Hat OpenShift compatibility
- Kubernetes level backup integration
 - Generic external backup interface

Operator Capabilities Levels



Features

- Automated failover
- Services for RW and RO workloads
- Affinity control
- Backup and Recovery
- Rolling updates
- Scale up/down of read replicas
- Fencing and hibernation
- Native Prometheus exporters
- Log in JSON format to stdout
- OpenShift compatibility
- TDE (in EDB Postgres for Kubernetes)
- ... and much more



DEMO



Demo Architecture

EDB My App

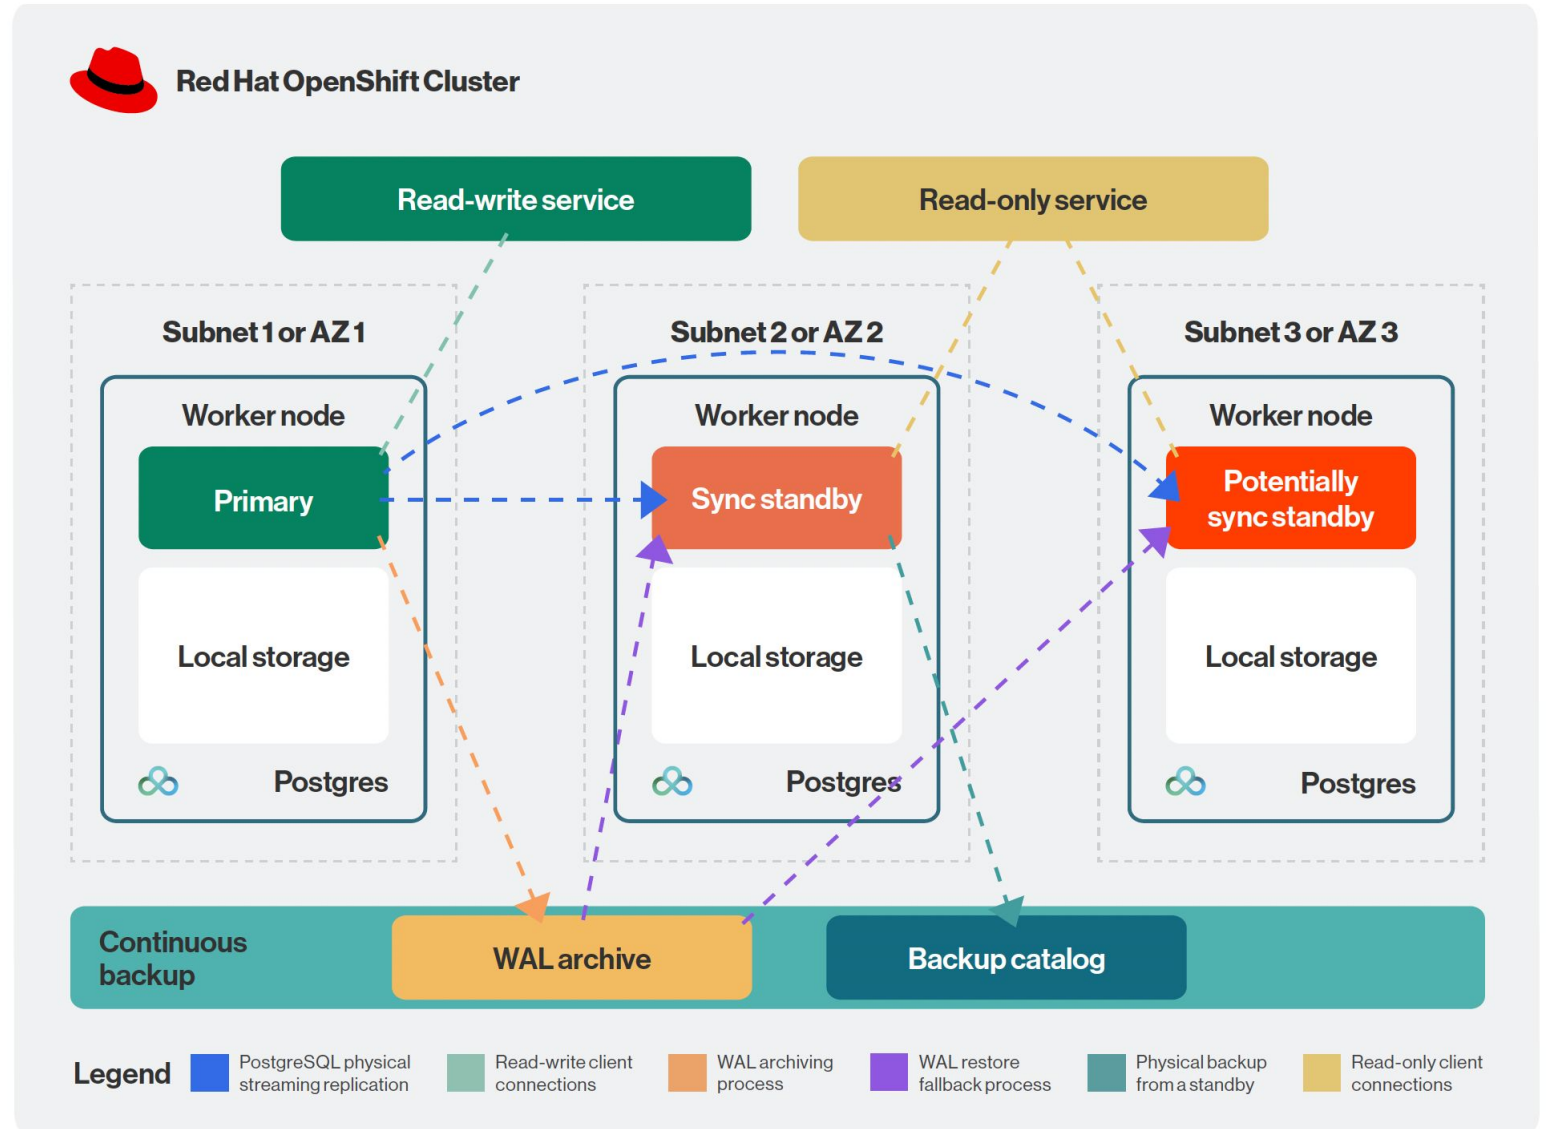
CLUSTER NAME	IP	REPLICATION ROLE	CONNECTED TO PRIMARY
cluster-example-1	10.1.35.207	Primary	10.1.35.207
cluster-example-2	10.1.35.209	Standby	
cluster-example-3	10.1.35.211	Standby	

Search
Search name

Customers [Create new](#)

NAME
Test 1
Test 2
Test 3

< Page 1 >



Features shown during the demo

- Kubernetes plugin install
- CloudNativePG operator install
- Postgres cluster install
- Insert data in the cluster
- Switchover (promote)
- Failover
- Backup
- Recovery
- Scale out/down
- Point In Time Recovery (PITR)
- Fencing
- Monitoring
- Rolling updates (minor and major)



Red Hat OpenShift



Red Hat OpenShift Operator

OpenShift Local cluster is for development and testing purposes. DON'T use it for production.


Red Hat OpenShift

Project: openshift-operators

Installed Operators

Installed Operators are represented by ClusterServiceVersions within this Namespace. For more information, see the [Understanding Operators documentation](#). Or create an Operator and ClusterServiceVersion using the [Operator SDK](#).


Name Search by name...

Name	Managed Namespaces	Status	Provided APIs
 EDB Postgres for Kubernetes 1.23.2 provided by The EDB Postgres for Kubernetes Contributors	All Namespaces	✓ Succeeded Up to date	Backups Cluster Image Catalog Cluster Image Catalog View 2 more...




Red Hat OpenShift Cluster

OpenShift Local cluster is for development and testing purposes. DON'T use it for production.

 kubeadmin ▾

Project: openshift-operators ▾



[Installed Operators](#) > Operator details

 **EDB Postgres for Kubernetes** Actions ▾
1.23.2 provided by The EDB Postgres for Kubernetes Contributors

Details | **YAML** | Subscription | Events | All instances | Backups | Cluster Image Catalog | **Cluster** | Image Catalog | Pooler | Scheduled Bac

Clusters Show operands in: All namespaces Current namespace only [Create Cluster](#)

Name ▾ Search by name... /

Name ↑	Kind ↑	Namespace ↑	Status ↑	Labels ↑
 cluster-sample	Cluster	 openshift-operators	Phase: Cluster in healthy state	No labels



Red Hat OpenShift Cluster details

OpenShift Local cluster is for development and testing purposes. DON'T use it for production.

