



EDB
Postgres® for the AI Generation

CloudNativePG 101



Gabriele Quaresima,
12/05/2024



Gabriele Quaresima Staff SDE @ EDB

<3 OSS

<3 Coding

<3 Playing sax

<3 Paola and Maria



Kubernetes (K8s)



Kubernetes

- **Orchestrator** (runs and manages containerized applications on a cluster)
- Higher-level abstractions of workloads
- **Declarative** over imperative.
Via controller pattern, Kubernetes aligns the current state (of any applied resource) to the desired one



Kubernetes - Sample object

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: nginx-deployment
spec:
  selector:
    matchLabels:
      app: nginx
  replicas: 2 # tells deployment to run 2 pods matching the template
  template:
    metadata:
      labels:
        app: nginx
    spec:
      containers:
      - name: nginx
        image: nginx:1.14.2
        ports:
        - containerPort: 80
```



Cloud Native Applications

- Small loosely coupled services (microservices)
- Introduce a set of principles and a culture shift paradigm for **devops** and **CI/CD** approach





Kubernetes Operators

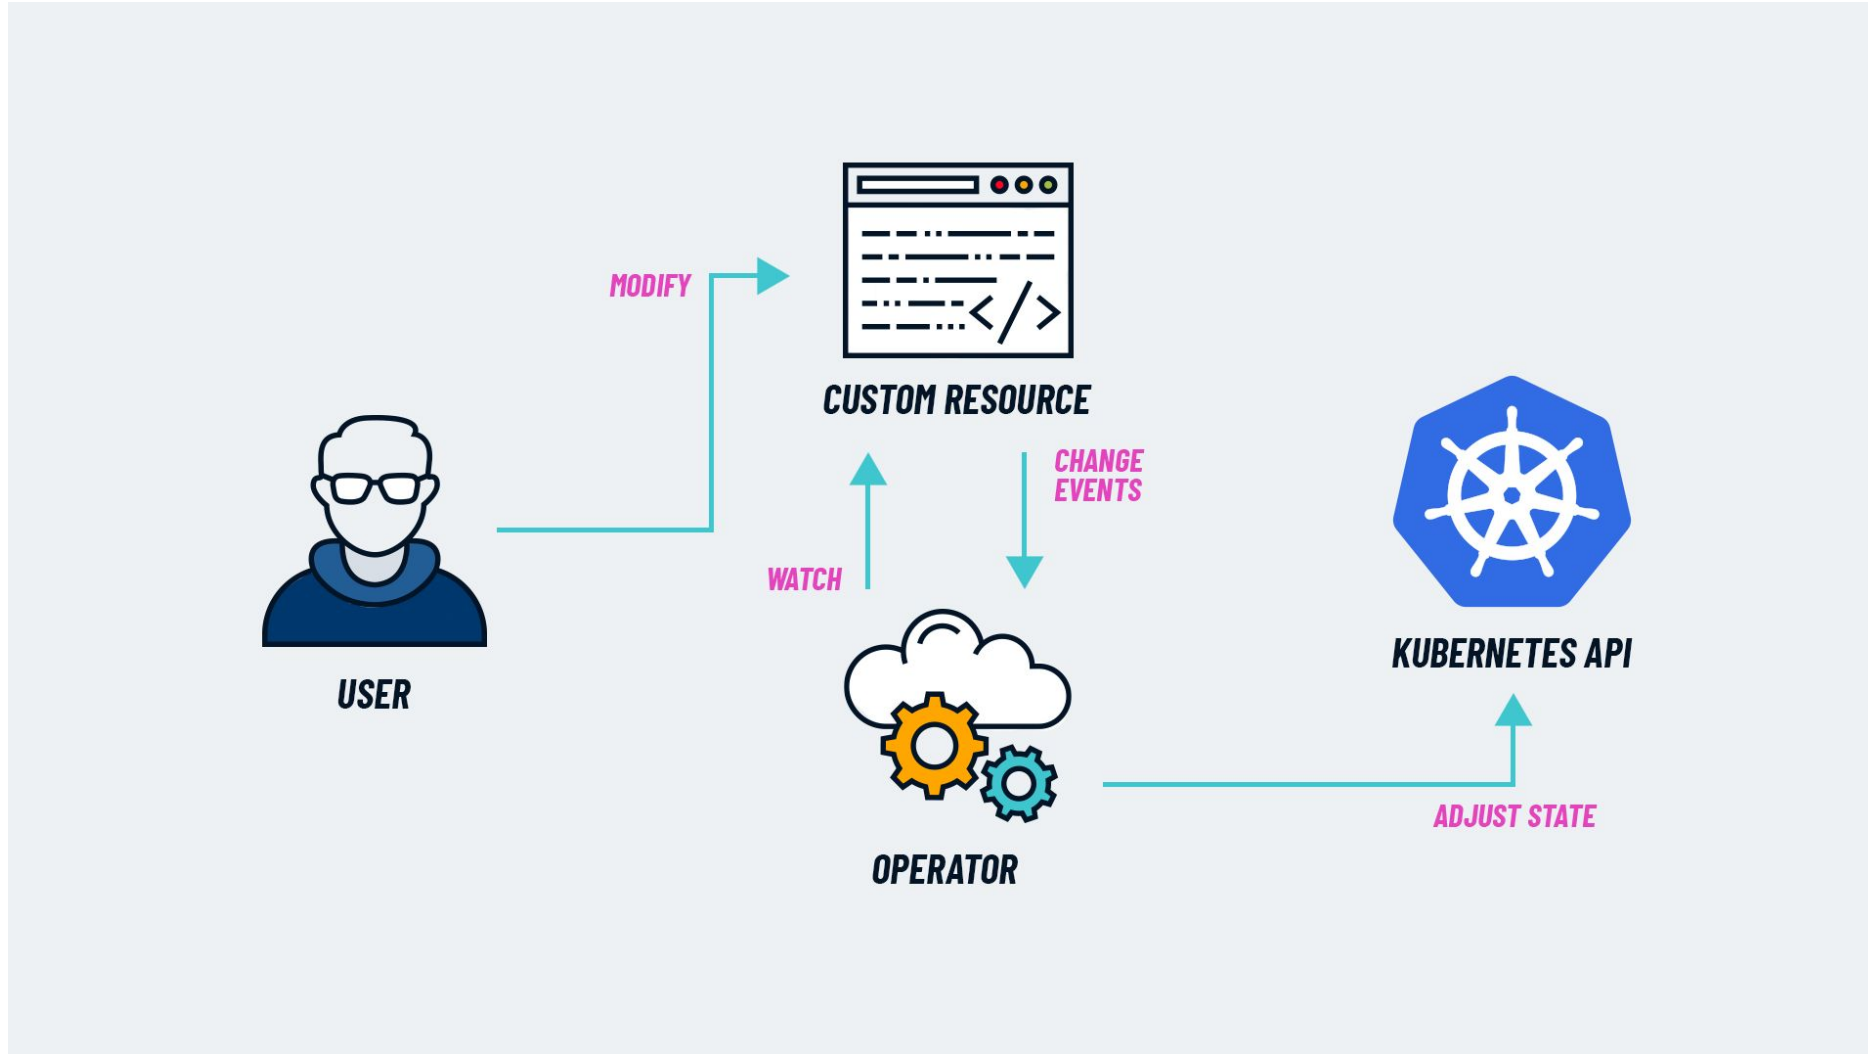


Kubernetes Operator

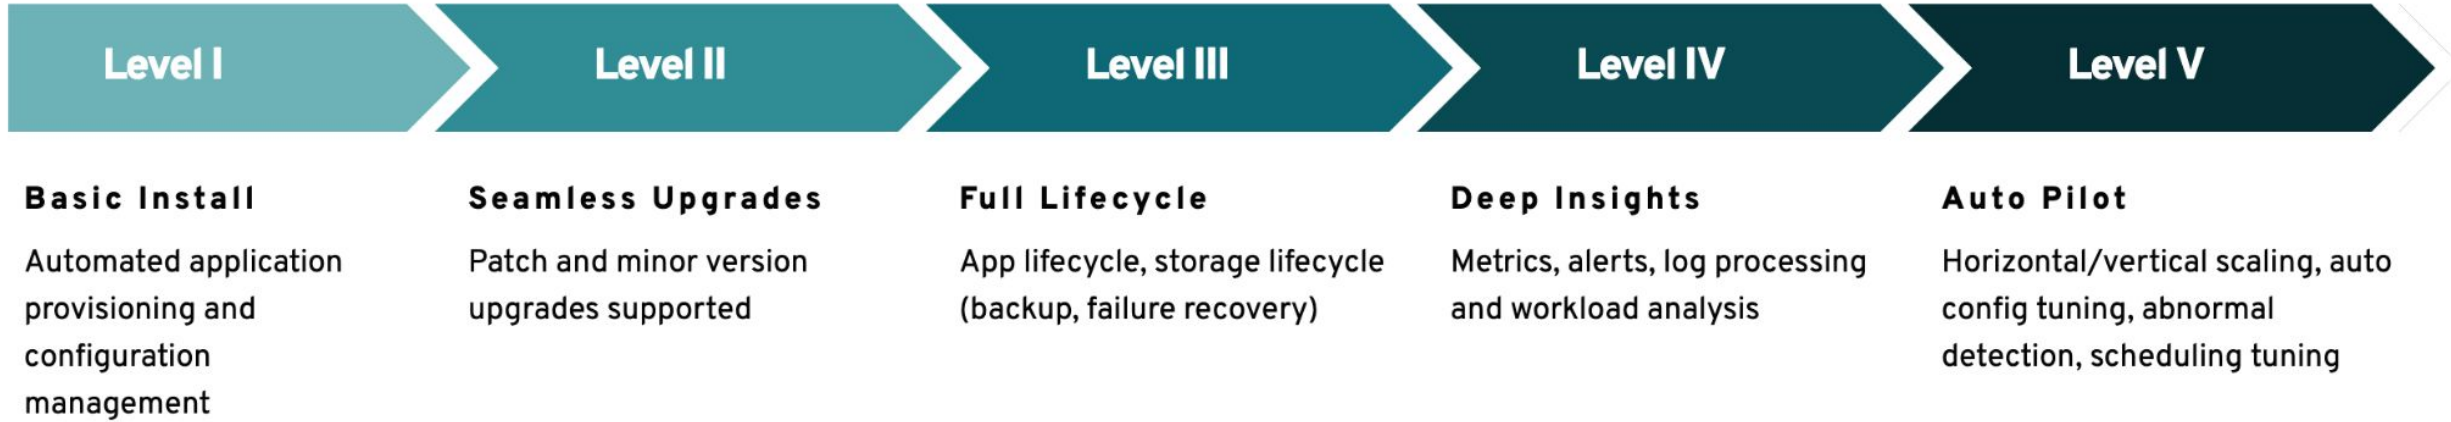
- Software extensions to Kubernetes
- Use of **custom resources** to manage applications and their components
- Operators follow Kubernetes principles, notably the control loop
- Operators implement and automate common Day-1 (installation, configuration, etc) and Day-2 (re-configuration, update, backup, failover, restore, etc.) activities
