

# POSTGRES ON OPENSIFT WORKSHOP



David Martini - Senior Solution Architect, Red Hat

Lucie Zeng - Sales Engineer, EDB

Raphaël Chir - Senior Sales Engineer, EDB

Eric Pillon - Strategic Account Executive, EDB

24th February, 2026



# Agenda

Commencer	Fin	Session
12:30	13:30	Accueil Cocktail Lunch
13:30	13:45	Partenariat entre Red Hat OpenShift et EDB (Red Hat - David Martini)
13:45	14:00	Présentation du marché Postgres et d'EDB (EDB - Eric Pillon)
14:00	14:30	Architecture de référence de l'opérateur CNPG (EDB - Raphaël Chir)
14:30	16:30	Atelier pratique (EDB - Lucie Zeng & Raphaël Chir)
16:30	17:00	Pause
17:00	17:30	Conclusion et perspectives



# Red Hat OpenShift with EDB

David Martini  
Senior Solution Architect

# The world's leading provider of open source enterprise IT solutions

More than  
**90%**  
of the  
Fortune  
**500**  
use  
**Red Hat**  
products and  
solutions<sup>1</sup>

**~22,000**  
employees

**105+**  
offices

The first  
**\$3**  
**billion**  
open  
source  
company  
in the world<sup>2</sup>

Sources: 1-Red Hat client data and [Fortune 500 list](#), September 2024.

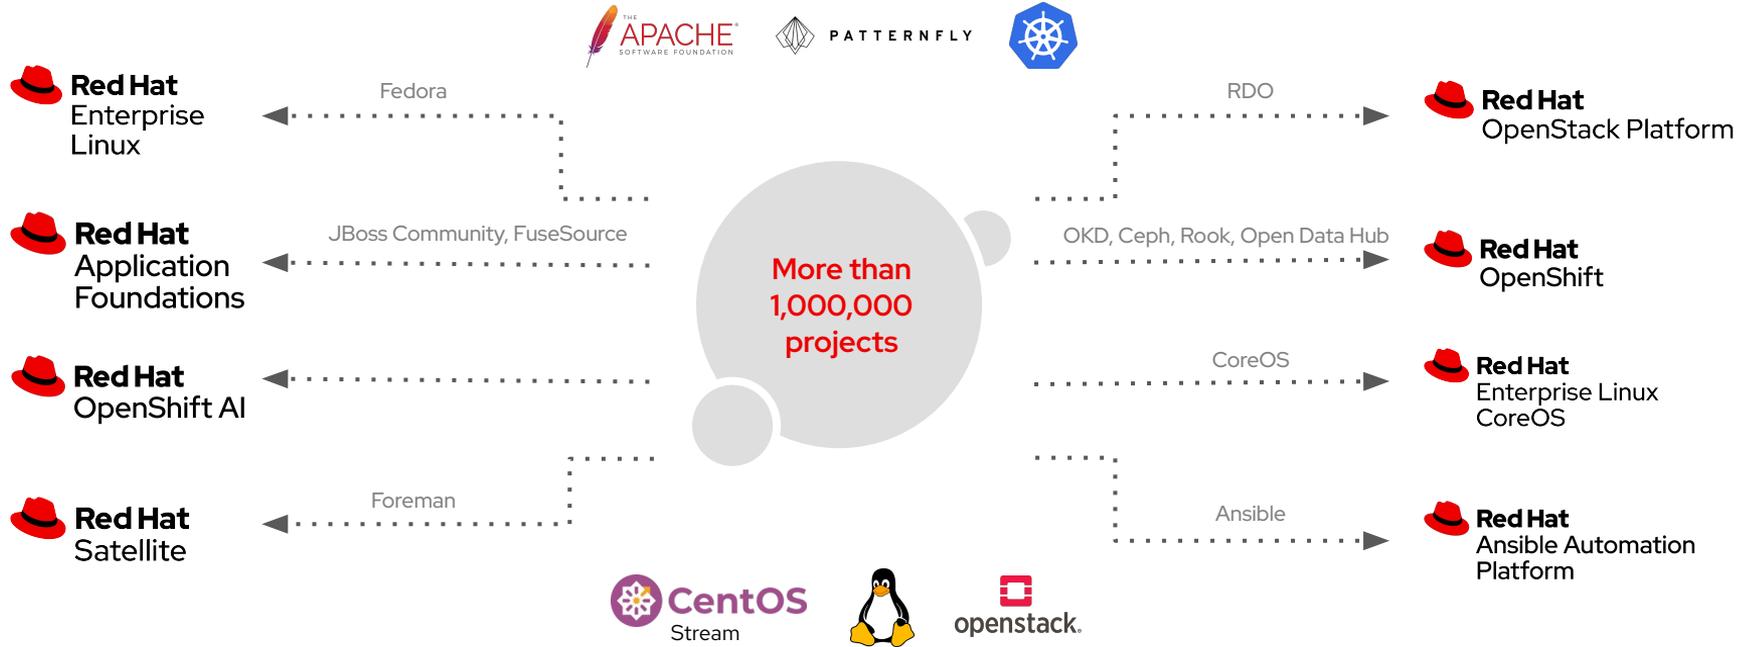
2 - [Red Hat SEC filings](#) prior to the acquisition by IBM. Note: Currency in U.S. dollars.

## Open source

Open source is about more than developing software.  
It's how we built our company.

And it's why we have been so successful.

# From communities to enterprise



## Linux contributions



 Platinum members

 Red Hat has been working on Linux since 1993

 Linux is the foundation of all Red Hat products

 The first Patent promise was issued by Red Hat in 2002

## OpenInfra contributions



 Historical member

 First Red Hat OpenStack release in 2013

 Most represented member on the board (3/21)

 Red Hat & IBM = more than 3 times contributions as the second contributor

## Cloud-native contributions



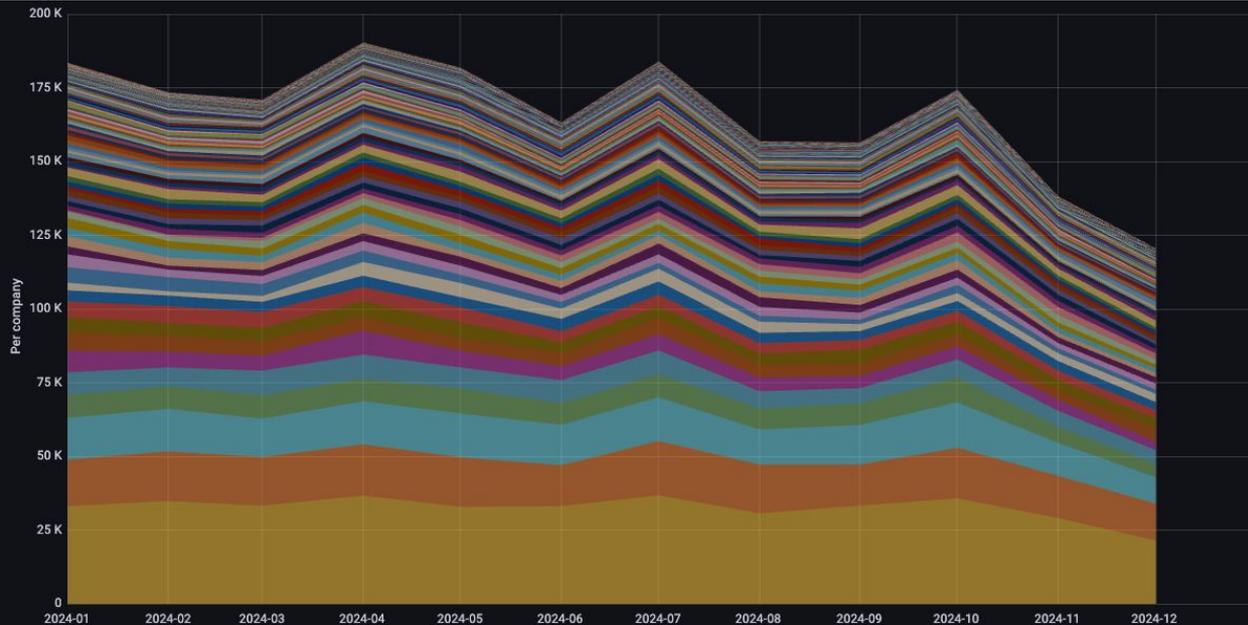
 Co-founder in 2015

 Platinum members

 Second Kubernetes contributor after Google

 First CNCF contributor in 2022, 2023, 2024, etc...