Red Hat OpenShift

Project: openshift-operators

Installed Operators > cloud-native-postgresql.v1.23.2 > Cluster details

cluster-sample Cluster in healthy state Actions

[Details](#) [YAML](#) [Resources](#) [Events](#)

Cluster overview

Working Pods

3

instancesStatus

Name
cluster-sample

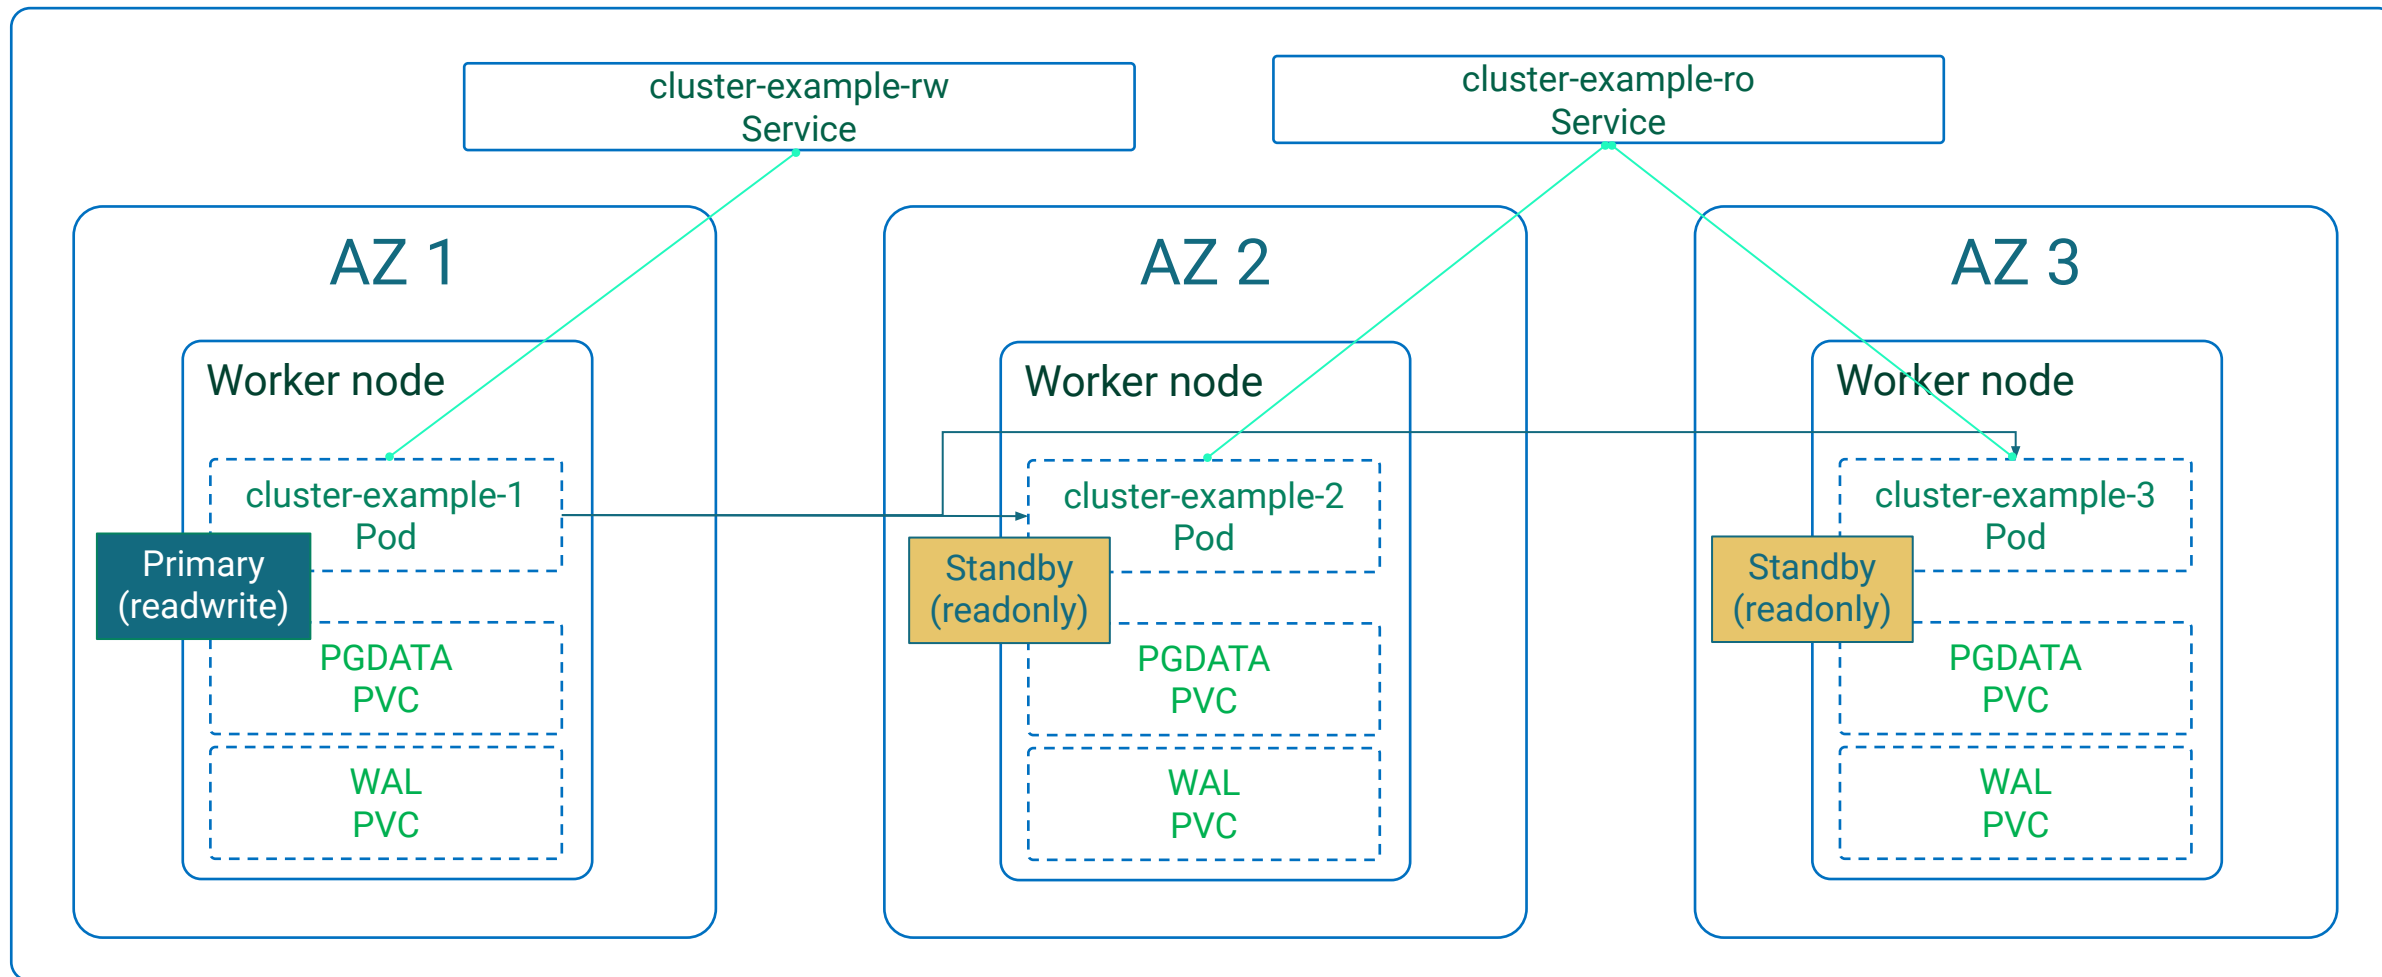
Current Primary Pod
cluster-sample-1