- Via operators, you can stop treating an application as a collection of primitives like Pods, Deployments, Services or ConfigMaps, but instead as a single object that only exposes the knobs that make sense for the application



Kubernetes Operator



Kubernetes Operator - Capability levels



Kubernetes Operator - List of existing operators

Artifact **HUB** DOCS STATS SIGN UP SIGN IN

1 - 20 of 15320 results Show: 20

FILTERS

- Official
- Verified publishers
- CNCF

KIND

- Argo templates (11)
- Containers images (188)
- Falco rules (23)
- Gatekeeper policies (45)
- Headlamp plugins (8)
- Helm charts (12825)
- Helm plugins (43)
- Inspektor gadgets (33)
- KCL modules (329)
- KEDA scalers (4)
- Keptn integrations (41)
- Knative client plugins (5)
- Krew kubectl plugins (284)



kube-prometheus-stack

Prometheus prometheus-community

★ 892 Helm chart

Updated 2 hours ago

Version 65.4.0

kube-prometheus-stack collects Kubernetes manifests, Grafana dashboards, and Prometheus rules combined with documentation and scripts to ...

Monitoring and logging



cert-manager

cert-manager cert-manager

★ 696 Helm chart

Updated 15 days ago

Version 1.16.1

A Helm chart for cert-manager

Security



ingress-nginx

Kubernetes ingress-nginx

★ 675 Helm chart

Updated 9 days ago

Version 4.12.0-beta.0

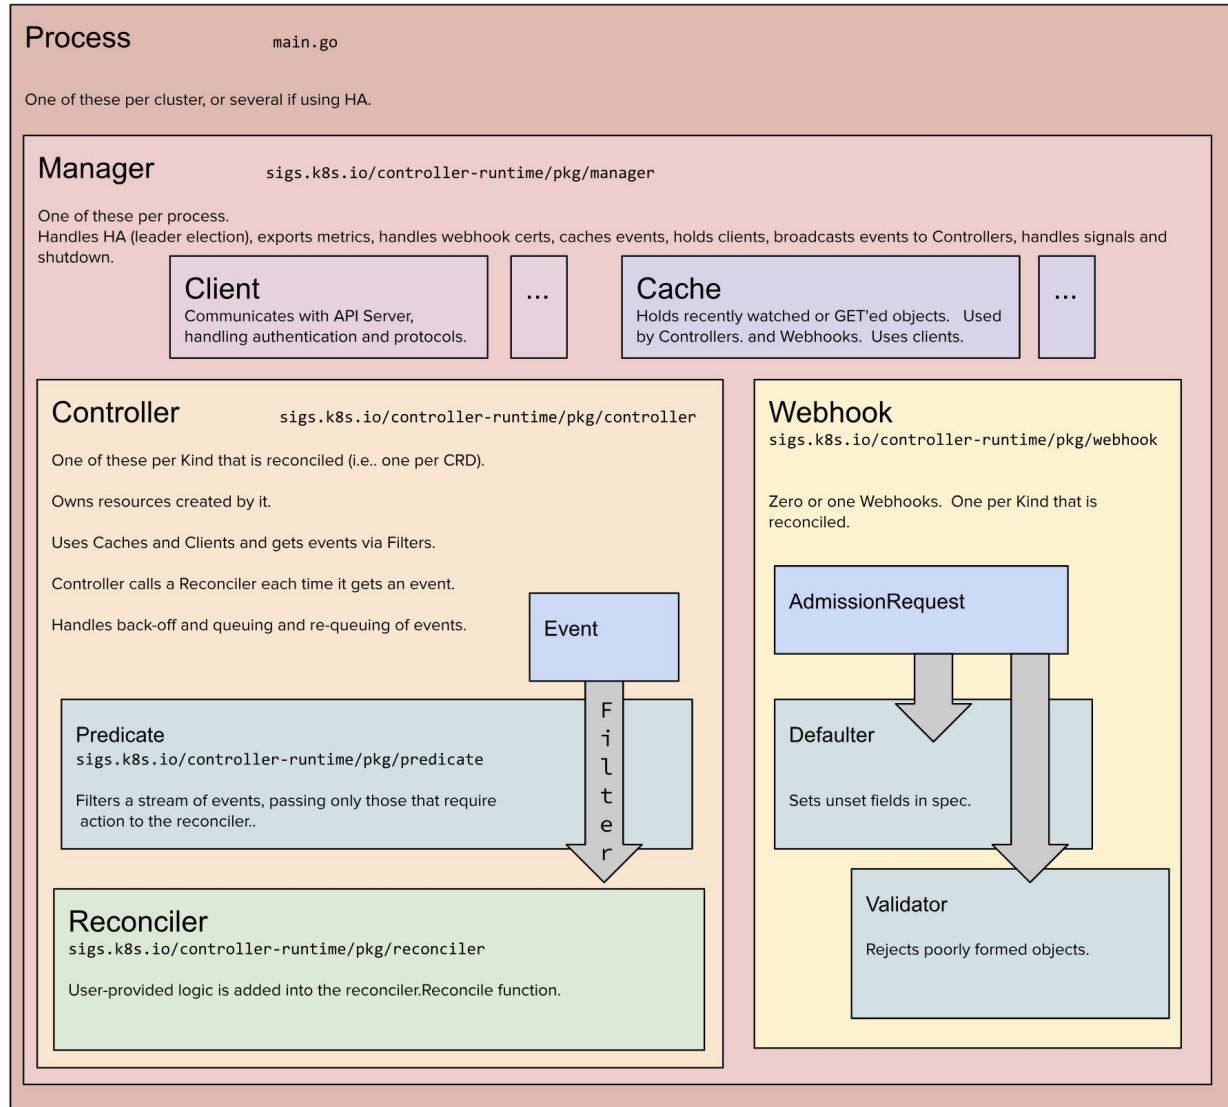


Kubernetes Operator - Write your own one

- **SDK** for the major languages
- Custom kubernetes **controllers**
- Design declarative **APIs** for operators
- Usually interact/control a custom object (Custom Resource Definitions (**CRD**))



Kubernetes Operator - Kubebuilder Architecture