# CNCF 2024 Global contribution



	min	max	avg	current	total
Red Hat Inc.	21.49 K	36.82 K	32.59 K	21.49 K	391.11 K
Microsoft Corporation	12.51 K	18.41 K	15.83 K	12.51 K	189.98 K
Google LLC	8.94 K	15.43 K	13.37 K	8.94 K	160.43 K
Independent	4.51 K	8.49 K	7.29 K	4.51 K	87.50 K
Isovalent	4.74 K	8.30 K	6.74 K	4.74 K	80.87 K
VMware Inc.	2.82 K	8.13 K	5.13 K	2.82 K	61.59 K
SUSE LLC	2.95 K	6.06 K	4.54 K	4.54 K	54.52 K
Splunk Inc.	3.40 K	5.53 K	4.50 K	3.40 K	54.02 K
International Business Machines Corporation	2.79 K	5.65 K	4.12 K	2.79 K	49.47 K
Amazon	2.69 K	4.52 K	3.56 K	2.69 K	42.78 K
Solo.io	1.41 K	4.79 K	3.12 K	2.80 K	37.45 K
Dynatrace LLC	1.43 K	5.26 K	2.88 K	1.43 K	34.51 K
DaoCloud Network Technology Co. Ltd.	1.81 K	4.37 K	2.75 K	2.11 K	32.97 K
Huawei Technologies Co. Ltd	1.33 K	3.67 K	2.53 K	2.08 K	30.40 K
Kong Inc.	1.48 K	3.49 K	2.51 K	1.50 K	30.10 K
Spotify AB	1.53 K	3.58 K	2.40 K	1.53 K	28.82 K
CNCF	1.31 K	3.40 K	2.28 K	1.31 K	27.41 K
Ericsson	1.67 K	2.78 K	2.10 K	1.73 K	25.19 K
Elasticsearch Inc.	498.00	3.10 K	1.90 K	2.03 K	22.79 K
Grafana Labs	1.42 K	2.30 K	1.89 K	1.42 K	22.66 K
Tetrate.io	1.15 K	2.45 K	1.88 K	1.59 K	22.57 K
Cisco	1.29 K	2.25 K	1.82 K	1.29 K	21.87 K

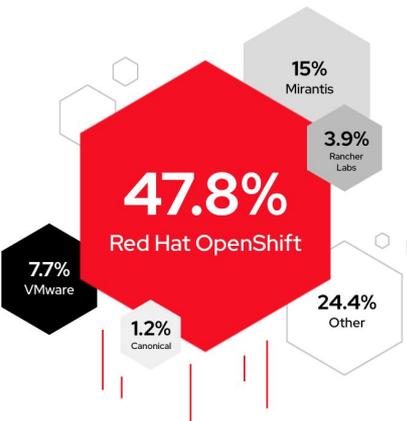
Red Hat is the **#1** company for contributions to **CNCF** in **2024**.

## Red Hat portfolio

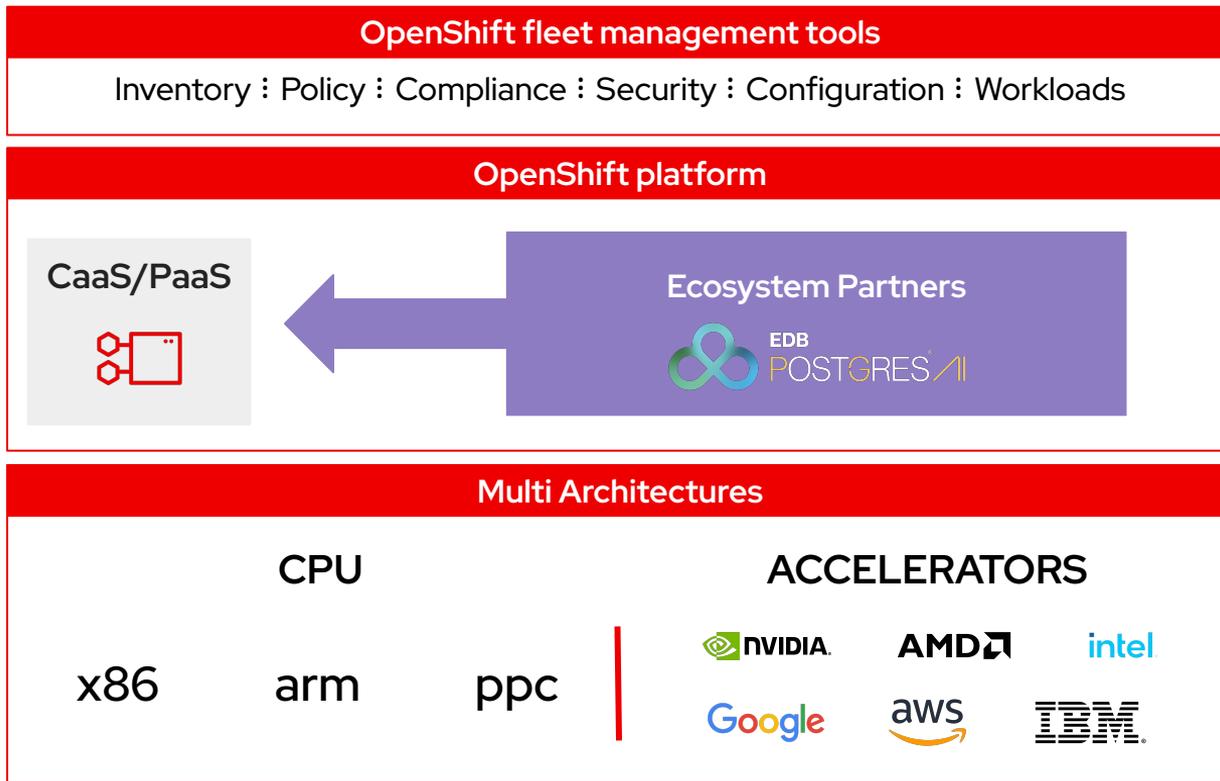
Red Hat's open source solutions work  
on bare metal and the public cloud ...

... and everything in between.

# A multi-faceted platform for strategic goals



**RED HAT OPENSIFT**  
Container platform market  
share leader



Physical

Virtual

Private cloud

Sovereign cloud

Public cloud

Edge



Red Hat  
OpenShift

# Red Hat OpenShift Application Platform

Trusted, comprehensive and consistent across hybrid cloud

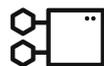
Application  
Build



Developer  
Tools



Runtimes &  
languages



Serverless &  
functions



Cloud  
Services  
APIs



Migration  
tools

Application  
Delivery



CI/CD &  
GitOps



Container registry



Observability



Logging



Virtual machines

Application  
Security



Service mesh



Vulnerability  
scanning



Security  
policies



Cert & secret  
tooling



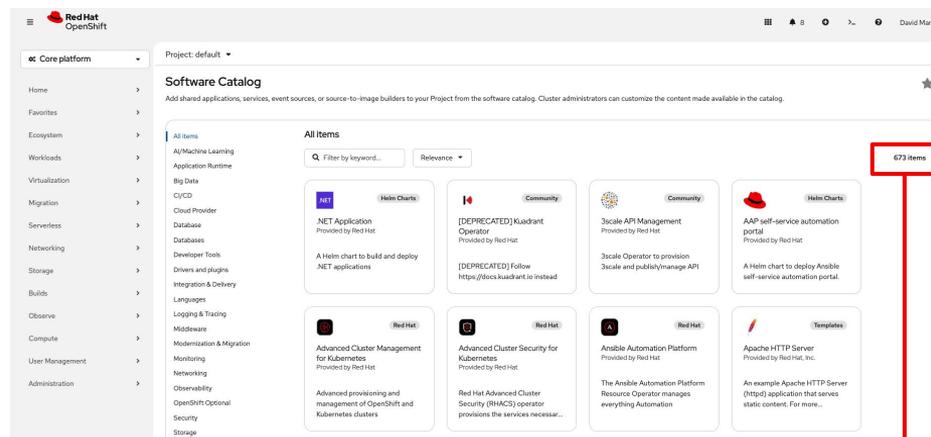
Auth and SSO  
services

# OpenShift is also content !

## Application runtimes

Included into OpenShift subscription (software collection)	PHP	Python	Java	NodeJS	Perl	Ruby	.NET Core	Third-party Language Runtimes
	MySQL	PostgreSQL	MongoDB	Redis	... virtually any available container image !			Third-party Databases
	Apache HTTP Server	Nginx	RH-SSO	JBoss Web Server	Tomcat			Third-party App Runtimes
Available through additional bundles	Spring Boot	Thorntail	Vert.x	JBoss EAP	JBoss AMQ Streams	JBoss AMQ	JBoss Fuse	Third-party Middleware
	3SCALE API mgmt	JBoss DAM	JBoss PAM	JBoss Data Virt	JBoss Data Grid	RH Mobile	JBoss Fuse Online	Third-party Middleware

## Marketplace



+600 items

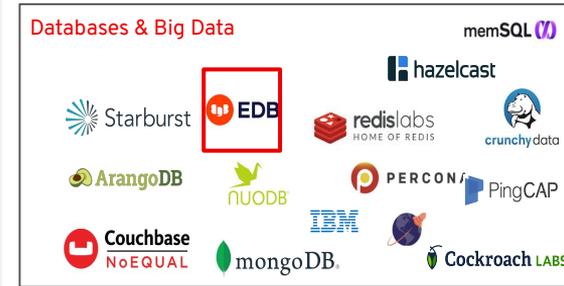
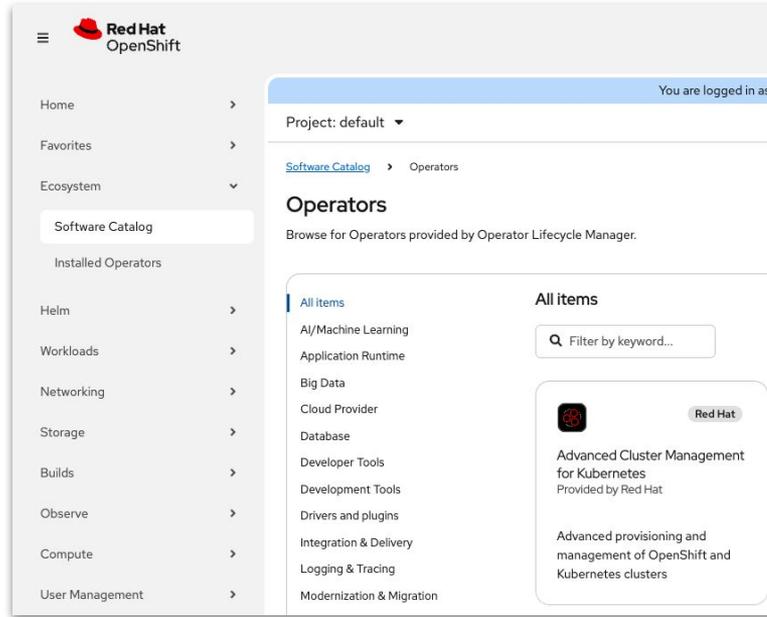
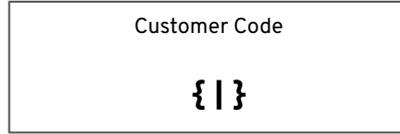
Health Index ⓘ



This image does not have any unapplied Critical or Important security updates.