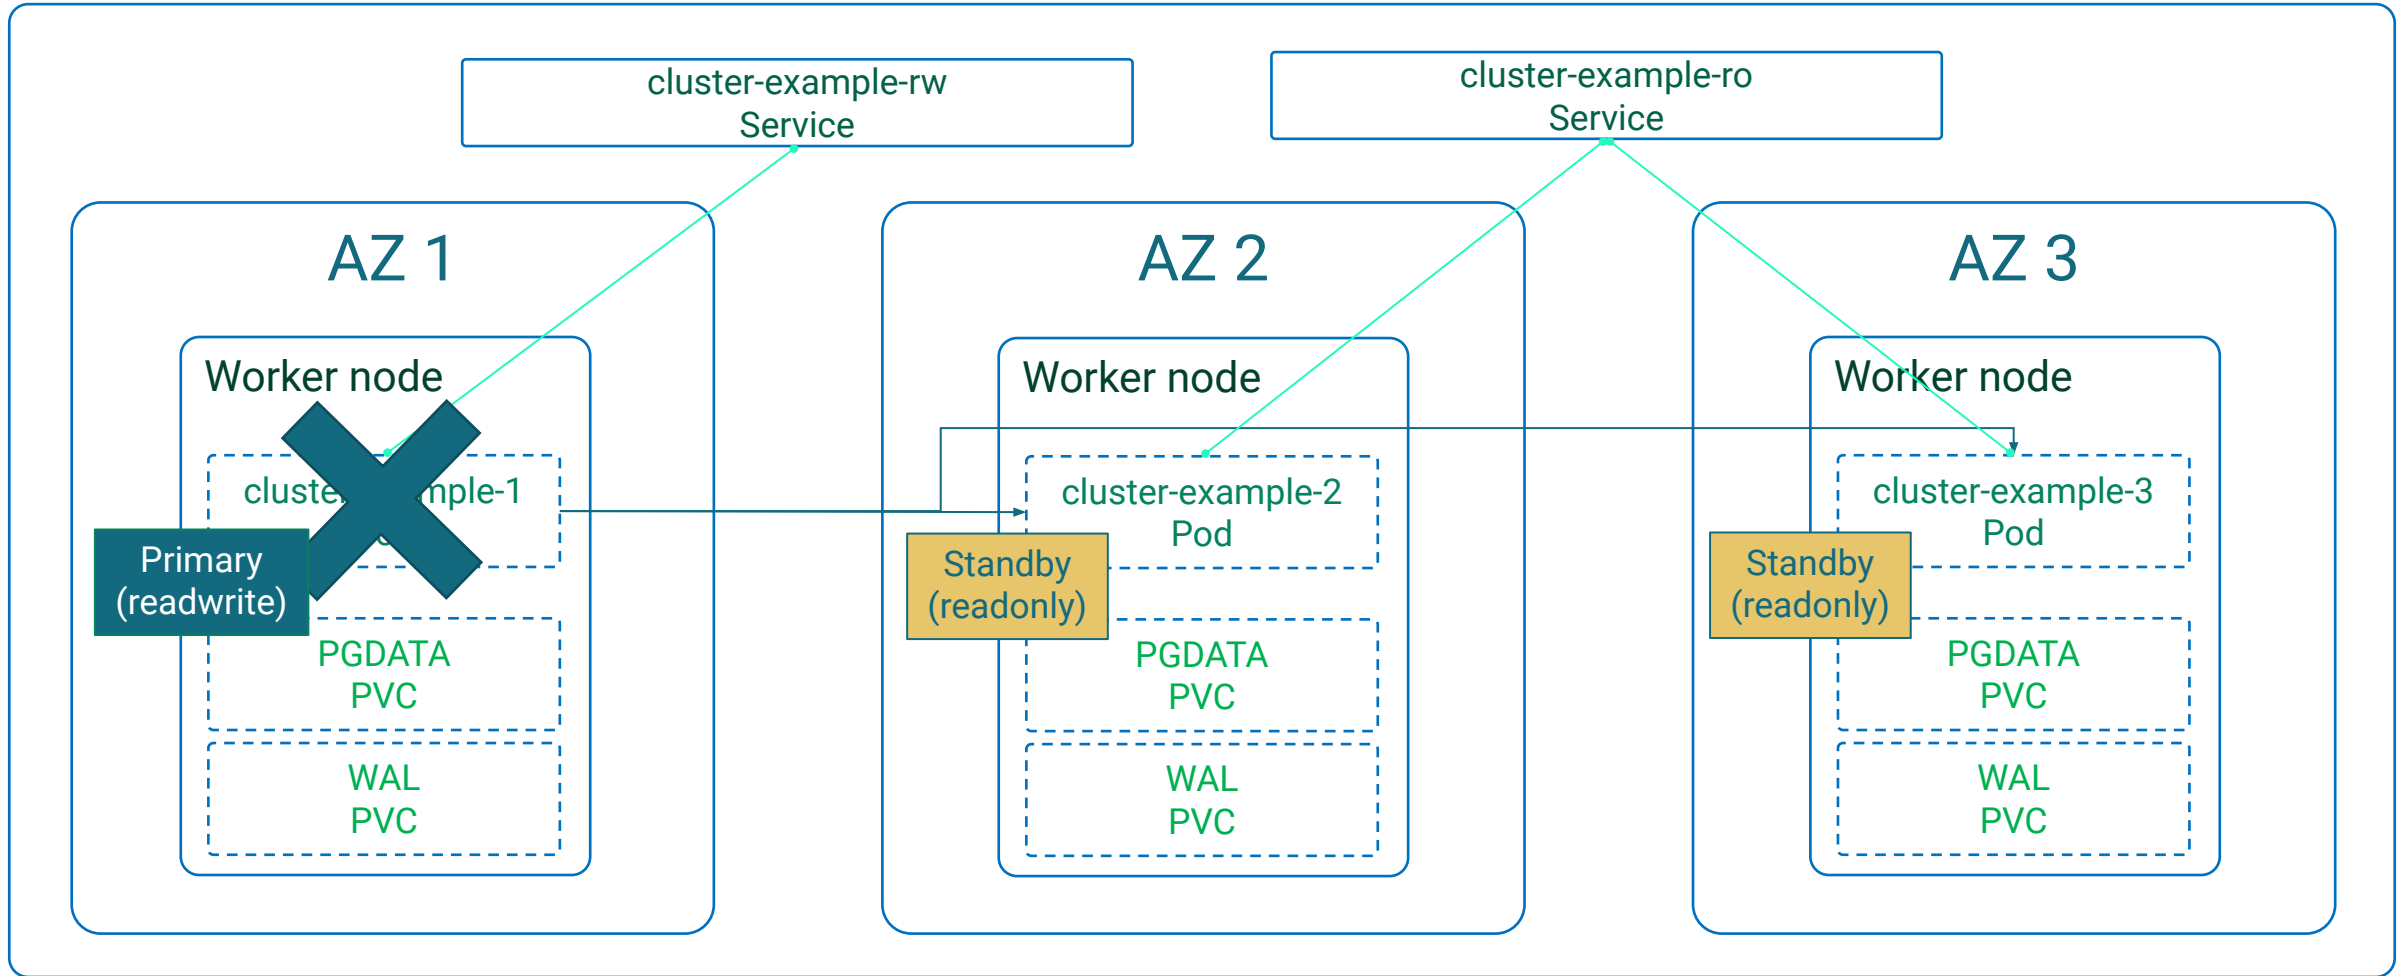
What happens under the hood



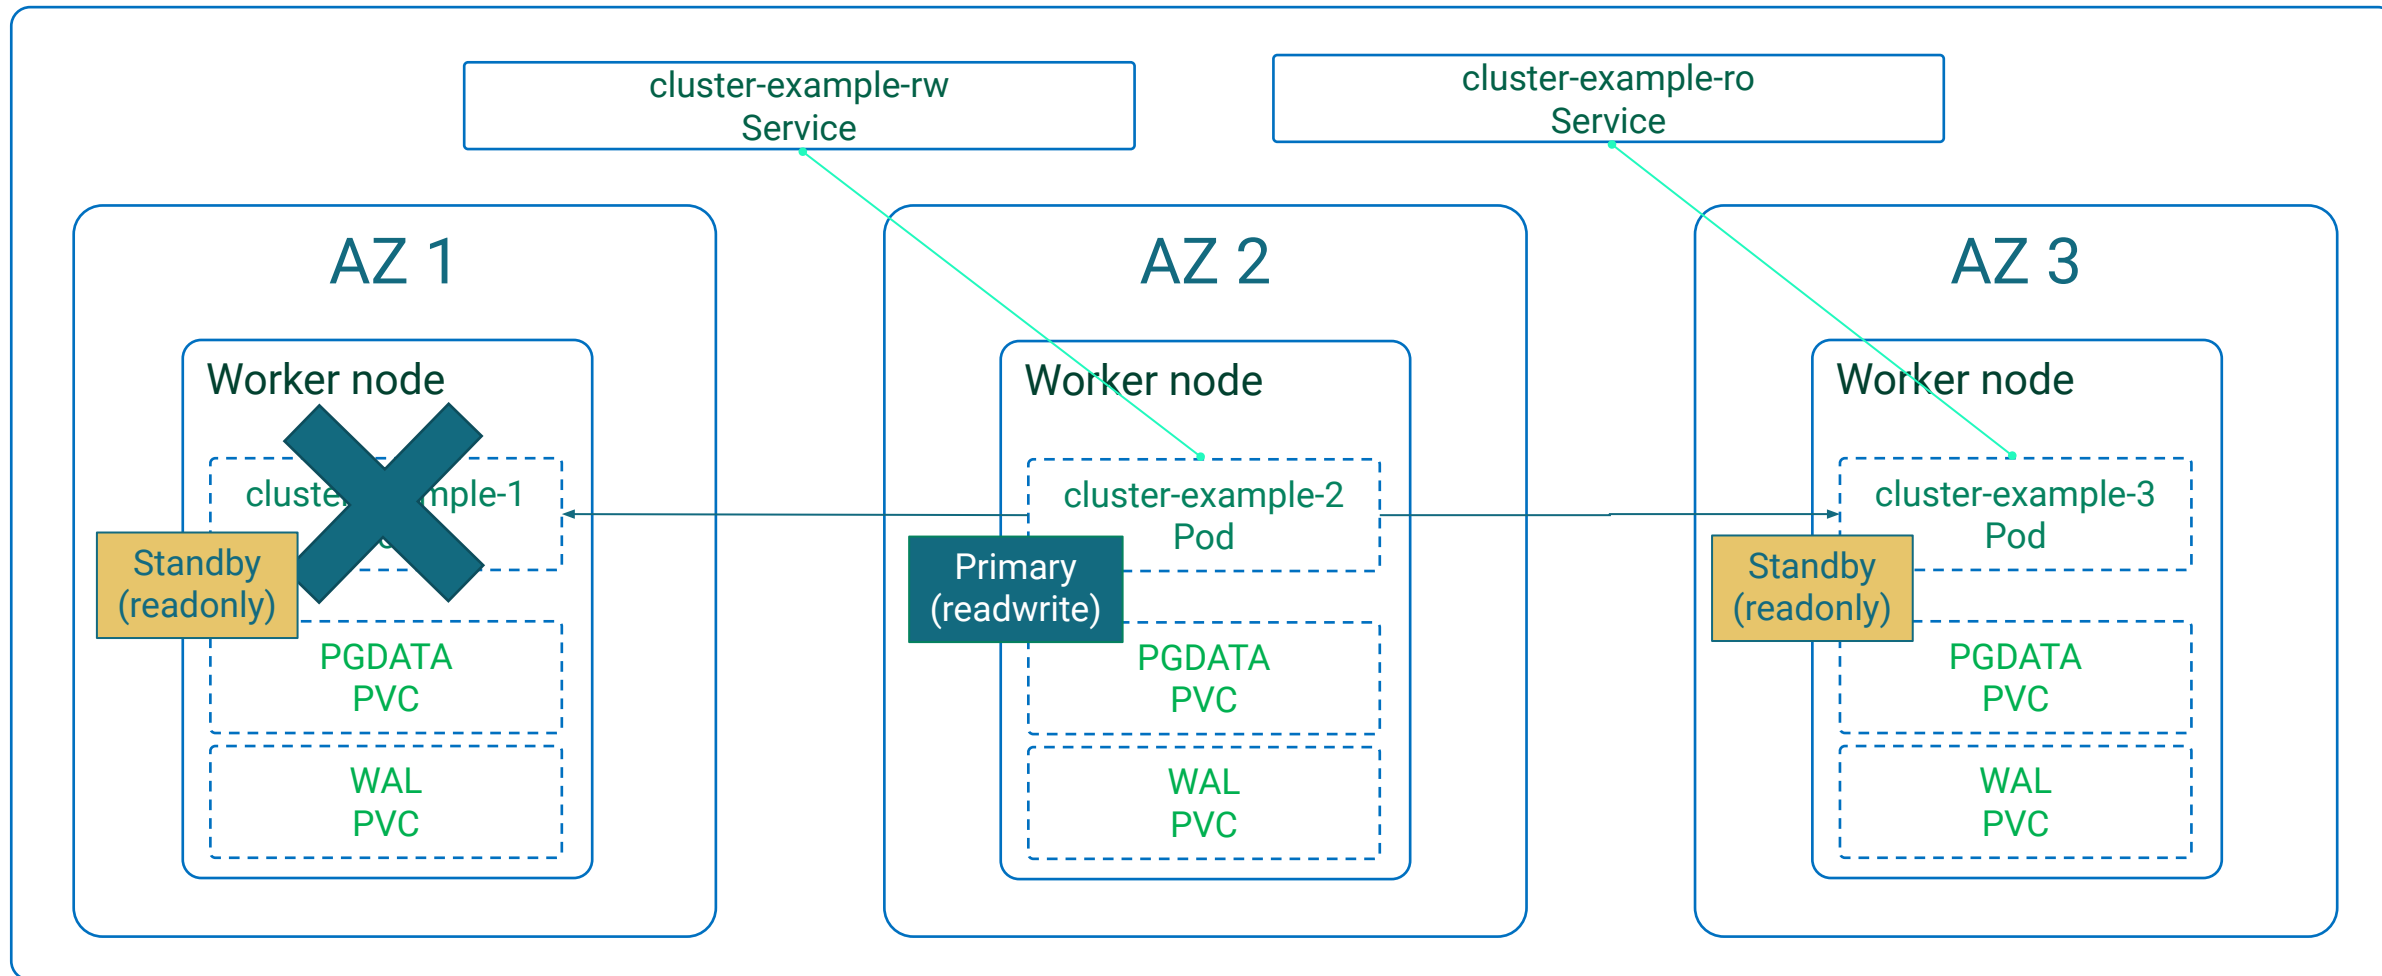
This is what happens under the hood



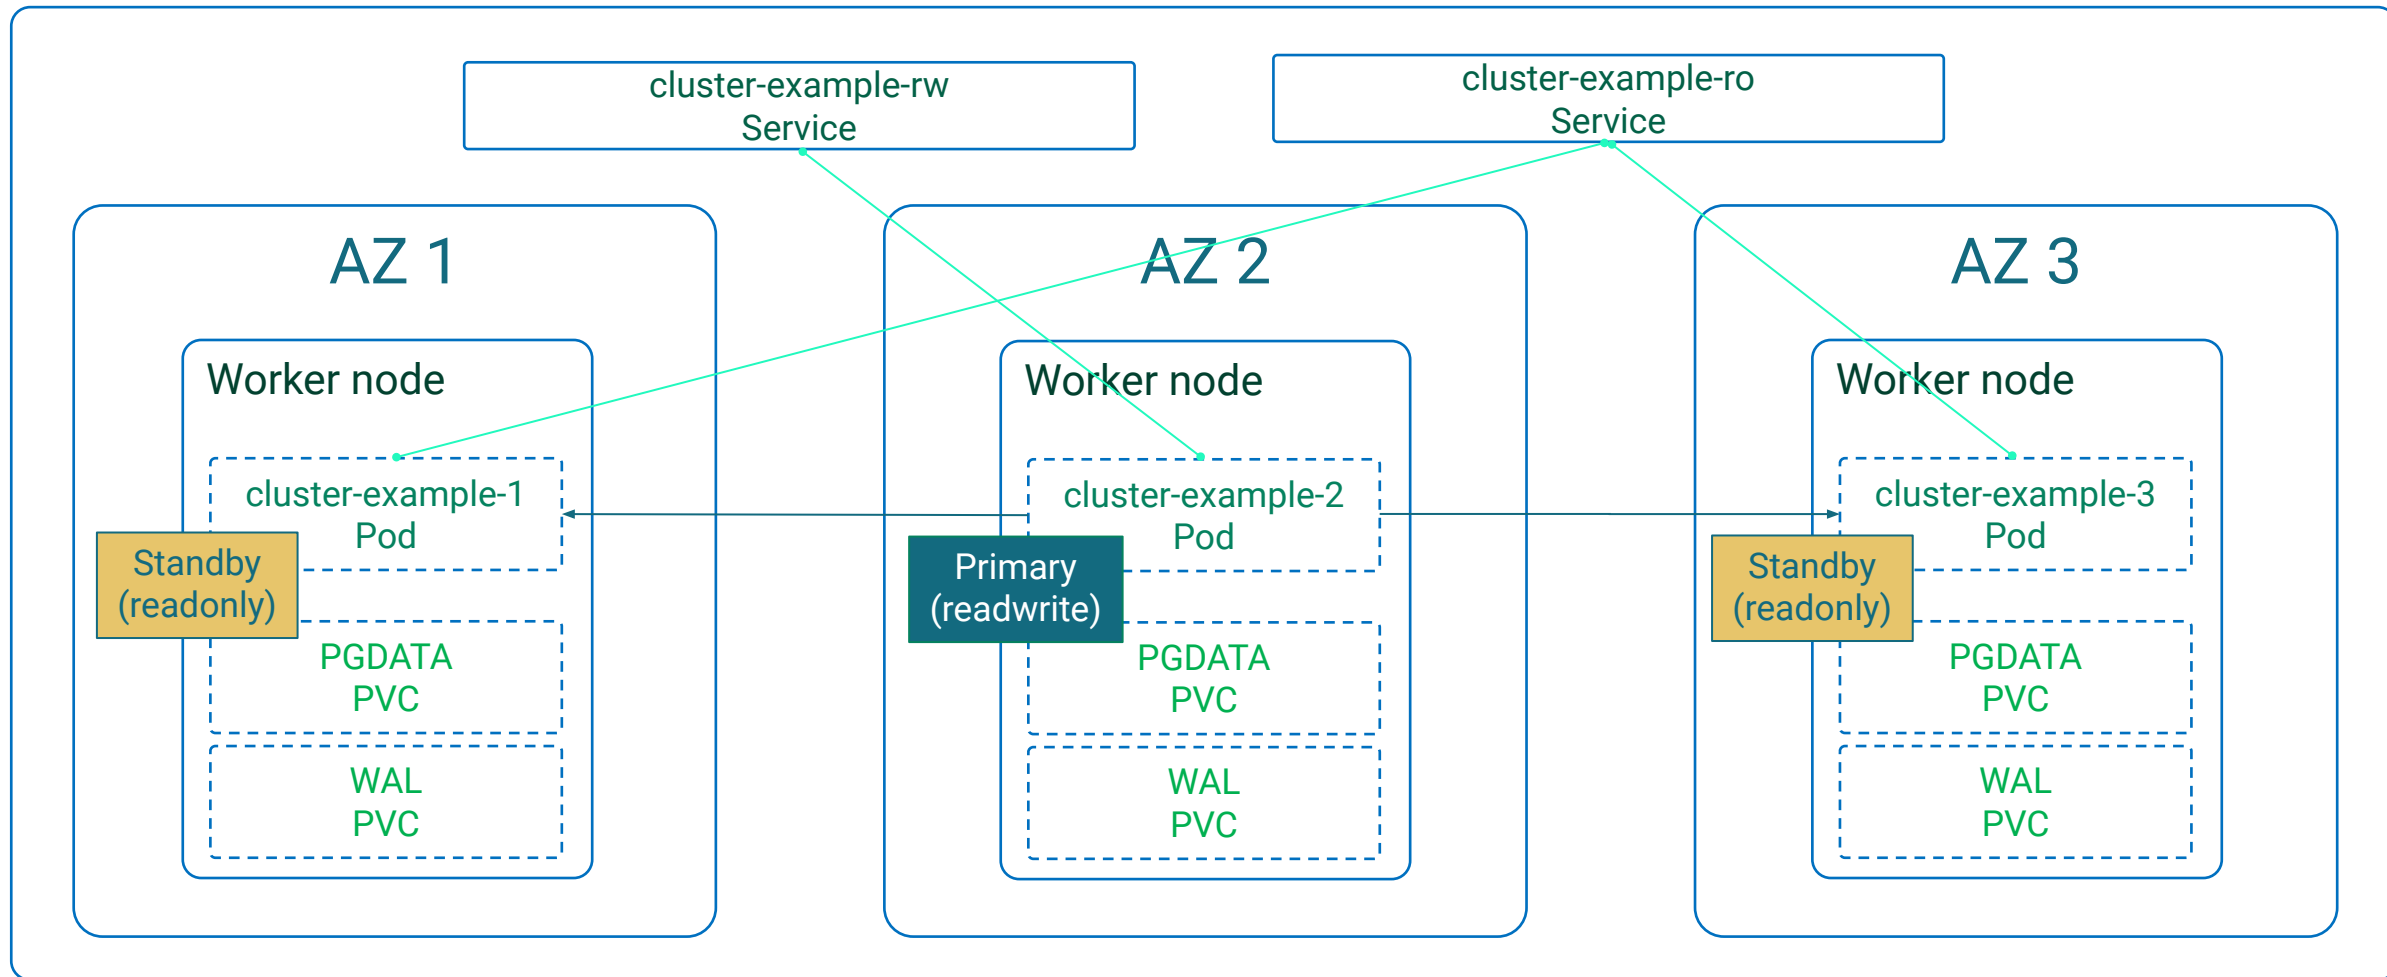
This is what happens under the hood



This is what happens under the hood



This is what happens under the hood



Subscriptions



Subscription = Software Plan + Support Plan




Capabilities and Tooling