Link to Github repo of **KubeBuilder**

<https://github.com/kubernetes-sigs/kubebuilder>






Kubernetes Operators for PostgreSQL



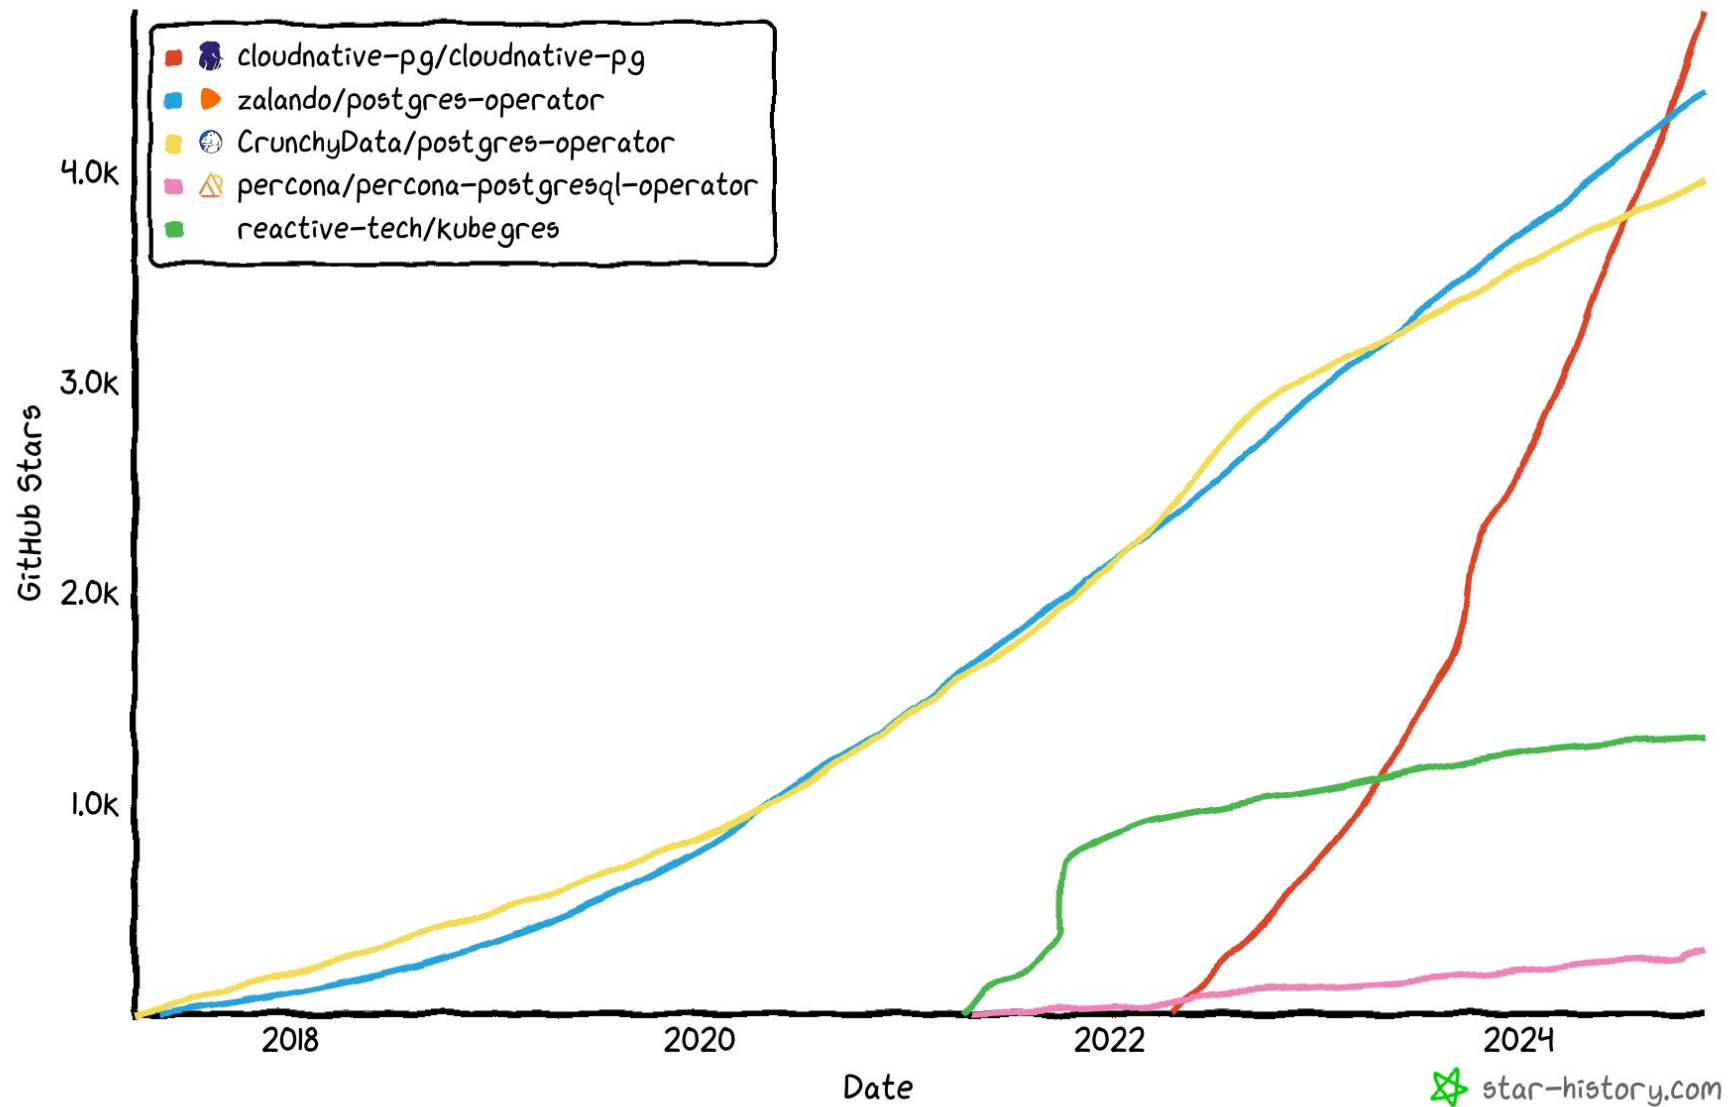
Bring Kubernetes to Postgres, bring Postgres to Kubernetes!



kubernetes



Star History



star-history.com






CloudNativePG



CloudNativePG

- **CloudNativePG** is an open-source operator designed to manage **PostgreSQL** workloads on any supported **Kubernetes** cluster
- It defines a new K8s resource called **Cluster**, representing a PostgreSQL cluster made up of a single primary and an optional number of replicas



CloudNativePG



CloudNativePG - Cluster CRD example

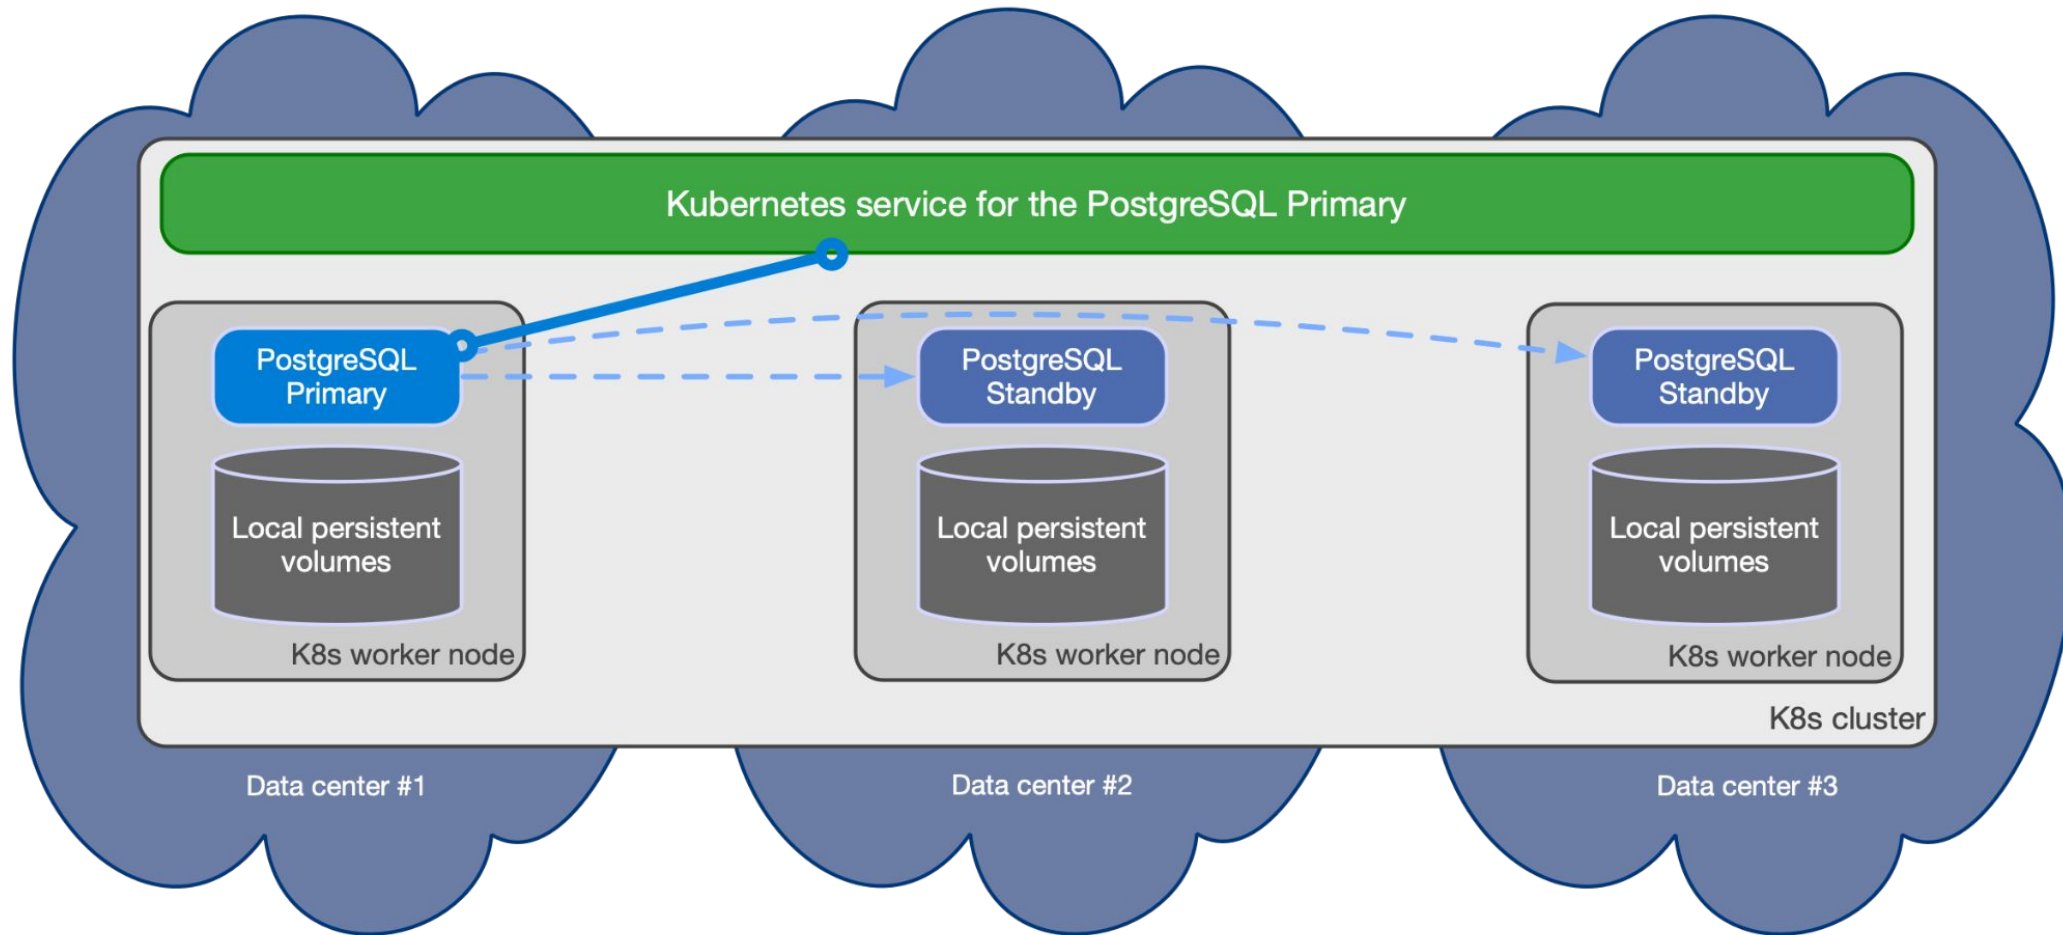
```
> batcat cloudnative-pg/docs/src/samples/cluster-example.yaml
```