The Container Health Index analysis is based on RPM packages signed and created by Red Hat, and does not grade other software that may be included in a container image.

# A large ecosystem



local-cluster8dmartini@redhat.com

Developer

+Add

Topology

Observe

Search

Functions

Builds

Pipelines

Environments

Helm

Project

ConfigMaps

Secrets

Project: dev-l2

### Developer Catalog

**All services**  
Browse the catalog to discover, deploy and connect to services

**Database**  
Browse the catalog to discover database services to add to your application

**Operator Backed**  
Browse the catalog to discover and deploy operator managed services

**Helm Chart**  
Browse the catalog to discover and install Helm Charts

**Virtual Machines**  
Create a Virtual Machine from a template

**Samples**  
Create an application from a code sample

### Eventing

**Event Type**  
Event Types available that can be consumed with Serverless Functions or regular Deployments

**Event Source**  
Create an Event source to register interest in a class of events from a particular system

**Event Sink**  
Create an Event sink to receive incoming events from a particular source

**Broker**  
Create a Broker to define an event mesh for collecting a pool of events and route those events based on attributes, through triggers

**Channel**  
Create a Knative Channel to create an event forwarding and persistence layer with in-memory and reliable implementations

OpenShift Self Managed Services

Red Hat OpenShift AI

### Serverless function

**Import from Git**  
Create and deploy stateless, Serverless functions

**Samples**  
Build from pre-built function sample code

**Sharing**  
Project access allows you to add or remove a user's access to the project

**Helm Chart repositories**  
Add a Helm Chart Repository to extend the Developer Catalog

**Pipelines**  
Create a Tekton Pipeline to automate delivery of your application

### Git Repository

**Import from Git**  
Import code from your Git repository to be built and deployed

**Container images**  
Deploy an existing Image from an Image registry or Image stream tag

**From Local Machine**

**Import YAML**  
Create resources from their YAML or JSON definitions

**Upload JAR file**  
Upload a JAR file from your local desktop to OpenShift

# Why Red Hat OpenShift for EDB: Operator Certification

The screenshot displays the Red Hat Ecosystem Catalog interface. At the top, the navigation bar includes the Red Hat logo, 'Red Hat Ecosystem Catalog', and links for 'Solutions', 'Products', 'Artifacts', and 'Partners'. A search bar is present with the text 'All Search Ecosystem Catalog'. The breadcrumb trail reads 'Home > Software > All software results > Containerized applications'. The main content area features the EDB Postgres for Kubernetes logo and a 'Certified' badge. Below this, the text reads 'PostgreSQL Operator for mission critical databases in Openshift Container Platform'. A navigation bar below the main content has tabs for 'Overview', 'Resources', 'Certifications' (which is active), 'Deploy & use', and 'FAQs'. Under the 'Certifications' tab, there is a section titled 'Certifications' with a link to 'Learn about Red Hat Certification and Partner Validation'. Below that is a section titled 'Certified components' with a search bar and a filter for 'Image type'. The search results show 'EDB Postgres for Kubernetes (formerly Cloud Native PostgreSQL Operator) Container Images' with a green 'A' icon, the image name 'enterprisedb/cloud-native-postgresql', and a description: 'Container images for EDB Postgres for Kubernetes (formerly Cloud Native PostgreSQL operator)'. Additional details include 'EnterpriseDB', 'Published 10 days ago', 'Container image', '1.26.0-rc2-ubi9', and 's390x'.

EDB Postgres for Kubernetes is a certified Level 5 Operator for Red Hat OpenShift

- ▶ This is designed to streamline Day 2 operations of PostgreSQL databases
- ▶ Enhanced Database Management
- ▶ Supports point-in-time recovery (PITR)
- ▶ Ensures robust data protection and recovery options
- ▶ Integration with business continuity solutions such as Red Hat OpenShift API for Data Protection (OADP) and Veeam Kasten, Trilio, Portworx Backup, IBM Fusion, and others

# EDB on OpenShift use cases

- ▶ Cloud-Native Database Deployment
- ▶ Database as a Service (DBaaS)
- ▶ High Availability and Disaster Recovery (HA & DR)
- ▶ DevOps and Continuous Integration/Continuous Deployment (CI/CD)
- ▶ Microservices and Application Modernization
- ▶ Move from VMWare to OpenShift
- ▶ Data Security and Compliance (using **TDE** and Advanced Security provided by EPAS)
- ▶ Hybrid and Multi-Cloud Deployments
- ▶ Multi-Tenant Applications (isolation)

# Euro Information

## Company profile

Euro-Information is the fintech company of the Crédit Mutuel group. Euro-Information manages the IT systems of 16 federations of Crédit Mutuel as well as those of CIC and of all the financial, insurance, property, consumer credit, private banking, financing, telephony and technological subsidiaries.



- Red Hat OpenShift
- EDB Postgres for Kubernetes
- PostgreSQL
- EPAS



- EDB considerably reduces IT costs associated with database maintenance.
- 280 cores: Enterprise Plan + Production Support

## Summary

Use Case

On prem DBaaS (in Production)

Workload

Transactional

Application Name

All internal Postgres applications

EDB Tools of Interest

PostgreSQL and EDB Postgres for Kubernetes

## Problem

- Fast database deployment
- Adopt a supported and secure Open Source platform
- Onprem DBaaS
- Align to in-house RDBMS standardization

## Solution

- Use Postgres capabilities to build and maintain local applications
- Use Red Hat OpenShift platform to accelerate the provisioning of databases and applications

## Results

- Applications running with PostgreSQL databases in a centralized environment
- Massive reduction of TCO of database service operations



# La Poste

## Company profile

La Poste is a postal service company in France, operating in Metropolitan France, the five French overseas departments and regions and the overseas collectivity of Saint Pierre and Miquelon. Under bilateral agreements, La Poste also has responsibility for mail services in Monaco through La Poste Monaco and in Andorra alongside the Spanish company Correos.



- Red Hat OpenShift
- EDB Postgres for Kubernetes
- PostgreSQL



- EDB considerably reduces IT costs associated with database maintenance.
- 12 Cores: Standard Plan + Premium Support

## Summary

Use Case

On prem DBaaS with HA and DR  
**(In Production)**

Workload

Transactional

Application Name

Portail XaaS

EDB Tools of Interest

PostgreSQL and EDB Postgres for Kubernetes

## Problem

- Provide a database HA solution for Ansible Automation Platform (AAP)
- Database must be in HA and DR

## Solution

- Use EDB Postgres for Kubernetes to provide a HA and DR solution for PostgreSQL databases
- Deploy in 2 OpenShift clusters our operator

## Results

- La Poste developer can use their internal 'La Post Service Portal' to provision more than 64 backends.
- Reduce risk deploying EDB solutions.

# Airbus

## Company profile

Airbus SE is a European aerospace corporation. The company's primary business is the design and manufacturing of commercial aircraft but it also has separate defence and space and helicopter divisions.



- Red Hat OpenShift
- EDB Postgres for Kubernetes
- EPAS
- TDE



- Improve database deployment speed
- Reduce DB support
- Cost reduction

## Summary

Use Case	On prem DBaaS (Production)
Workload	Transactional
Application Name	All VMWare PostgreSQL databases
EDB Tools of Interest	EDB Postgres Advanced Server with Oracle and TDE (optional)

## Problem

- Flexibility
- Cost reduction
- New managed service in OpenShift

## Solution

- EDB Postgres for Kubernetes with EPAS. Depending of the applications needs, EPAS and/or TDE will be activated

## Results

- POC done
- Decision is taken
- Number of cores not yet communicated

# Banque de France

## Company profile

The Banque de France is France's central bank. A two-hundred-year-old institution, privately-owned when it was founded on January 18, 1800 under the Consulate by General Bonaparte, it became state-owned on January 1, 1946 when it was nationalized by General de Gaulle.



- Red Hat OpenShift
- CloudNativePG
- PostgreSQL



- 100 cores
- Subscription plan:
  - Community360 plan + Production Support

## Summary

Use Case

OnPrem DBaaS (in production)