Management/Monitoring

Postgres Enterprise Manager
pgAdmin



High Availability

EDB Postgres Distributed
Failover Manager
Repmgr
Patroni



Backup and Recovery

Barman
pgBackRest



Migration

Migration Portal
Migration Toolkit
Replication Server



Integration

Connectors
Foreign Data Wrappers
Connection Poolers

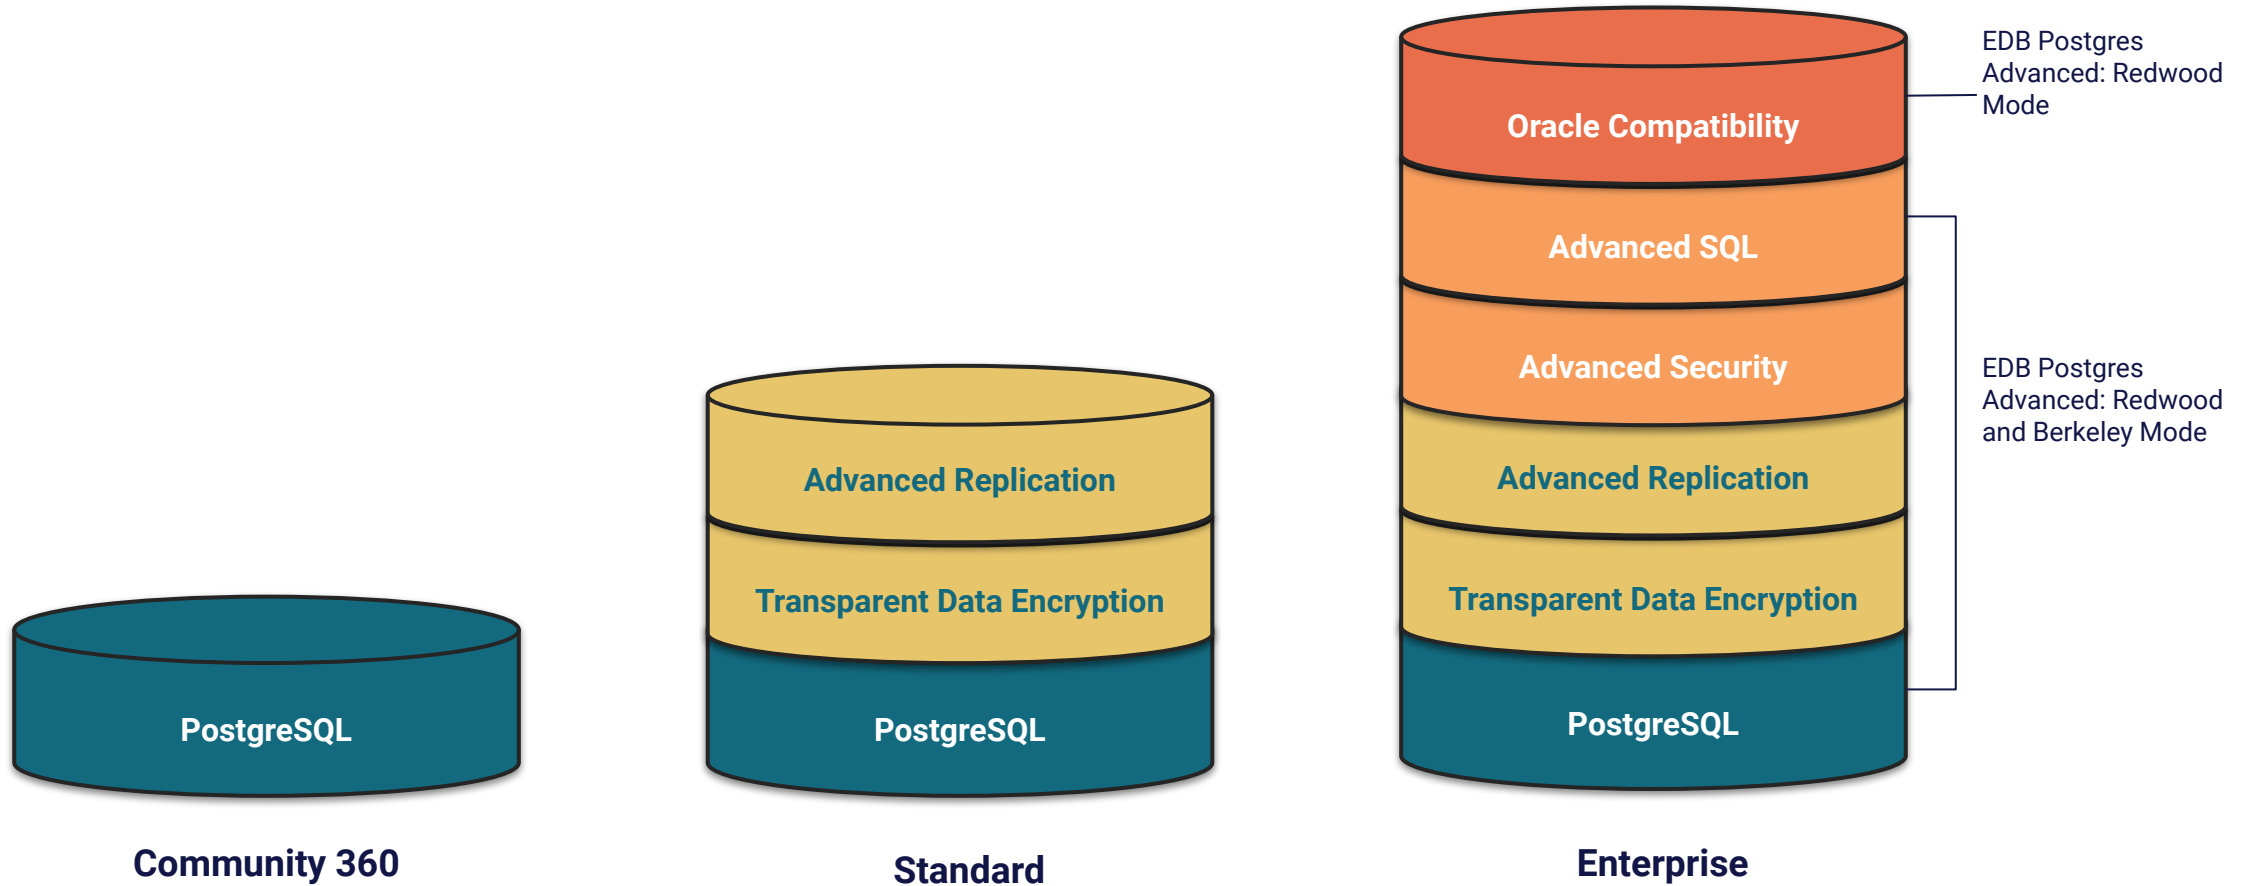


Kubernetes

EDB Postgres for Kubernetes
CloudNativePG



EDB Databases & Subscription Plans



Software Plans

Software	Community 360	Standard	Enterprise
PostgreSQL (Open Source)	✓	✓	✓
CloudNativePG	✓	✓	✓
Open Source Tools	✓	✓	✓
Technical Support	✓	✓	✓
EDB Tools (HA, Monitoring, Backup & Recovery, Migration)		✓	✓
EDB Postgres for Kubernetes (Red Hat OpenShift)		✓	✓
TDE (Transparent Data Encryption) for VM/Baremetal		✓	✓
TDE (Transparent Data Encryption) for Kubernetes/ Red Hat OpenShift		✓	✓
EDB Postgres Advanced Server (EPAS – Oracle Compatibility)			✓
EDB Postgres Distributed (Extreme High Availability)		Add-on	Add-on
Technical Account Manager	Optional	Optional	Optional



Support Plan

Support Plan	Community 360	Standard	Enterprise
Premium Support	✓	✓	✓
Production Support	✓	✓	✓
Basic Support	✓	✓	✓