```
File: cloudnative-pg/docs/src/samples/cluster-example.yaml
```

```
1  apiVersion: postgresql.cnpg.io/v1
2  kind: Cluster
3  metadata:
4    name: cluster-example
5  spec:
6    instances: 3
7
8    storage:
9      size: 1Gi
```



CloudNativePG - Suggested architecture



CloudNativePG - PVC management

- **No Statefulsets**
- **The operator manages directly PVCs:**
 - It creates a PVC for each PostgreSQL instance, with the goal of storing the PGDATA
 - It then mounts it into each pod
- **CSI** for PVCs allows
 - **Volume Snapshot** backup and restore
 - Online resize
 - PVC cloning



CloudNativePG - Main features

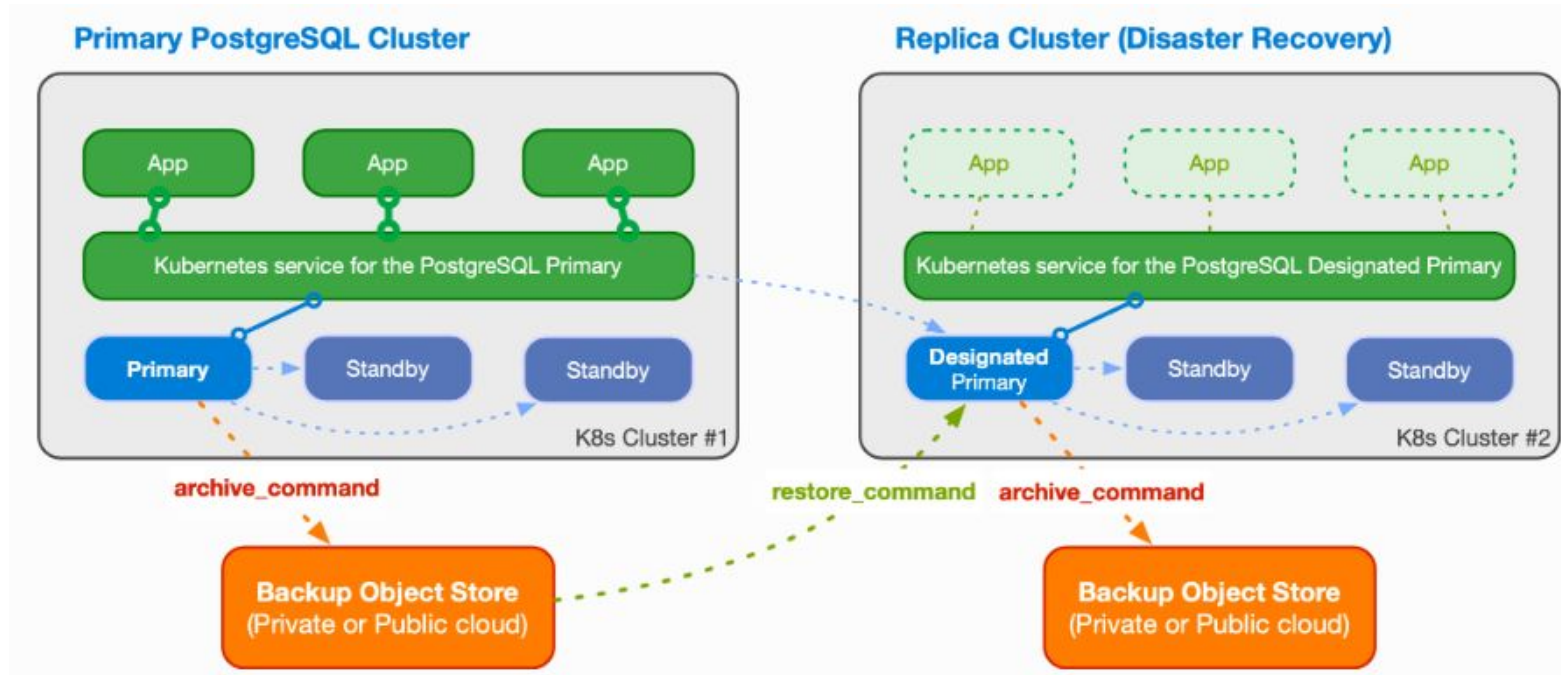
- Direct integration with K8s Api Server for self healing
- High Availability capability
- Planned switchover
- Scale up/down capabilities
- Definition of an arbitrary number of instances, rw and ro services
- Declarative management of PostgreSQL configuration
- Support for local PV with PVC templates