Workload

Transactional

Application Name

Multiple applications

EDB Tools of Interest

PostgreSQL, CloudNativePG

## Problem

- Fast database deployment
- Provide containerized Postgres DBaaS

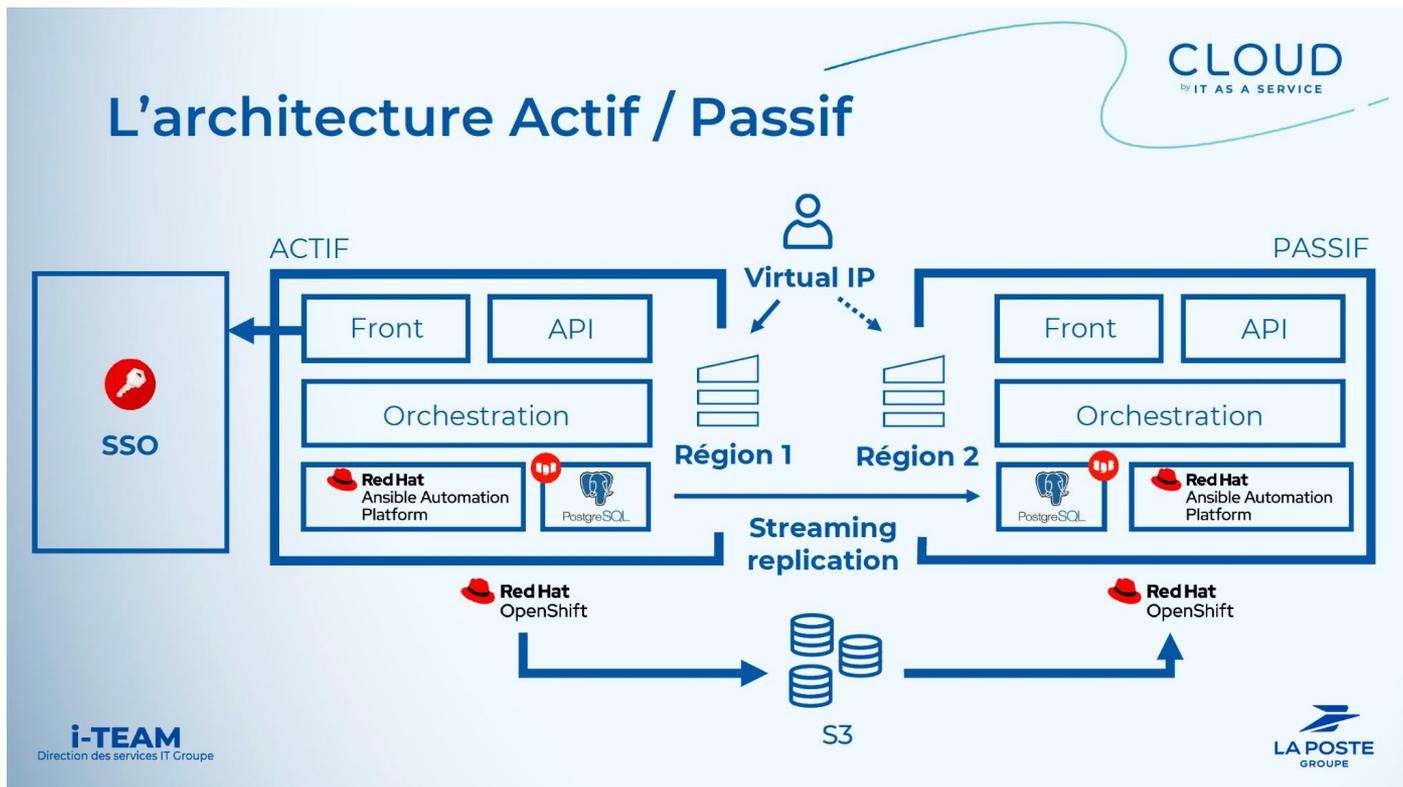
## Solution

- Use OpenShift to provide this service with the operator
- Fast deployment and with Open Source database

## Results

- OpenShift based PostgreSQL cluster deployments expand the internal offering alongside traditional VM based database cluster deployments

# La Poste Architecture



# Thank you

Red Hat is the world's leading provider of enterprise open source software solutions. Award-winning support, training, and consulting services make Red Hat a trusted adviser to the Fortune 500.



[linkedin.com/company/red-hat](https://www.linkedin.com/company/red-hat)



[youtube.com/user/RedHatVideos](https://www.youtube.com/user/RedHatVideos)



[facebook.com/redhatinc](https://www.facebook.com/redhatinc)



[twitter.com/RedHat](https://twitter.com/RedHat)

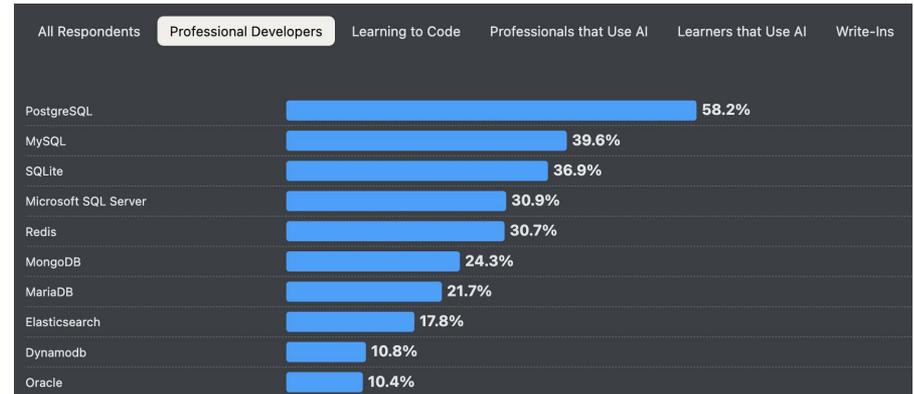
# Introduction to Postgres and EDB



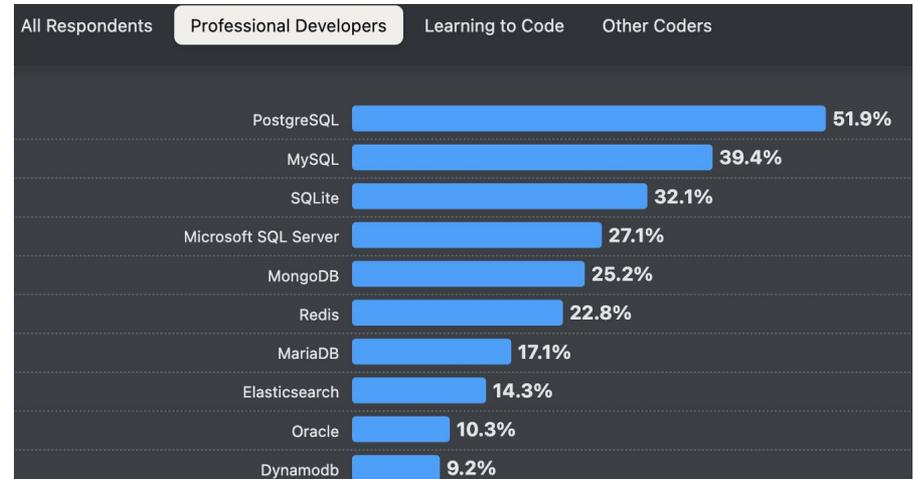
# PostgreSQL most loved database

Source: [Stackoverflow](#)

