Using Kubernetes to Streamline Deployment at Scale

Andrew Mumford | Sr Solution Consultant

EDB

EDB IN SUMMARY

EDB is the world's largest software, support, and services company focused exclusively on PostgreSQL. We are proud to serve some of the world's leading financial services, government, media & communications, and information technology organisations. Our 16 offices worldwide enable us to deploy our global expertise in all your business locations.

POSTGRESQL **COMMUNITY** LEADERSHIP

- **40%** of Postgres code contributed
- **>300** Dedicated Postgres engineers
- **3 of 7** Postgres Core Team Members

EDB **PLATFORM** (SOFTWARE & TOOLS)

- Databases: PostgreSQL, EPAS
- Tools: Variety of supported open source and proprietary tools for High availability, backup, monitoring and migration

EDB **SUPPORT**

- 24/7 world-class support
- Experienced support engineers, with the world's leading Postgres contributors
- Cloud/Remote DBA Service, Technical Account Management, CTO Office

EDB **SERVICES**

- PostgreSQL deployment, design, migration
- Postgres Optimization: Best practices
- Enterprise Strategy: Use-case driven PostgreSQL architectures
- Embedded PostgreSQL experts

EDB™

ENABLING THE SAME POSTGRES EVERYWHERE

From self-managed to fully managed DBaaS in the cloud

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

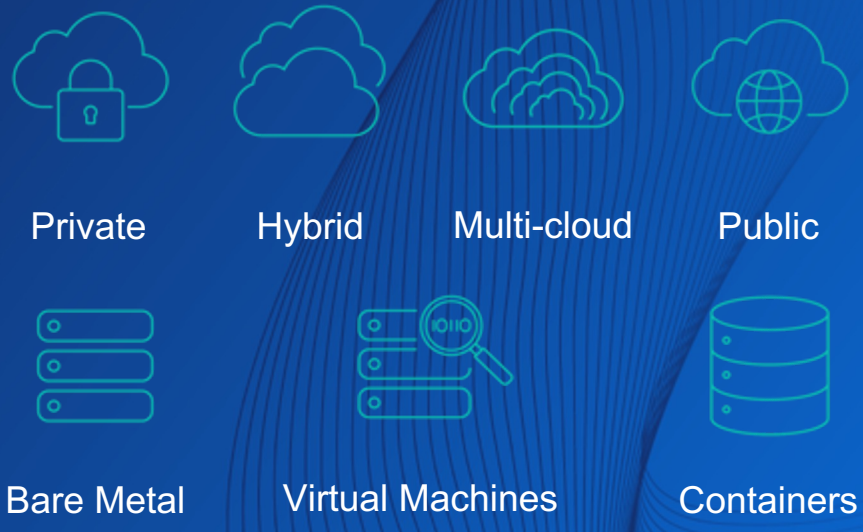

EDB

Why Kubernetes?

Services, Load Balancing, and Networking 	Health checking 	Storage management 
 Automated Scheduling	 Scalability: scale-up/down	 Rolling Deployments

EDB

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

EDB

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
    - 3000+ 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

EDB

What happens under the hood?