Thank you



Annexes

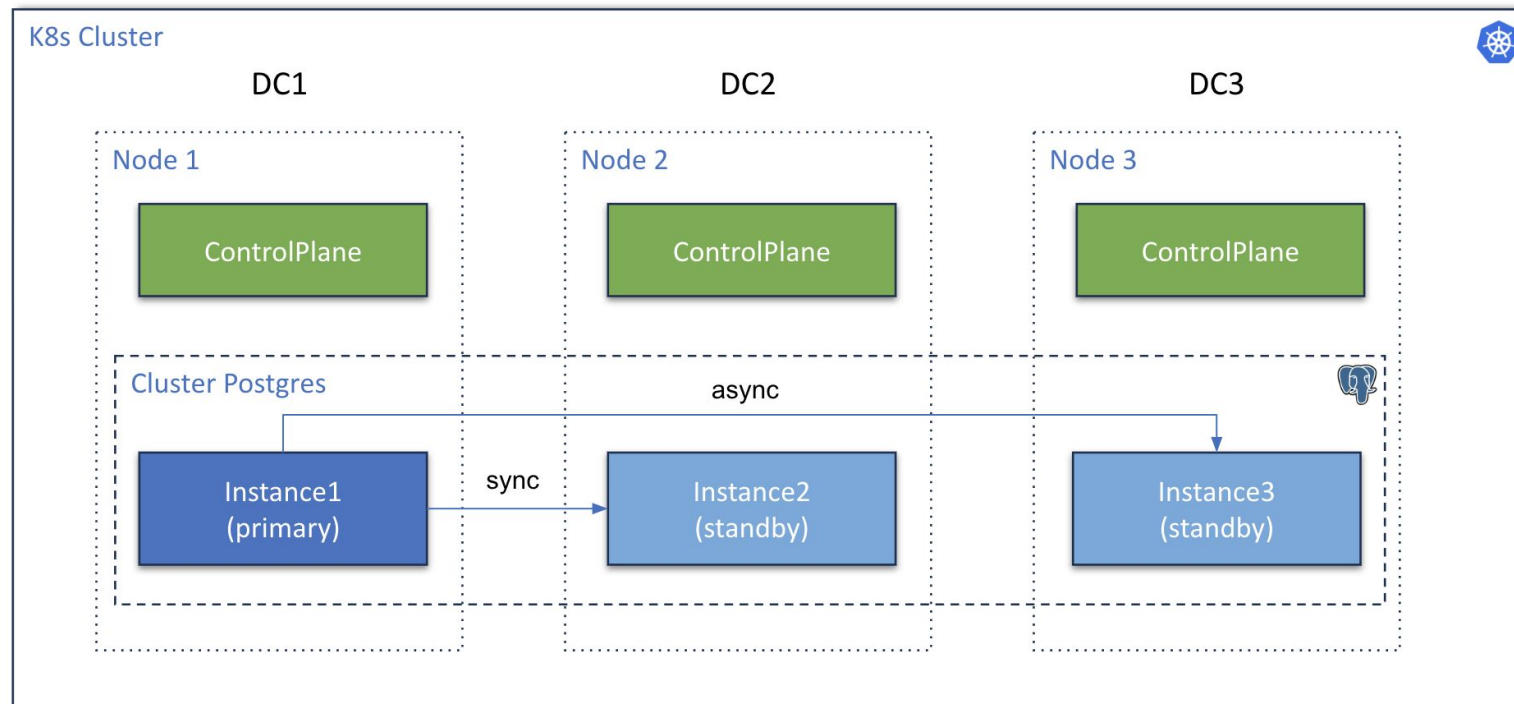


Architectures



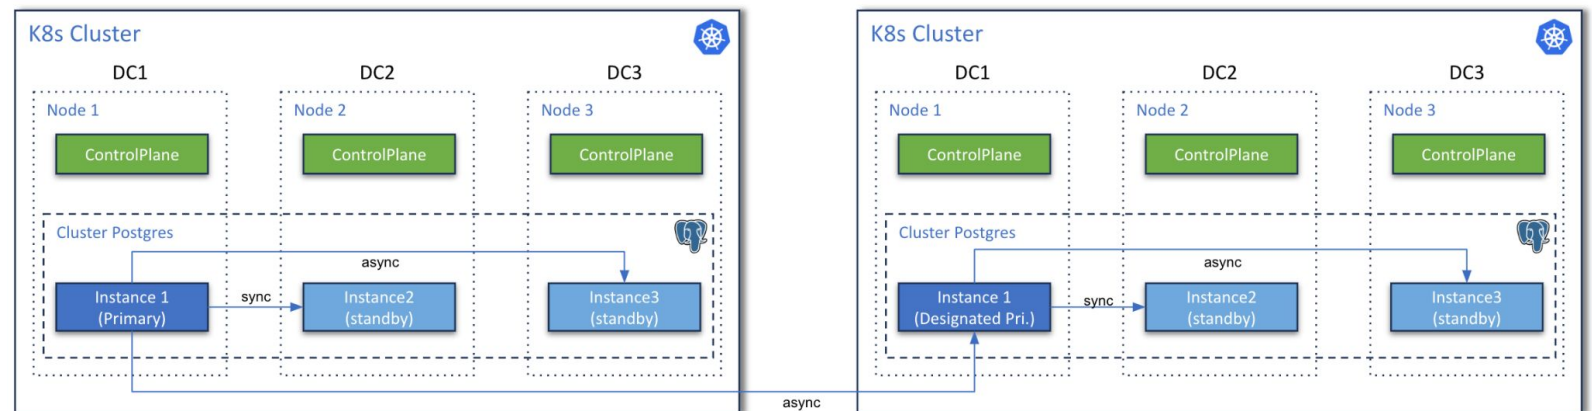
Architectures - Multi-availability zone Kubernetes clusters

- The multi-availability zone Kubernetes architecture with three (3) or more zones is the one that we recommend for PostgreSQL usage. This scenario is typical of Kubernetes services managed by Cloud Providers.



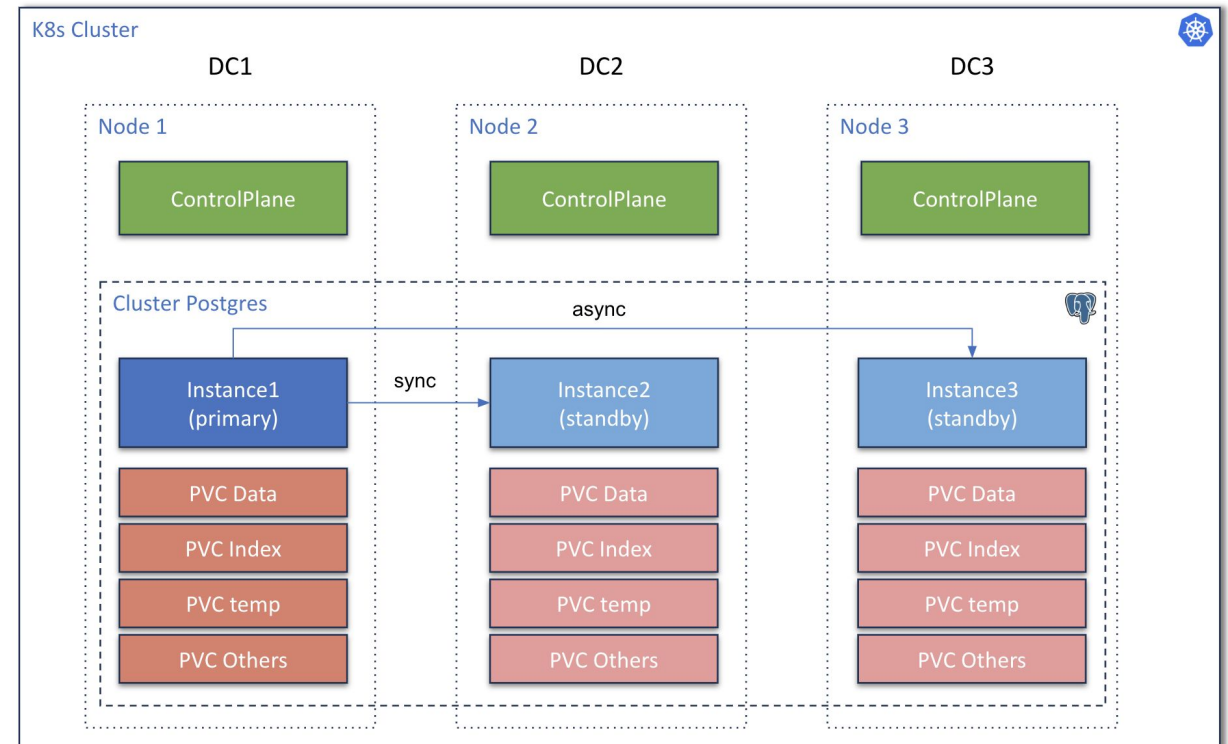
Architectures - Single availability zone Kubernetes clusters

- If your Kubernetes cluster has only one availability zone, EDB Postgres for Kubernetes still provides you with a lot of features to improve HA and DR outcomes for your PostgreSQL databases, pushing the single point of failure (SPoF) to the level of the zone as much as possible - i.e. the zone must have an outage before your EDB Postgres for Kubernetes clusters suffer a failure.
- This scenario is typical of self-managed on-premise Kubernetes clusters, where only one data center is available.



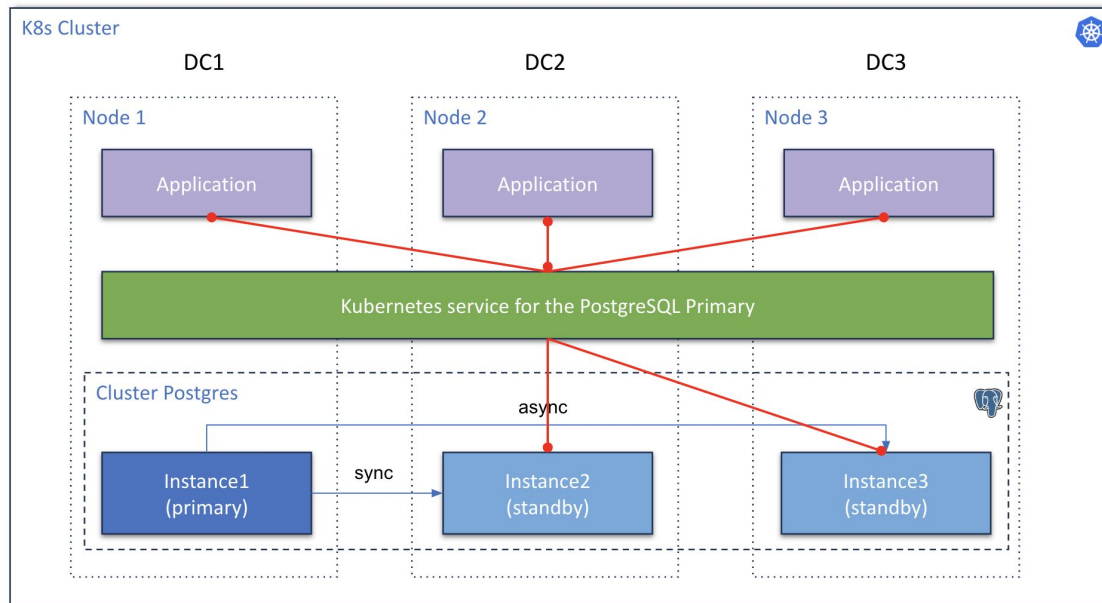
PostgreSQL architecture

- EDB Postgres for Kubernetes supports clusters based on asynchronous and synchronous streaming replication to manage multiple hot standby replicas within the same Kubernetes cluster, with the following specifications:
 - One primary, with optional multiple hot standby replicas for HA
 - Available services for applications:
 - -rw: applications connect only to the primary instance of the cluster
 - -ro: applications connect only to hot standby replicas for read-only workloads
 - -r: applications connect to any of the instances for read-only workloads
 - Shared-nothing architecture recommended for better resilience of the PostgreSQL cluster