## 2025



## 2024



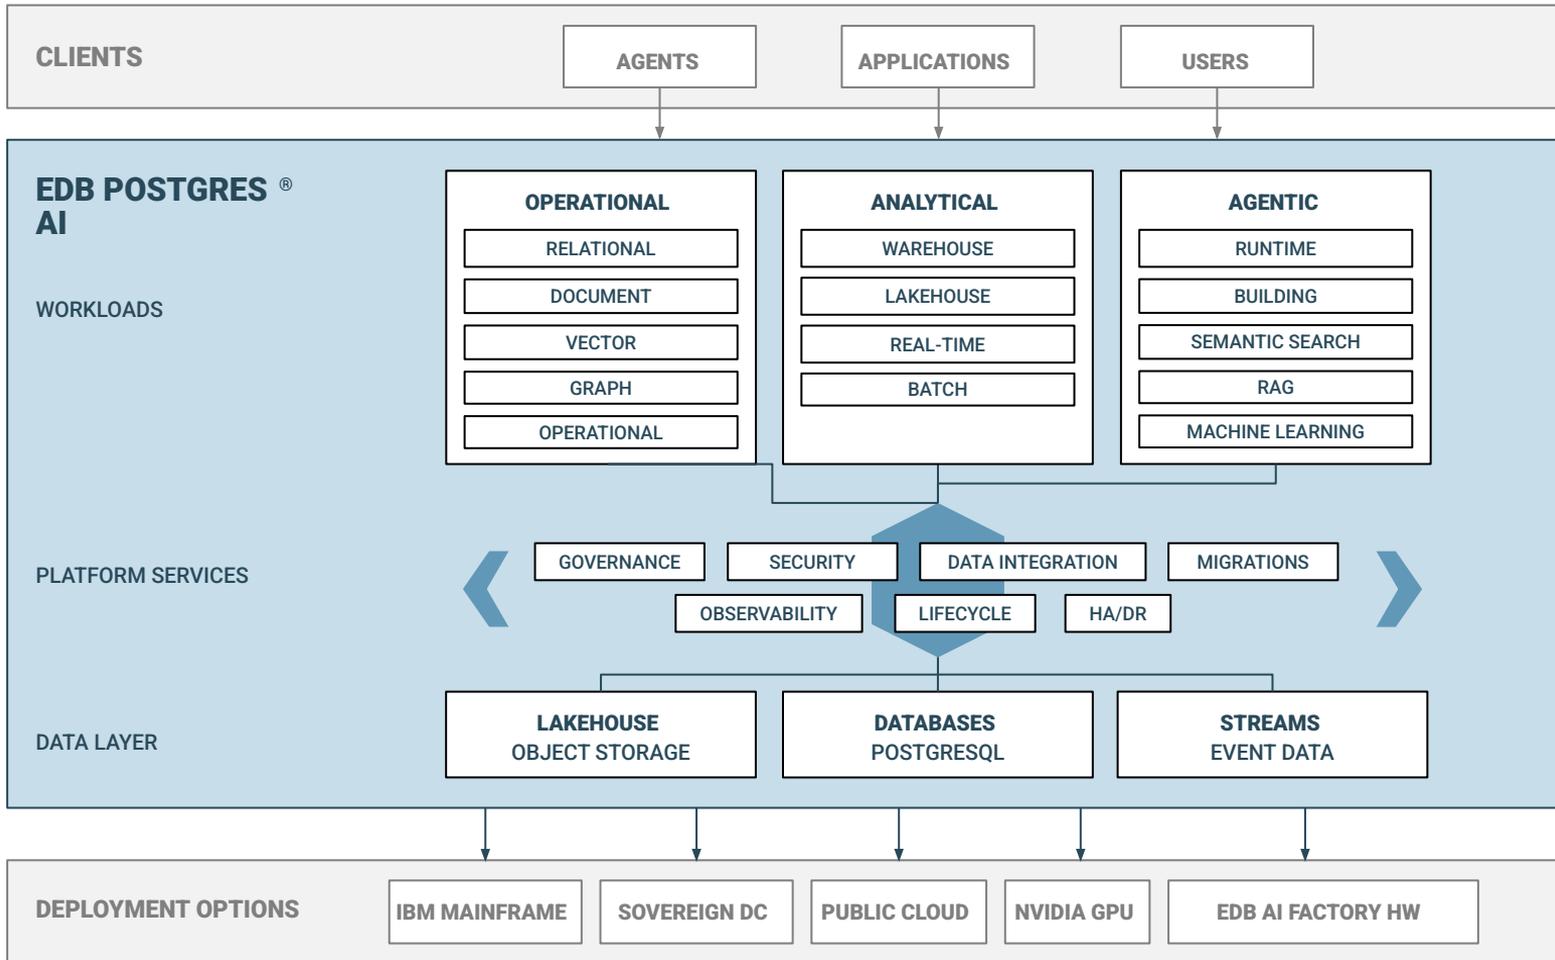


**EDB**  
Postgres® for the AI Generation

## 20+ years of Postgres innovation & adoption

- Number one contributor to Postgres, fastest-growing and most loved Database in the world
  - 2 Core Team members, 7 Committers, 9 Major Contributors, 10 Contributors, #1 site for desktop downloads
- Over 700 employees in more than 30 countries
- EDB Postgres AI
  - The industry's first platform that can be deployed as cloud, software or physical appliance
  - Secure, compliant and enterprise grade performance guaranteed





# CNPG Operator: Reference Architecture and functionalities



# 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
- 2025, January: CloudNativePG was recognized as an official **#CNCF** project



# 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

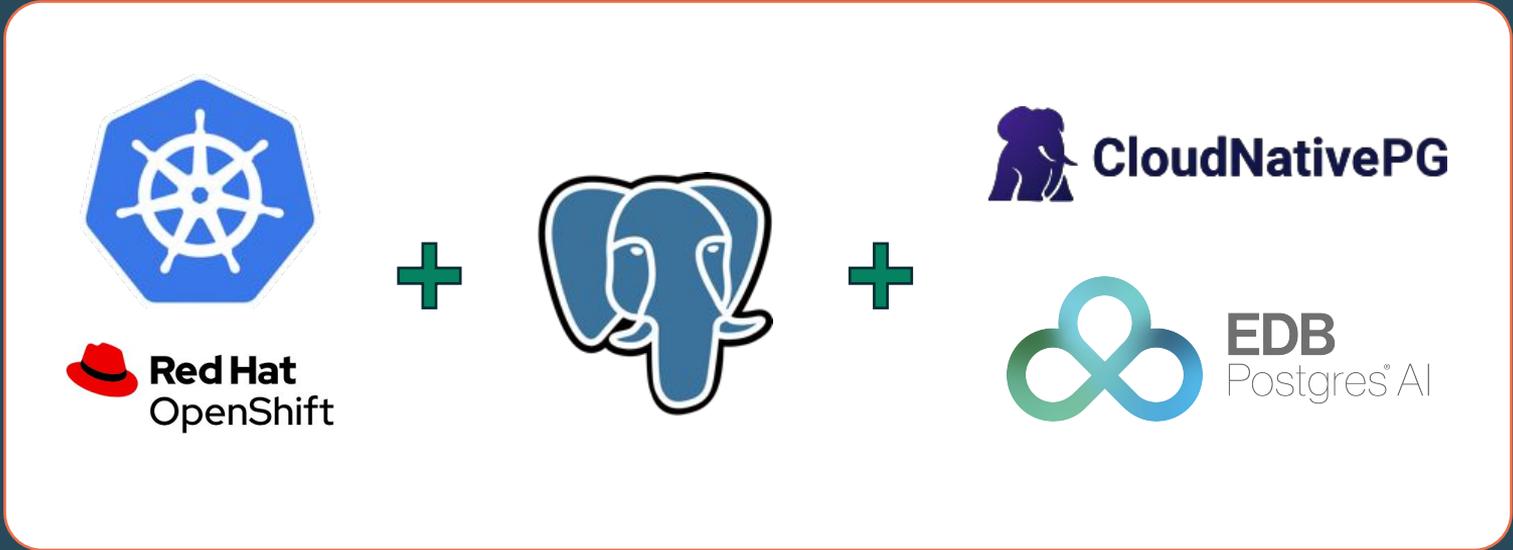


CNPG tries to encode years of experience managing PostgreSQL clusters into **an Operator which should automate all the known tasks a user could be willing to do**

Our PostgreSQL operator must simulate the work of a DBA



# Win Technology



# Autopilot

It automates the steps that a human operator would do to deploy and to manage a Postgres database inside Kubernetes, including automated failover.



# Security

CloudNativePG is secured by default.

**SECURITY**



It doesn't rely on statefulsets and uses its own way to manage persistent volume claims where the PGDATA is stored.

## Data persistence



## Designed for Kubernetes

It's entirely declarative, and directly integrates with the Kubernetes API server to update the state of the cluster — for this reason, it does not require an external failover management tool.



# Features

Deployment	Administration	Backup & Recovery	Monitoring	Security	High Availability
Kubernetes operator	Single node	Backup	Prometheus	TDE	Switchover
Kubernetes plugin	Cluster (Multi node)	Recovery	Grafana dashboards	Certificates	Failover
EDB Postgres (EPAS)	PostgreSQL configuration	PITR	Postgres Enterprise Manager	Data redaction	Scale out / scale down
PostGIS	Pooling	Volume Snapshots	Logging	Password management	Minor / Major updates



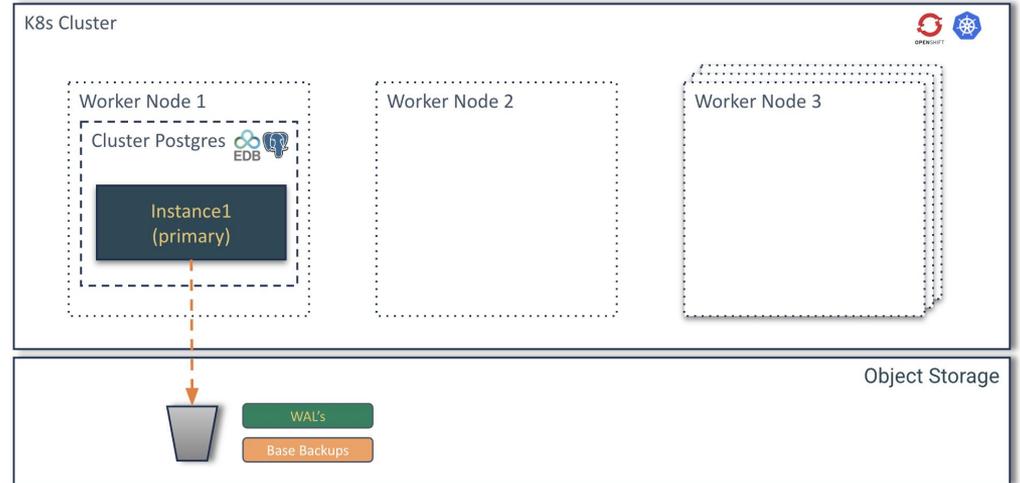
# Use Cases



# Use case 1 architecture

A single database is the simplest setup, involving one instance of a database server.

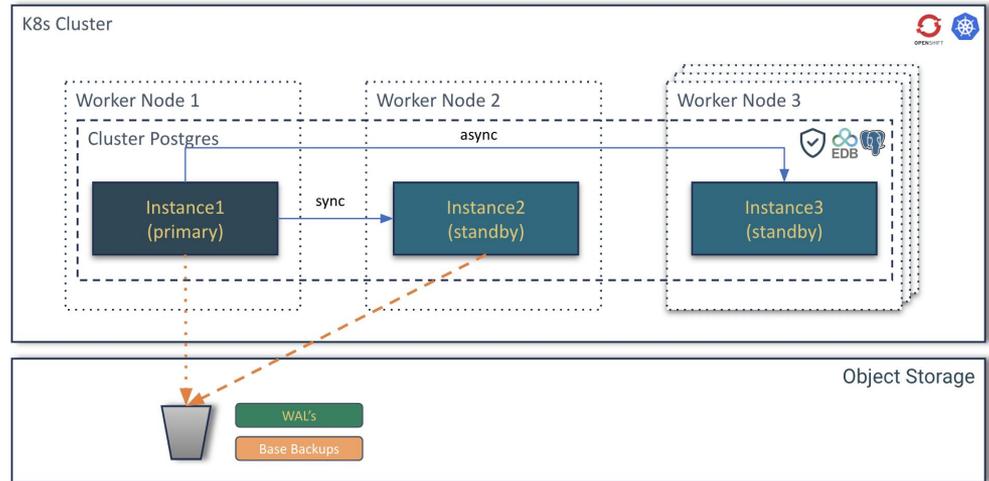
- Development and testing environments
- Small applications with low traffic
- Non-critical data analysis
- Applications with high tolerance for downtime
- Cost-sensitive projects



# Use case 2 architecture

An HA database setup aims to minimize downtime by having redundant components. If one component fails, another takes over automatically or with minimal intervention. This usually involves techniques like clustering, replication, or mirroring within the same data center or availability zone.