EDB

This is what happens under the hood

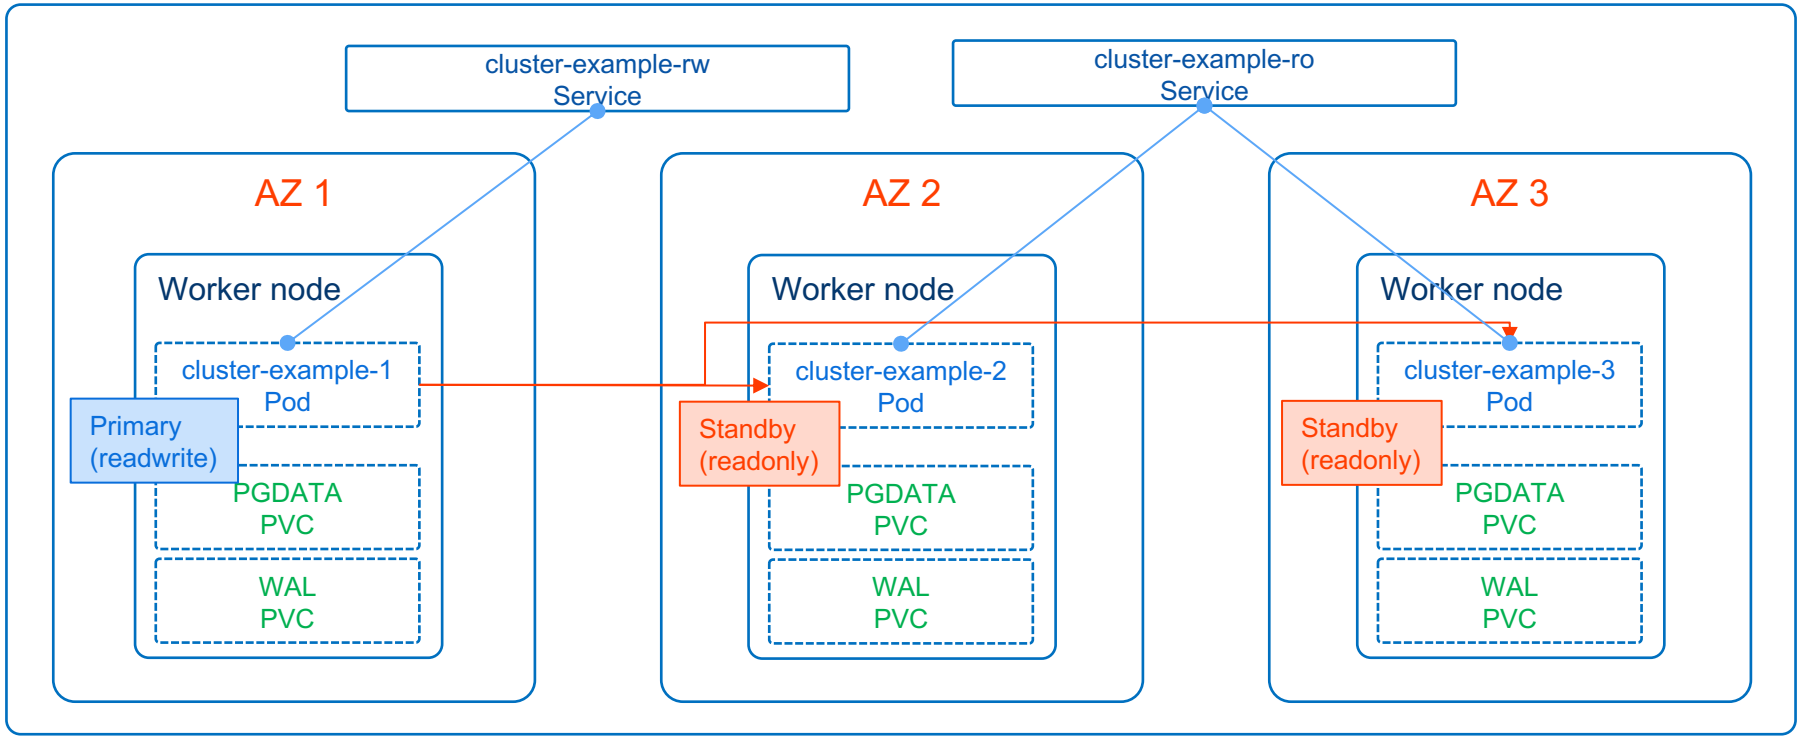

This is what happens under the hood

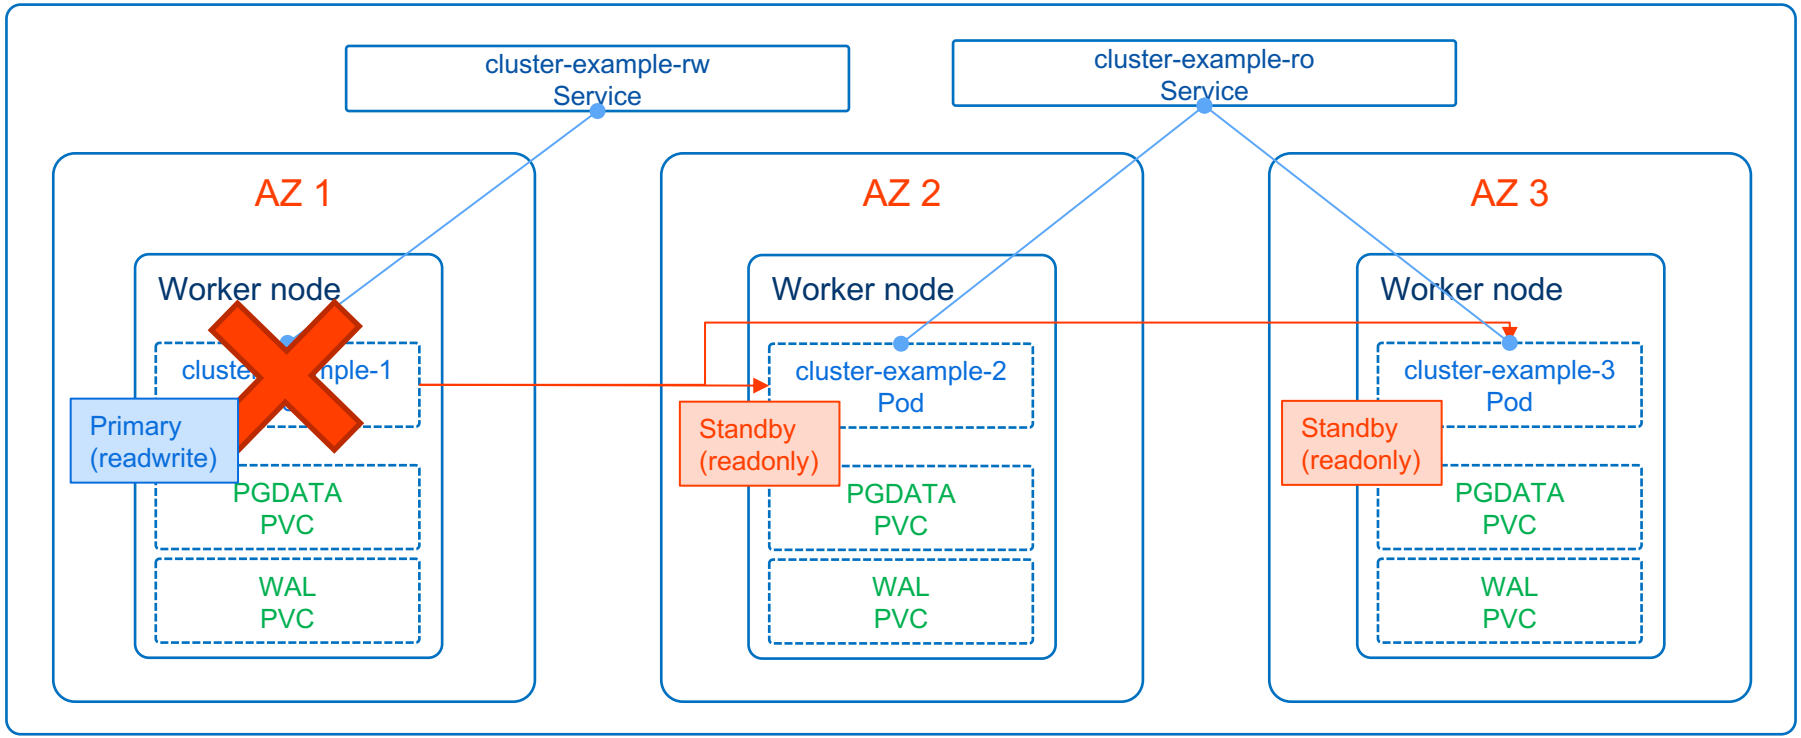

This is what happens under the hood

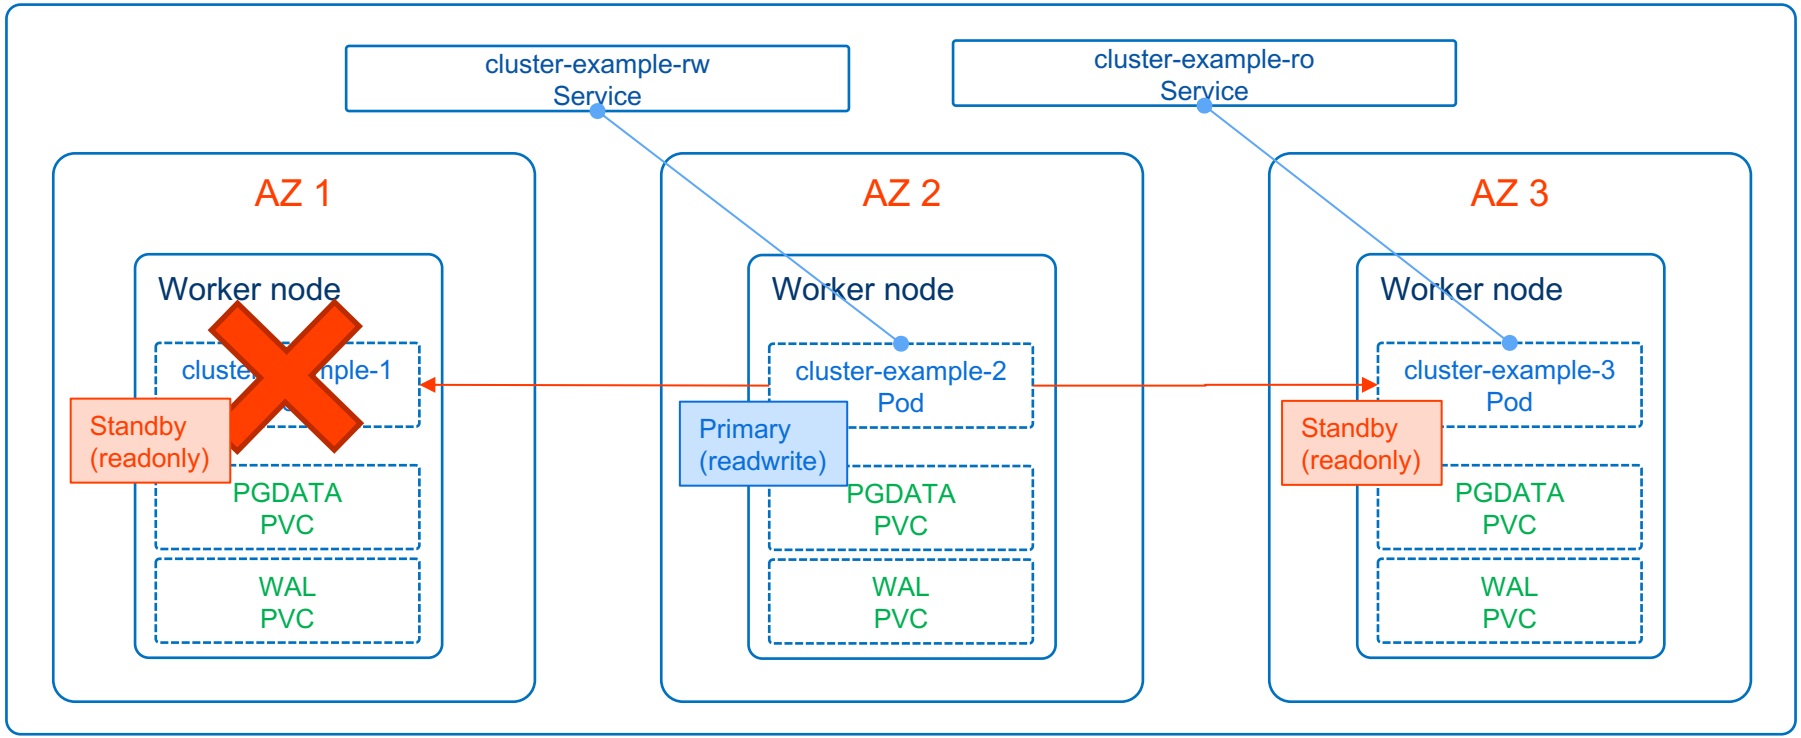

This is what happens under the hood

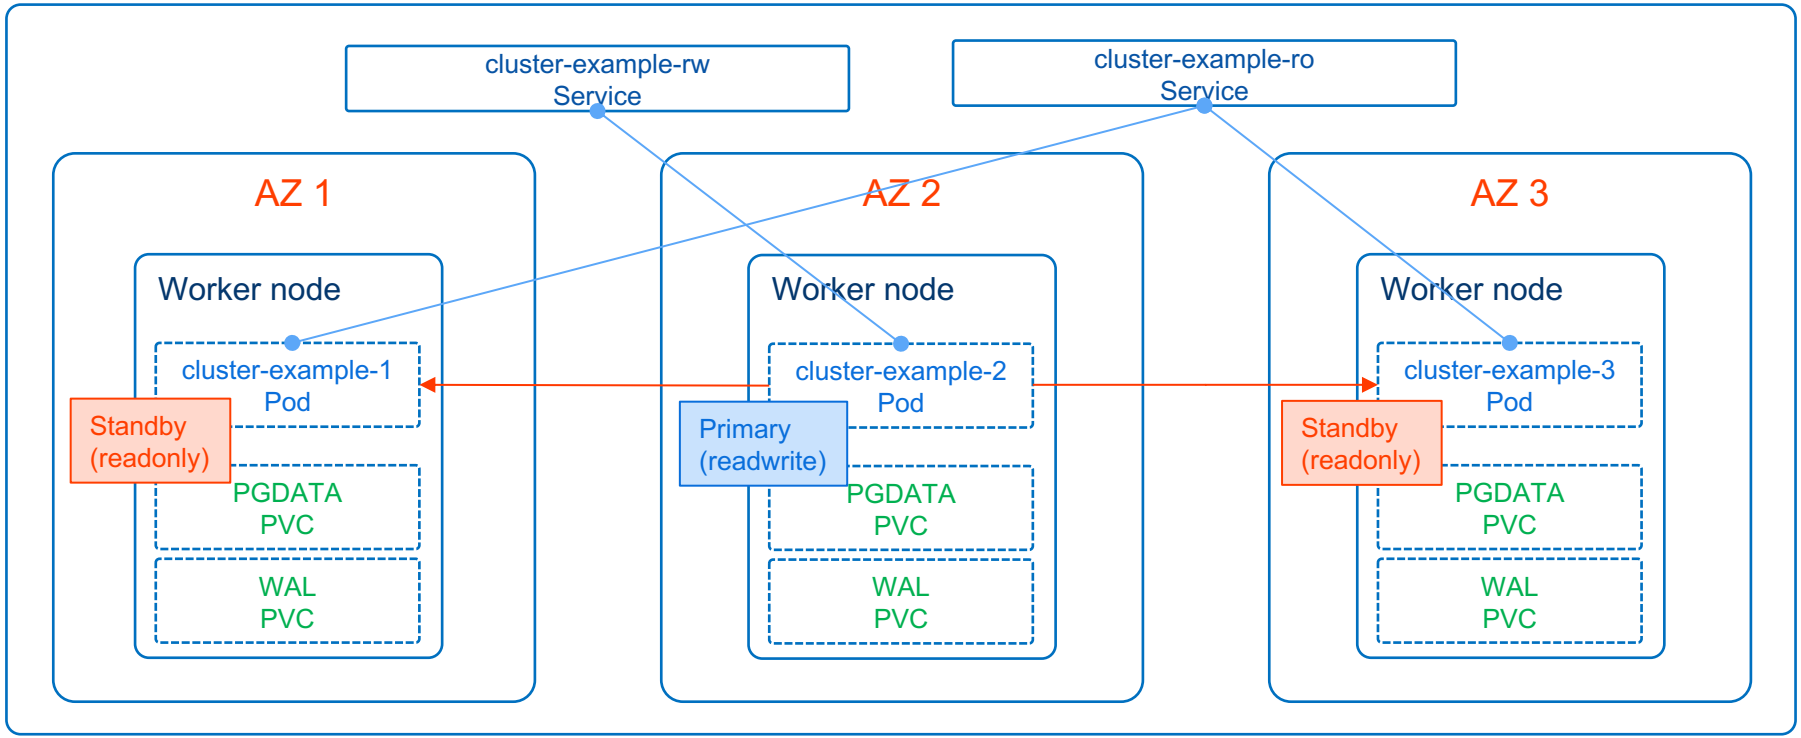

Demo

EDB

Features shown during the demo