- Backup on VolumeSnapshot and object stores
- Full recovery or PITR from an existing backup
- kubectl plugin for cluster management
- Declarative tablespaces and databases
- WIP: CNPG-I: define your own plugin!



CloudNativePG - Deployment across Kubernetes clusters

Distributed PostgreSQL topology with **replica clusters** enables multi-cluster deployments in private, public, hybrid, and multi-cloud contexts.

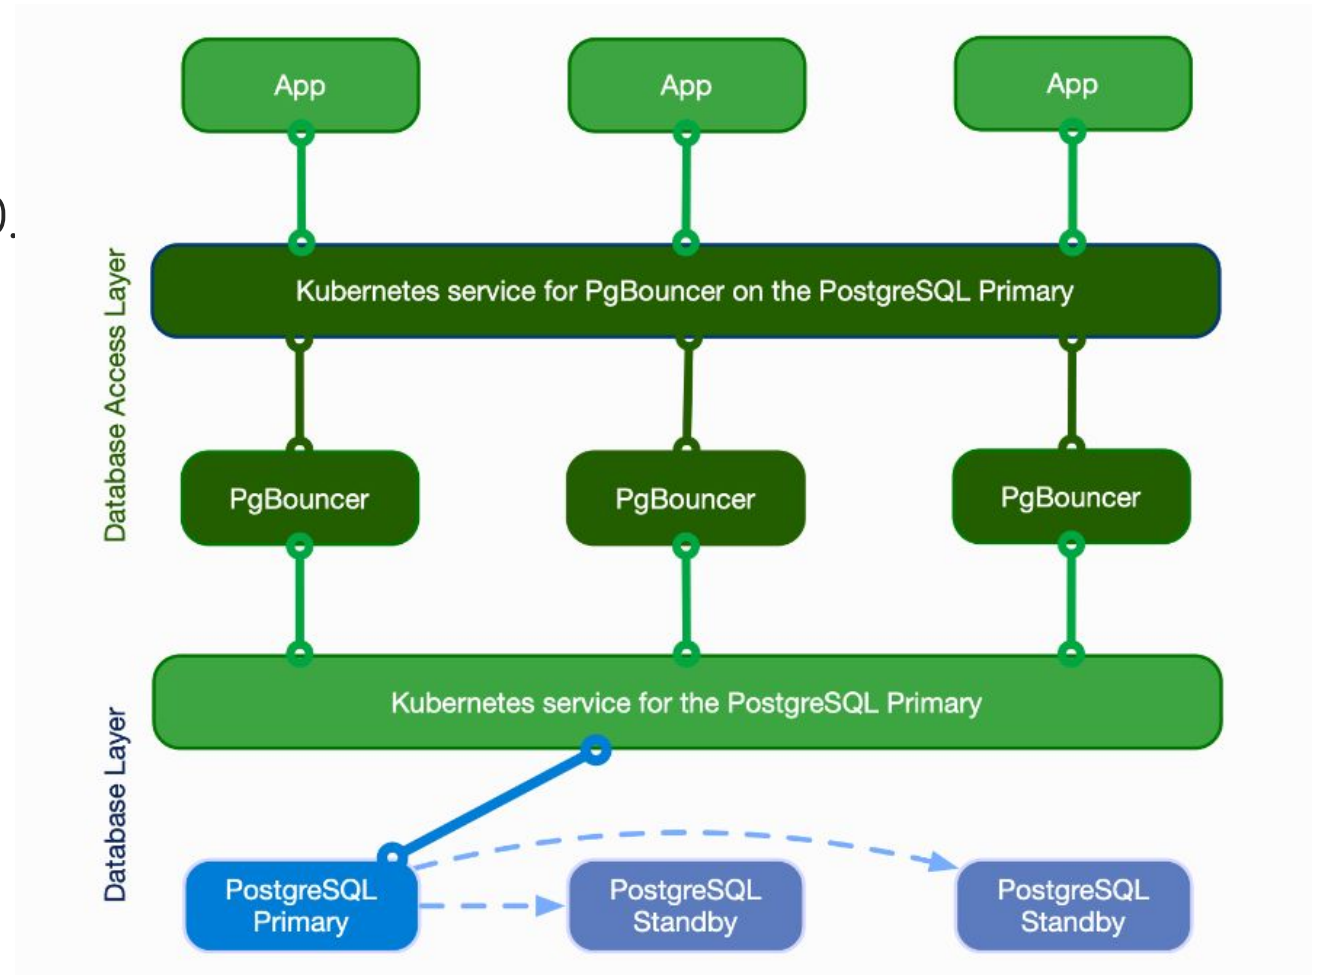
A replica cluster is a separate **Cluster** resource that is in continuous recovery, replicating from another source, either via WAL or streaming replication.



CloudNativePG - Connection pooling

CloudNativePG provides native support for connection pooling with PGBouncer, through the **Pooler** CRD.

A pooler is a deployment of PgBouncer pods that sits between your application and a PostgreSQL service. It creates a separate, scalable, configurable, and highly available database access layer.



CloudNativePG - Installation