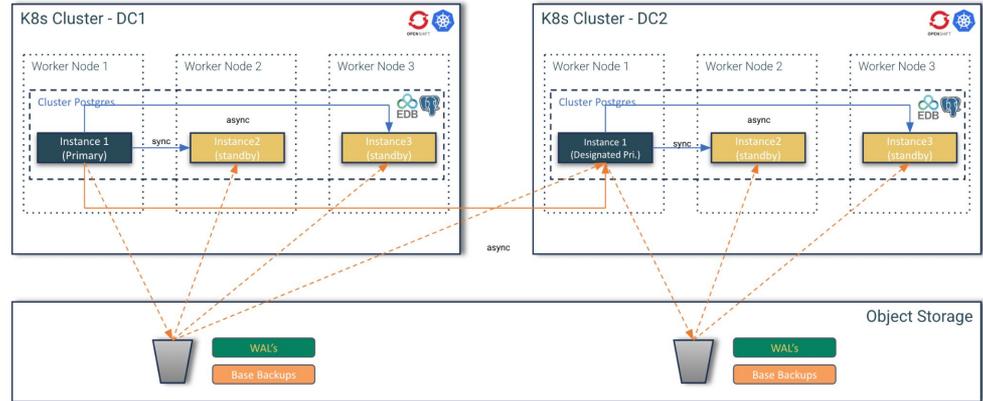
- Business critical Applications
- Applications with stringent SLAs
- Real-time systems
- Improving user experience
- Minimizing planned downtime



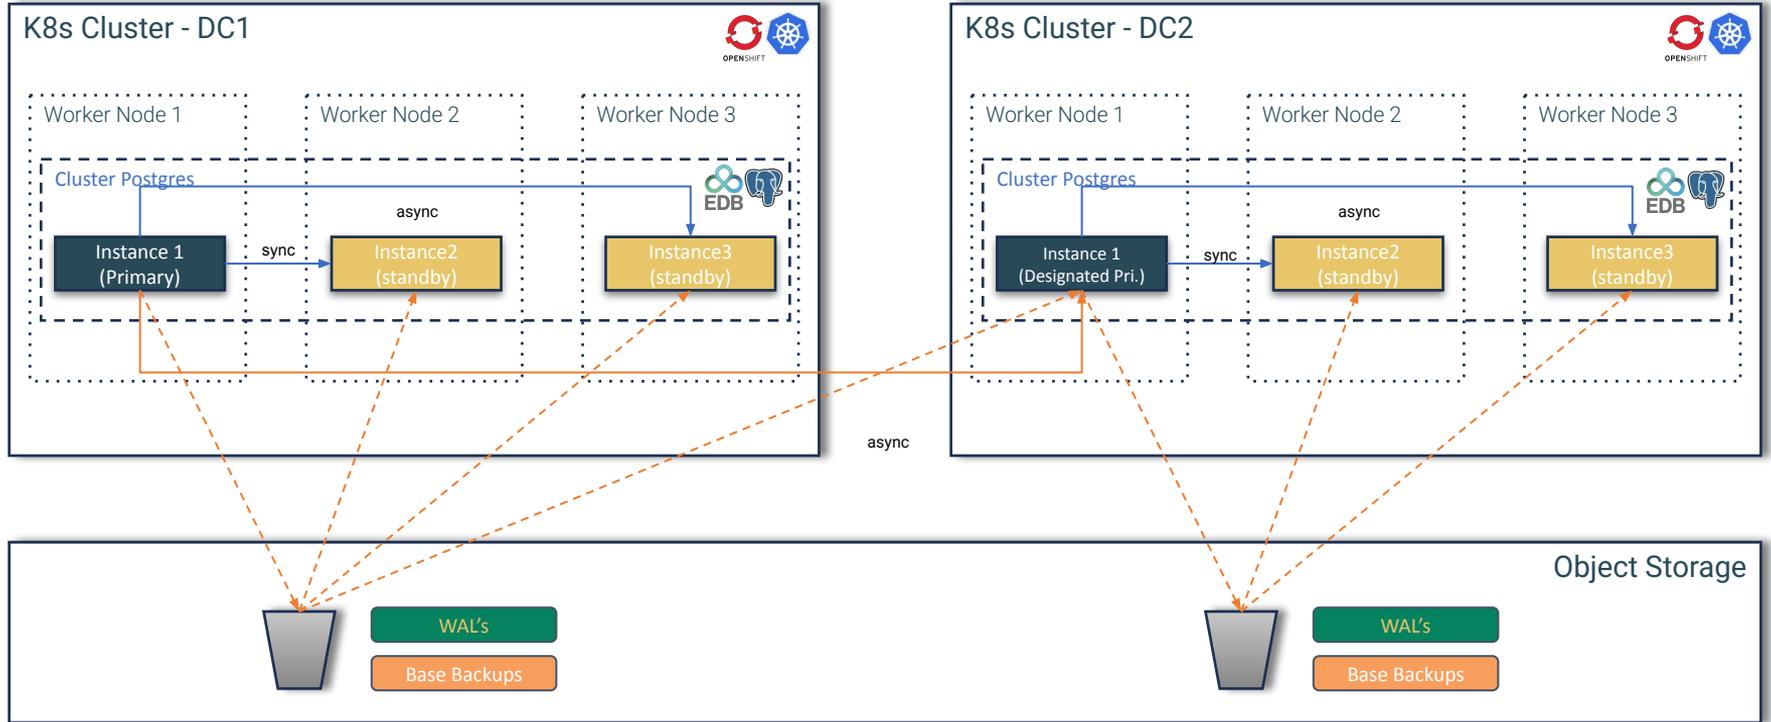
# Use case 3 architecture

A DR database setup focuses on protecting data and ensuring business continuity in the event of a large-scale disaster affecting an entire data center or region (e.g., natural disasters, power outages, cyberattacks). This typically involves replicating data to a geographically separate location.

- Regulatory compliance
- Protecting against catastrophic data loss
- Ensuring business continuity for mission-critical systems



# Use case 3 architecture



Interactive session  
It's time to go hands-on!

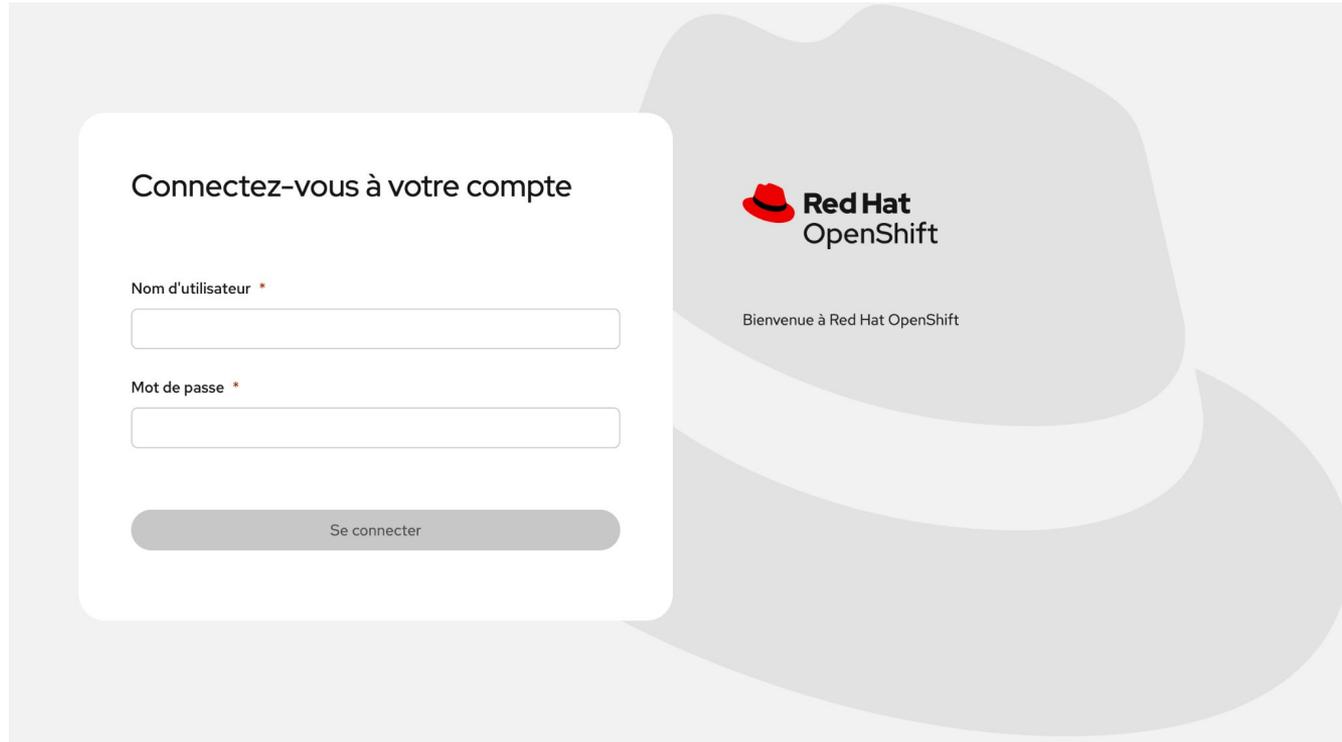


# Environment

<b>OpenShift Console</b>	<a href="https://console-openshift-console.apps.cluster-ldqng.ldqng.sandbox2860.opentlc.com">https://console-openshift-console.apps.cluster-ldqng.ldqng.sandbox2860.opentlc.com</a>
<b>OpenShift User name</b>	user1..user30
<b>OpenShift Password</b>	edbredhat
<b>Minio UI</b>	<a href="https://minio-ui-minio.apps.cluster-ldqng.ldqng.sandbox2860.opentlc.com/">https://minio-ui-minio.apps.cluster-ldqng.ldqng.sandbox2860.opentlc.com/</a>
<b>Minio internal service</b>	<a href="http://minio-service.minio.svc.cluster.local:9000">http://minio-service.minio.svc.cluster.local:9000</a>
<b>Minio API</b>	<a href="https://minio-api-minio.apps.cluster-ldqng.ldqng.sandbox2860.opentlc.com/">https://minio-api-minio.apps.cluster-ldqng.ldqng.sandbox2860.opentlc.com/</a>
<b>Minio User</b>	minio
<b>Minio password</b>	edb-workshop



# Red Hat OpenShift Login



The image shows a login form for Red Hat OpenShift. The form is white and rounded, set against a background of a large, light gray fedora hat. The form contains the following elements:

- Header:** "Connectez-vous à votre compte" (Connect to your account)
- Username Field:** "Nom d'utilisateur \*" (Username) with a red asterisk, followed by a white input box.
- Password Field:** "Mot de passe \*" (Password) with a red asterisk, followed by a white input box.
- Login Button:** A gray button with the text "Se connecter" (Log in).