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

EDB

Demo Architecture

kubernetes

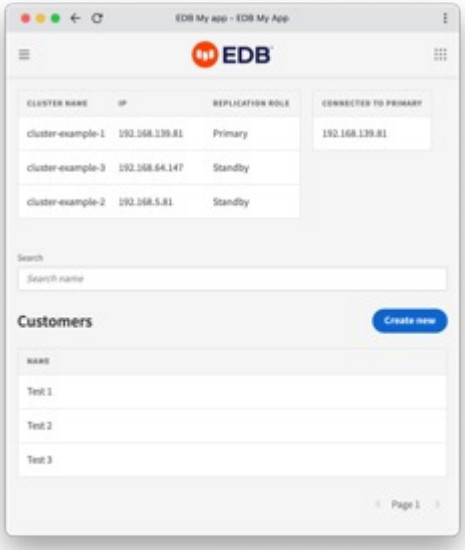

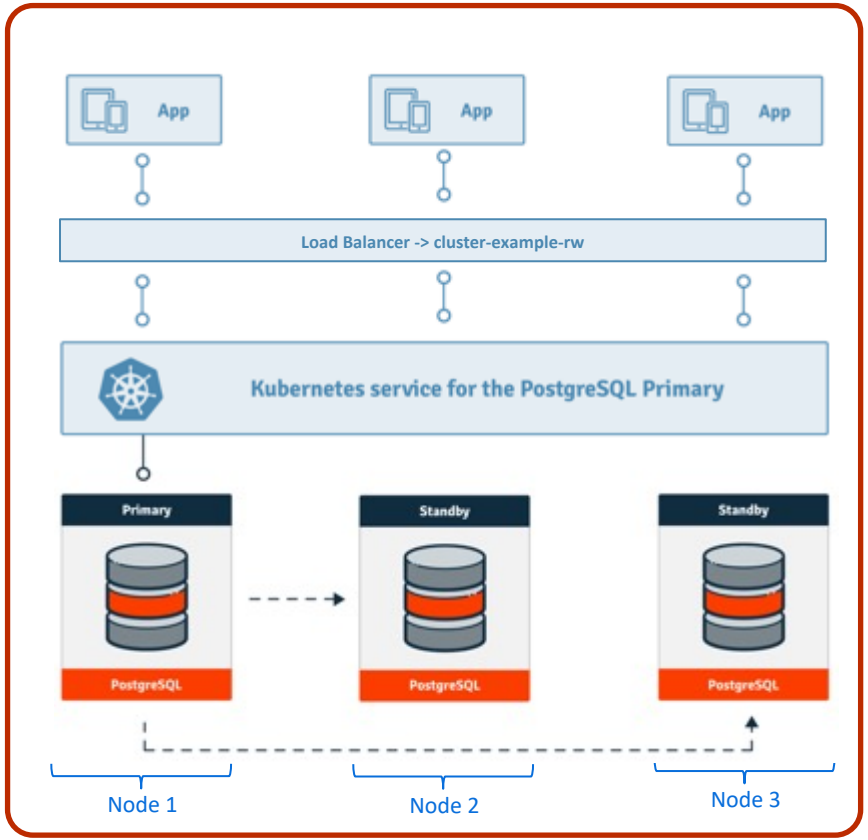

EDB

cluster-example.yaml

- Cluster name: cluster-example
- 3 Instances using replication slots
  - 1 Primary
  - 2 Standby's
- PostgreSQL 14.2
- Min 1 sync replica
- Activate pg_stat_statement extension
- 1GB disk
- Activate monitoring metrics
- CPU
  - Request: 1
  - Limit: 2

```yaml
apiVersion: postgresql.cnpg.io/v1
kind: Cluster
metadata:
  name: cluster-example
spec:
  instances: 3
  imageName: ghcr.io/cloudnative-pg/postgresql:14.2

  # Reduce the frequency of standby HA slots updates to once every 5 minutes
  replicationSlots:
    highAvailability:
      enabled: true

  minSyncReplicas: 1
  maxSyncReplicas: 1

  postgresql:
    parameters:
      pg_stat_statements.max: "10000"
      pg_stat_statements.track: all

  storage:
    size: 1Gi

  monitoring:
    enablePodMonitor: true

  resources:
    requests:
      memory: "512Mi"
      cpu: "1"
    limits:
      memory: "1Gi"
      cpu: "2"
```

EDB

Monitoring

- For each PostgreSQL instance, the operator provides an exporter of metrics for Prometheus via HTTP, on port 9187, named metrics.
- The operator comes with a predefined set of metrics, as well as a highly configurable and customizable system to define additional queries via one or more ConfigMap or Secret resources

Thank you