```
kubectl apply --server-side -f \  
https://raw.githubusercontent.com/cloudnative-pg/cloudnative-pg/main/releases/cnpg-1.24.1.yaml
```

Refer to the CloudNativePG documentation for Helm Chart and Operator Lifecycle Manager (OLM), directly from OperatorHub.io



CloudNativePG - Operand images

- The PostgreSQL operand **container images** are available for all PGDG supported versions of PostgreSQL, across multiple architectures, directly from the postgres-containers project's GitHub Container Registry.
- Daily jobs ensure that CVEs in the entire stack are promptly addressed
- Additionally, the community provides images for the PostGIS extension
- Check this blog article, if you want create your own **custom** container image.
At the end, Aln't you really want to add pgvector into your PG image? :)



CloudNativePG - Observability

- CloudnativePG exposes native **Prometheus** exporters. Some are already defined, but you can extend them and add new ones
- All logs are in stdout and in **JSON** format
- CNPG plugin **status** command



CloudNativePG - Kubectl Plugin

CloudnativePG provides a **plugin** for kubectl to manage a cluster in Kubernetes

```
> ./kubectl-cnpg --help
A plugin to manage your CloudNativePG clusters

Usage:
  kubectl cnpg [command]

Operator-level administration
  install      CloudNativePG installation-related commands

Troubleshooting
  logs        Logging utilities
  report      Report on the operator or a cluster for troubleshooting

Cluster administration
  destroy     Destroy the instance named [cluster]-[node] or [node] with the associated PVC
  fencing     Fencing related commands
  hibernate   Hibernation related commands
  maintenance Sets or removes maintenance mode from clusters
  promote     Promote the pod named [cluster]-[node] or [node] to primary
  reload      Reload a cluster
  restart     Restart a cluster or a single instance in a cluster

Database administration
  backup      Request an on-demand backup for a PostgreSQL Cluster
  certificate  Create a client certificate to connect to PostgreSQL using TLS and Certificate authentication
  publication  Logical publication management commands
  snapshot    DEPRECATED (use `backup -m volumeSnapshot` instead)
  status      Get the status of a PostgreSQL cluster
  subscription Logical subscription management commands

Miscellaneous
  fio         Creates a fio deployment, pvc and configmap
  pgadmin4    Creates a pgAdmin deployment
  pgbench     Creates a pgbench job
  psql        Start a psql session targeting a CloudNativePG cluster

Additional Commands:
  completion  Generate the autocompletion script for the specified shell
  help        Help about any command
  version     Prints version, commit sha and date of the build
```



CloudNativePG - Product

- CloudNativePG is **production ready!**
- We proudly reached around 4.8K Github stars in 2.5 years, 3.5K commits, ~60M downloads
- CloudNativePG became the de facto standard to deploy PostgreSQL in Kubernetes
- Community is very active inside the Slack workspace
- We recently applied for the **CNCF Sandbox**



CloudNativePG - License

CloudNativePG has been originally created **EDB**.

CloudNativePG is 100% **open source** and community-driven.

All components are available under the [Apache 2 license](#) on Github.



CloudNativePG - Adopters

- **Google Cloud**

<https://cloud.google.com/kubernetes-engine/docs/tutorials/stateful-workloads/cloudnativepg>

- **Microsoft Azure**

<https://learn.microsoft.com/en-us/azure/aks/postgresql-ha-overview>

- **IBM**

- **EDB**

EDB CloudNativePG Cluster, EDB CloudNativePG Global Cluster, EDB Postgres AI Cloud Service (DBaaS)

Full list [here](#)



CloudNativePG - Beyond the operator

CloudNativePG project contains a plenty of interesting repositories:

- Helm charts
- CNPG-I
- Grafana dashboards
- PostGis Containers
- Ciclops



CloudNativePG - Resources

- CloudNativePG [website](#) and [GitHub](#) project
- CloudNativePG Recipes on Gabriele Bartolini's [blog](#)
- CloudNativePG maintainers and contributors are active on [Data On Kubernetes](#) community
- Latest KubeCon talks mentions CloudNativePG, eg.
 - [“We tested and compared 6 Database Operators. The results are in!”](#)
 - [“Kubernetes on Multisites - A Story About Stateful App, Hybrid Clouds and High Availability”](#)
- CNCF blog article [“Cloud Neutral Postgres Databases with Kubernetes and CloudNativePG”](#)



Join the
Community





Demo



Q&A



Thanks!