To the right of the form, the Red Hat logo (a red fedora) is displayed above the text "Red Hat OpenShift". Below this, the text "Bienvenue à Red Hat OpenShift" (Welcome to Red Hat OpenShift) is visible.



# Red Hat OpenShift environment

The screenshot shows the Red Hat OpenShift console interface. At the top left is the Red Hat OpenShift logo. The top right shows the user 'kube:admin' and a notification bell with '6' alerts. A blue banner at the top of the main content area reads: 'You are logged in as a temporary administrative user. Update the [cluster OAuth configuration](#) to allow others to log in.'

Below the banner, the current project is 'openshift-image-registry'. The main section is titled 'Installed Operators' with a star icon. It contains a paragraph: '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](#).

There is a search bar with a dropdown menu set to 'Name' and the text 'Search by name...'. Below this is a table of installed operators.

Name	Managed Namespaces	Status	Provided APIs
<a href="#">EDB Postgres for Kubernetes</a> 1.28.0 provided by EDB	<a href="#">openshift-image-registry</a> The operator is running in openshift-operators but is managing this namespace	Succeeded Up to date	<a href="#">Backups</a> <a href="#">Cluster Image Catalog</a> <a href="#">Cluster</a> <a href="#">Postgres Database</a> <a href="#">View 6 more...</a>



# Red Hat OpenShift operator

The screenshot shows the Red Hat OpenShift console interface. At the top, it indicates the user is logged in as a temporary administrative user. The main content area displays the details for the 'EDB Postgres for Kubernetes' operator, version 128.0. The 'Provided APIs' section lists several APIs with their descriptions and 'Create instance' links:

- Backups**: PostgreSQL backup (physical base backup). [Create instance](#)
- Cluster Image Catalog**: A cluster-wide catalog of PostgreSQL operand images. [Create instance](#)
- Cluster**: PostgreSQL cluster (primary/standby architecture). [Create instance](#)
- Postgres Database**: Declarative creation and management of a database on a Cluster. [Create instance](#)
- Fallover Quorum**: FalloverQuorum contains the information about the current fallover quorum status of a PG cluster. [Create instance](#)
- Image Catalog**: A catalog of PostgreSQL operand images. [Create instance](#)
- Pooler**: Pooler for a Postgres Cluster (with PgBouncer). [Create instance](#)
- Postgres Publication**: Declarative creation and management of a Logical Replication Publication in a PostgreSQL Cluster. [Create instance](#)
- Scheduled Backups**: Backup scheduler for a given Postgres cluster. [Create instance](#)

Additional information on the right side includes the provider (EDB), creation time (Dec 17, 2025, 10:38 AM), links to documentation and maintainers, and a list of maintainers with their contact information.



# Use case

# The environment



# Features shown during the demo

- Kubernetes plugin install
- Check the CloudNativePG operator status
- Postgres cluster install
- Insert data in the cluster
- Failover
- Backup
- Recovery
- Scale out/down
- Fencing
- Hibernation
- Monitoring
- Rolling updates (minor and major)

Deployment

Administration

Backup and  
Recovery

High Availability

Monitoring

Last CloudNativePG tested version is 1.25



This demo is in 

[https://github.com/raphael-chir/edb-postgres-f  
or-kubernetes-in-openshift](https://github.com/raphael-chir/edb-postgres-f<br/>or-kubernetes-in-openshift)

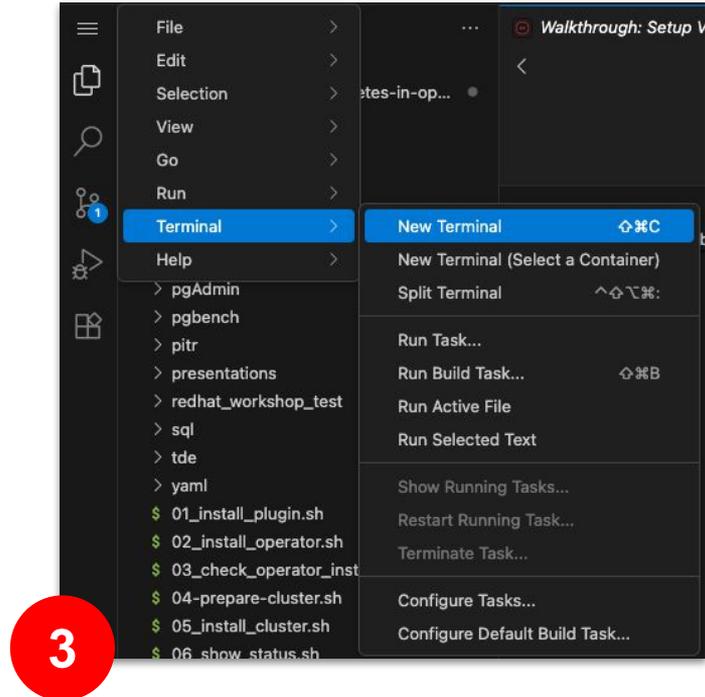
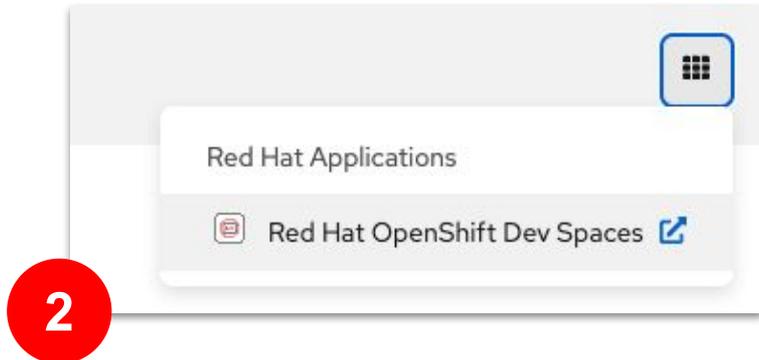
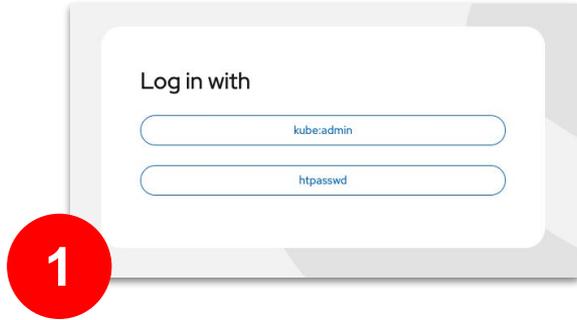


# Hands-on



# OpenShift and DevSpaces access

<https://github.com/raphael-chir/edb-postgres-for-kubernetes-in-openshift>



# Install Postgres cluster

## Scripts

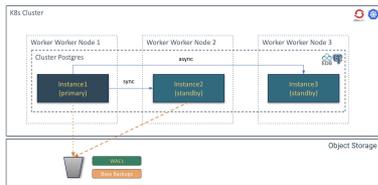
```
04-prepare-cluster.sh
05_install_cluster.sh
06_show_status.sh
07_insert_data.sh
```

## Prepare Postgres cluster template

Setup credentials (AWS or Minio) and prepare context for your users

## Deploy Postgres cluster in HA

Different types of architectures can be deployed. In our demo, 3 instances are deployed with HA available.



## Insert Data

Insert data in test table

Doc [link](#)

## Code

```
export cluster_name="cluster-userXX"
export TMP=$TMPDIR

# Get template
envsubst < templates/${cluster_name}-template.yaml >
$TMP/${cluster_name}.yaml

# Deploy Postgres cluster
${kubectl_cmd} apply -f $TMP/${cluster_name}.yaml
-----
ApiVersion: postgresql.k8s.enterprisedb.io/v1
kind: Cluster
metadata:
  name: cluster-example
spec:
  instances: 3
  imageName: "docker.enterprisedb.com/k8s/edb-postgres-advanced:17.6"
  enableSuperuserAccess: true

  replicationSlots:
    highAvailability:
      enabled: true
```



# Monitoring

## Show status

### Cluster Summary

```
Name: default/cluster-sergio-1
System ID: 7566239054312464401
PostgreSQL Image: docker.enterprisedb.com/k8s/edb-postgres-advanced:17.6
Primary instance: cluster-sergio-1
Primary promotion time: 2025-10-28 11:52:15 +0000 UTC (4m30s)
Status: Cluster in healthy state
Instances: 3
Ready instances: 3
Size: 127M
Current Write LSN: 0/6032300 (Timeline: 1 - WAL File: 000000010000000000000006)
```

### Continuous Backup status (Barman Cloud Plugin)

```
ObjectStore / Server name: object-store/cluster-sergio
First Point of Recoverability: -
Last Successful Backup: -
Last Failed Backup: -
Working WAL archiving: OK
WALs waiting to be archived: 0
Last Archived WAL: 000000010000000000000005.00000028.backup @ 2025-10-28T11:54:11.908195Z
Last Failed WAL: -
```