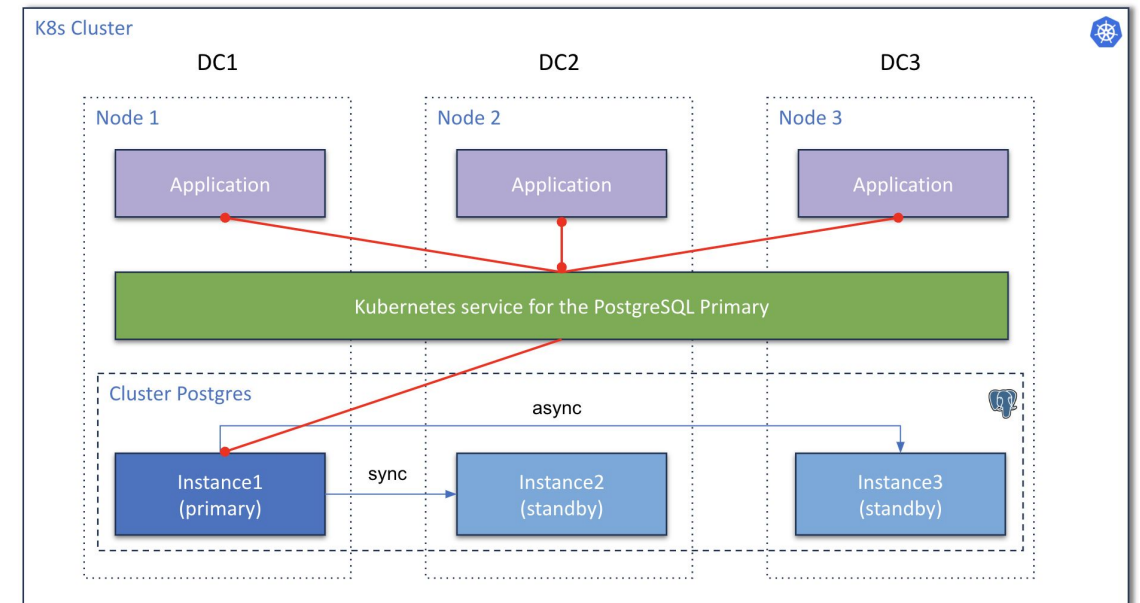


PostgreSQL workloads

Read-write workloads

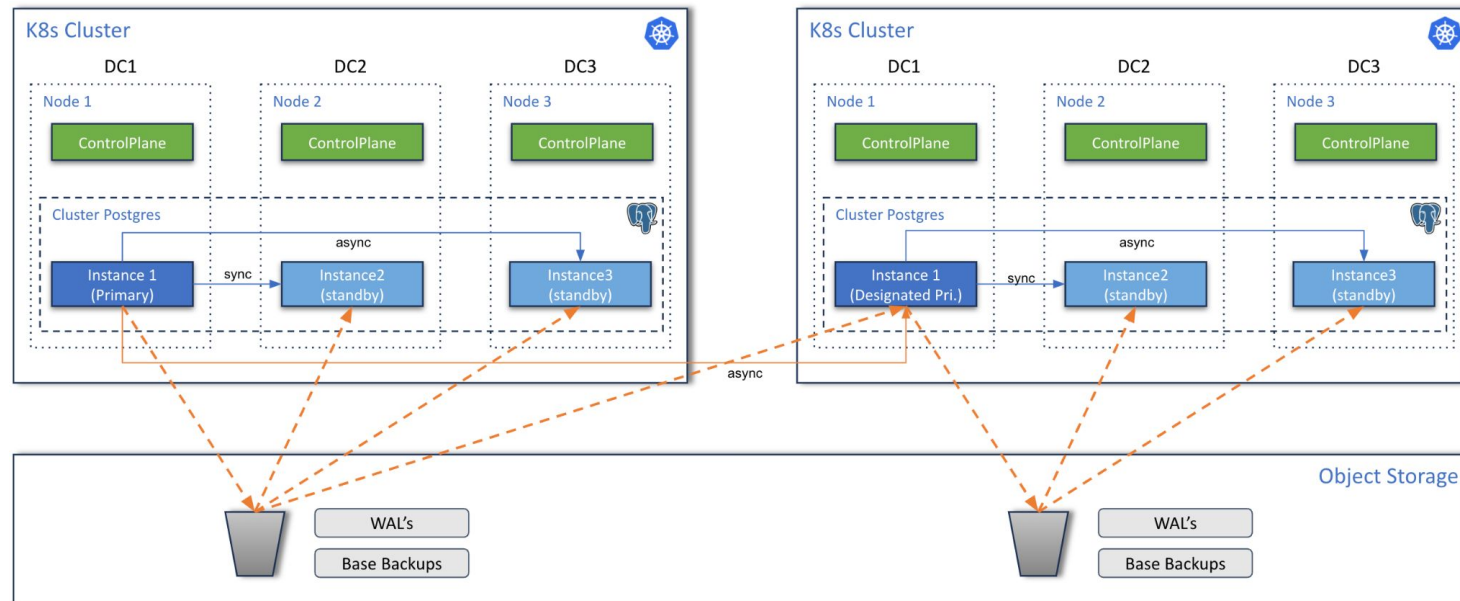


Read-only workloads



PostgreSQL Disaster Recovery

- PostgreSQL cluster spanning over two different Kubernetes clusters, where the primary cluster is in the first Kubernetes cluster and the replica cluster is in the second. The second Kubernetes cluster acts as the company's disaster recovery cluster, ready to be activated in case of disaster and unavailability of the first one.
- A replica cluster can have the same architecture of the primary cluster. In place of the primary instance, a replica cluster has a designated primary instance, which is a standby server with an arbitrary number of cascading standby servers in streaming replication (symmetric architecture).



Recommended architectures

<https://www.cncf.io/blog/2023/09/29/recommended-architectures-for-postgresql-in-kubernetes/>



The screenshot shows the top of a web page from the Cloud Native Computing Foundation. The header includes the logo, navigation links for 'About', 'Projects', 'Training', 'Community', and 'Blog & News', a 'Join' button, and a search icon. The main heading is 'Recommended architectures for PostgreSQL in Kubernetes'. Below this, it says 'By Gabriele Bartolini' and 'September 29, 2023'. A sub-headline reads 'Member post by Gabriele Bartolini, VP of Cloud Native at EDB'. The main text begins with a quote from Kelsey Hightower: 'You can run databases on Kubernetes because it's fundamentally the same as running a database on a VM', followed by a link to a tweet from Kelsey Hightower. The tweet text is: 'You can run databases on Kubernetes because it's fundamentally the same as running a database on a VM. The biggest challenge is understanding that rubbing Kubernetes on Postgres won't turn it into Cloud SQL.' The article text continues: 'Truth is that I agree with him now as much as I agreed with him back then. At that time, the holistic offering of storage capabilities in Kubernetes was still immature (local persistent volumes would become GA only the year after), the operator pattern – which in the meantime has proven to be crucial for stateful applications like databases – was yet to become widely accepted, and the Data on Kubernetes Community was more than two years away (second half of 2020). Nowadays, the situation is completely different. And I am sure that many people who've worked hard in the last few years to bring stateful workloads in Kubernetes agree with me that Kelsey's recent powerful words will contribute to reversing the public perception and facilitate our mission – provided we keep doing great.'