### Streaming Replication status

#### Replication Slots Enabled

Name	Sent LSN	Write LSN	Flush LSN	Replay LSN	Write Lag	Flush Lag	Replay Lag	State	Sync State	Sync Priority	Replication Slot
cluster-sergio-2	0/9000000	0/9000000	0/9000000	0/9000000	00:00:00	00:00:00	00:00:00	streaming	quorum	1	active
cluster-sergio-3	0/9000000	0/9000000	0/9000000	0/9000000	00:00:00	00:00:00	00:00:00	streaming	quorum	1	active

### Instances status

Name	Current LSN	Replication role	Status	QoS	Manager	Version	Node
cluster-sergio-1	0/9000000	Primary	OK	Guaranteed	1.28.0		ocp4-multiarch-worker1-ppc64
cluster-sergio-2	0/9000000	Standby (sync)	OK	Guaranteed	1.28.0		ocp4-multiarch-w5gg4-master-2
cluster-sergio-3	0/9000000	Standby (sync)	OK	Guaranteed	1.28.0		ocp4-multiarch-w5gg4-master-1



# Promote and Failover

## Scripts

08\_promote.sh

### Promote

The meaning of this command is to promote a pod in the cluster to primary, so you can start with maintenance work or test a switch-over situation in your cluster:

Doc [link](#)

## Code

```
# Promote
. ./config.sh

${kubect1_cnp} promote ${cluster_name} ${cluster-name}-2
```



# Minor upgrade

## Scripts

09\_upgrade.sh

### Minor upgrade

To upgrade to a newer minor version, simply update the PostgreSQL container image reference in your cluster definition, either directly or via image catalogs. CloudNativePG will trigger a rolling update of the cluster, replacing each instance one by one, starting with the replicas. Once all replicas have been updated, it will perform either a switchover or a restart of the primary to complete the process.

Doc [link](#)

## Code

```
# Minor upgrade

# From
apiVersion: postgresql.k8s.enterprisedb.io/v1
kind: Cluster
metadata:
  name: cluster-sergio1
spec:
  instances: 3
  imageName:
    "docker.enterprisedb.com/k8s/edb-postgres-advanced:17.6"
  ...

# To
apiVersion: postgresql.k8s.enterprisedb.io/v1
kind: Cluster
metadata:
  name: cluster-sergio1
spec:
  instances: 3
  imageName:
    "docker.enterprisedb.com/k8s/edb-postgres-advanced:17.7"
  ...
```



# Backups

## Scripts

```
10_backup_cluster.sh  
11_backup_describe.sh
```

## Backup

Backups are used using Barman cloud with the new Barman Plugin. In case the plugin is not compatible (like in Red Hat OpenShift), we will continue using Barman Cloud (without the plugin).

PostgreSQL natively provides first class backup and recovery capabilities based on file system level (physical) copy. These have been successfully used for more than 15 years in mission critical production databases, helping organizations all over the world achieve their disaster recovery goals with Postgres.

Doc [link](#)

## Code

```
cat $TMP/backup.yaml  
  
apiVersion: postgresql.k8s.enterprisedb.io/v1  
kind: Backup  
metadata:  
  name: cluster-example-backup-test  
spec:  
  cluster:  
    name: cluster-example  
  
${kubect1_cmd} apply -f ./yaml/backup.yaml
```



# Recovery

## Scripts

12\_restore\_cluster.sh

## Recovery

In PostgreSQL, recovery refers to the process of starting an instance from an existing physical backup. PostgreSQL's recovery system is robust and feature-rich, supporting Point-In-Time Recovery (PITR)—the ability to restore a cluster to any specific moment, from the earliest available backup to the latest archived WAL file.

Doc [link](#)

## Code

```
cat $TMP/restore.yaml

apiVersion: postgresql.k8s.enterisedb.io/v1
kind: Cluster
metadata:
  name: cluster-restore
spec:
  instances: 1
  imageName: quay.io/enterisedb/postgresql:16.4
  imagePullPolicy: IfNotPresent
  bootstrap:
    recovery:
      source: source
  plugins:
  - name: barman-cloud.cloudnative-pg.io
    isWALArchiver: true
    parameters:
      barmanObjectName: object-store
  externalClusters:
  - name: source
    plugin:
      name: barman-cloud.cloudnative-pg.io
      parameters:
        barmanObjectName: object-store
        serverName: cluster-example

${kubect1_cmd} apply -f $TMP/restore.yaml
```



# Promote and Failover

## Scripts

13\_failover.sh

### Automated Failover

In the case of unexpected errors on the primary for longer than the `.spec.failoverDelay` (by default 0 seconds), the cluster will go into failover mode. This may happen, for example, when:

- The primary pod has a disk failure
- The primary pod is deleted
- The postgres container on the primary has any kind of sustained failure

In the failover scenario, the primary cannot be assumed to be working properly.  
Doc [link](#)

## Code

```
# Automatic Failover (simulate failover killing pods of primary)
${kubect1_cmd} delete pvc/${primary} \
                    pvc/${primary}-wal \
                    pod/${primary} \
                    --force
```



# Scale out, scale down

## Scripts

```
14_scale_out.sh  
15_scale_down.sh
```

### Scale out and down

The operator allows you to scale up and down the number of instances in a PostgreSQL cluster. New replicas are started up from the primary server and participate in the cluster's HA infrastructure.

Doc [link](#)

## Code

```
# Scale out  
kubectl scale cluster cluster-example --replicas=4  
  
# Scale down  
kubectl scale cluster cluster-example --replicas=2
```



# Fencing and Hibernation

## Scripts

```
30_fencing_on.sh  
31_fencing_off.sh
```

## Fencing

Fencing is the process of protecting the data in one, more, or even all instances of a PostgreSQL cluster when they appear to be malfunctioning. When an instance is fenced, the PostgreSQL server process is guaranteed to be shut down, while the pod is kept running. This ensures that, until the fence is lifted, data on the pod isn't modified by PostgreSQL and that you can investigate file system for debugging and troubleshooting purposes.

Doc [link](#)

## Code

```
# Fencing on  
${kubect1_cnp} fencing on ${cluster_name} ${replica}  
  
# Fencing off  
${kubect1_cnp} fencing off ${cluster_name} ${replica}
```



# Fencing and Hibernation

## Scripts

```
32_hibernation_on.sh  
33_hibernation_off.sh
```

## Hibernation

CloudNativePG supports hibernation of a running PostgreSQL cluster in a declarative manner, through the `cnpg.io/hibernation` annotation. Hibernation enables saving CPU power by removing the database pods while keeping the database PVCs. This feature simulates scaling to 0 instances.

Doc [link](#)

## Code

```
# Hibernation on  
${kubect1_cmd} annotate cluster ${cluster_name} \  
    --overwrite k8s.enterprisedb.io/hibernation=on  
  
# Hibernation off  
${kubect1_cmd} annotate cluster ${cluster_name} \  
    --overwrite k8s.enterprisedb.io/hibernation=off
```



# Major upgrade (copying data)

## Scripts

20\_upgrade\_major\_version.sh

### Major upgrade (with native logical replication)

Major PostgreSQL releases introduce changes to the internal data storage format, requiring a more structured upgrade process.

CloudNativePG supports three methods for performing major upgrades:

- Logical dump/restore – Blue/green deployment, offline.
- Native logical replication – Blue/green deployment, online.
- **Physical with pg\_upgrade** – In-place upgrade, offline (covered in the "Offline In-Place Major Upgrades" section below).

Doc [link](#)

## Code

```
apiVersion: postgresql.k8s.enterisedb.io/v1
kind: Cluster
metadata:
  name: ${cluster_major_upgrade}
spec:
  instances: 1
  imageName: ${postgres_major_upgrade_image}

storage:
  size: 2Gi

bootstrap:
  initdb:
    import:
      type: microservice
      databases:
        - app
      source:
        externalCluster: ${cluster_name}

externalClusters:
  - name: ${cluster_name}
    connectionParameters:
      # Use the correct IP or host name for the source database
      host: ${cluster_name}-rw
      user: postgres
      dbname: postgres
    password:
      name: ${cluster_name}-superuser
      key: password
```



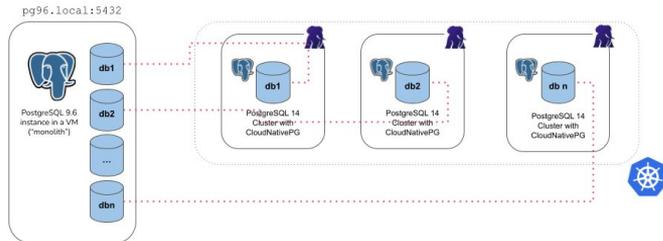
# Major upgrade (copying data)

## Scripts

```
20_upgrade_major_version.sh
```

## Major upgrade (with native logical replication)

- In this step we will migrate the app database from our existing cluster to the new cluster
- We will create the yaml file with the setting **"import"** in the bootstrap section
- The operator uses internally postgres tools `pg_dump` and `pg_restore`
- This method can be used to **migrate** another database or to run **out-of-the-place upgrade**
- Possible settings:
  - Microservices
  - Monolith



## Code

```
apiVersion: postgresql.k8s.enterisedb.io/v1
kind: Cluster
metadata:
  name: cluster-sergio-major
spec:
  instances: 1
  imageName: docker.enterisedb.com/k8s/postgresql:18.1
  imagePullSecrets:
    - name: postgresql-operator-pull-secret
  enableSuperuserAccess: true

  bootstrap:
    initdb:
      import:
        type: microservice
        databases:
          - postgres
        source:
          externalCluster: cluster-sergio

  externalClusters:
    - name: cluster-sergio
      connectionParameters:
        host: cluster-sergio-rw
        user: postgres
        dbname: postgres
      password:
        name: cluster-sergio-superuser
        key: password
    ...
```

# Thank you

For any queries, contact:

[sergio.romera@enterprisedb.com](mailto:sergio.romera@enterprisedb.com)

