



**EDB**

Postgres® for the AI Generation

# EDB Postgres AI Day

Driving The Future Of Data With Postgres AI

Milan, 24th October 2024





# AGENDA

- EDB Postgres AI
- EDB + Red Hat: Database The Cloud-Native Way
- Driving Innovation with EDB Lakehouse & AI
- Customer Testimonial
- Beyond the Cloud: EDB's Formula for Maximizing Postgres Value
- Demystify Complexity Through Reference Architecture and Automation
- Accelerating the Value of Postgres with EDB Services





**EDB**  
Postgres® for the AI Generation

**Andrea Rizzi, Senior Director Southern Europe & Middle East**

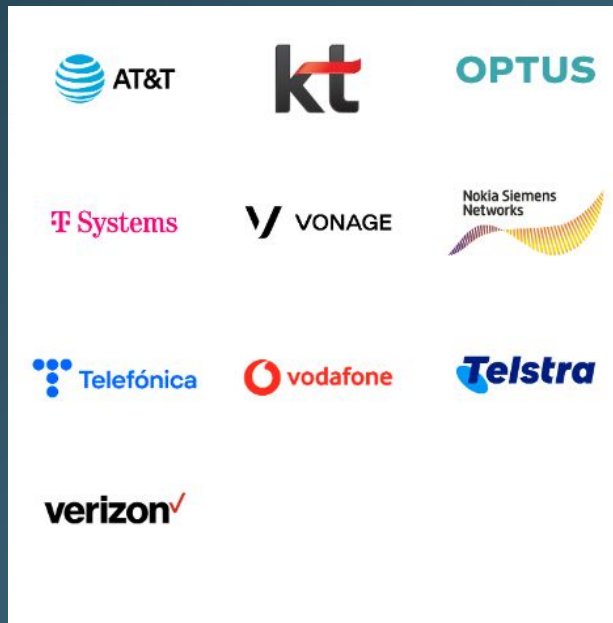


20+ years of Postgres innovation and adoption.  
EDB is the number one contributor  
to the community.





# 1,600+ Leading Enterprises With Complex, Compliant, And Scaling Data Needs Trust EDB



# OUR COMMITMENT: YOUR PEACE OF MIND

		PREMIUM	PRODUCTION	BASIC
	Availability	24 hours x 7 days	24 hours x 7 days	8am - 6pm / Mon - Fri
INITIAL RESPONSE Service Level Objective	Severity-1	15 minutes	30 minutes	N/A
	Severity-2	30 minutes	60 minutes	N/A
	Severity-3	60 minutes	2 hours	4 hours
	Severity-4	1 business day	1 business day	1 business day
REMEDY Service Level Objective	Severity-1	4 hours	24 hours	N/A
	Severity-2	8 hours	48 hours	N/A
	Severity-3	10 business days	15 business days	20 business days
	Severity-4	30 business days	45 business days	60 business days
RESOLUTION Service Level Objective	Severity-1	24 hours	N/A	N/A
	Severity-2	5 days	N/A	N/A





# The Sovereign Data and AI Platform

Merging the power of your data and AI in a scalable, secure and compliant manner to build your sovereign assets into your platform for success



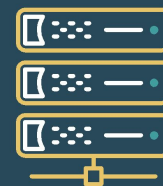
MULTI-CLOUD  
INFRASTRUCTURE



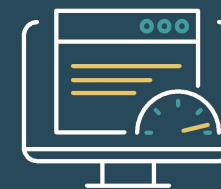
EASY AI INFUSION INTO  
MISSION-CRITICAL  
WORKLOADS



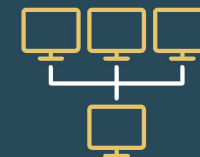
DATA ESTATE MANAGEMENT



APPLIANCE



SOFTWARE



AGILE AND OPEN  
SOURCE ENTERPRISE GRADE



ON-PREMISES



DATA ESTATE OBSERVABILITY



GEO-DISTRIBUTED  
EXTREME HIGH AVAILABILITY

# EDB Postgres AI is the solution to modernize, optimize, and evolve

DO MORE WITH POSTGRES — WHEREVER YOUR DATA IS

## Sovereign Data & 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

### PLATFORM TOOLS AND SERVICES

MIGRATION  
PORTAL

CONTINUOUS  
HIGH  
AVAILABILITY

BACKUP AND  
RECOVERY

### EXTENSIBILITY

CSP INTEGRATIONS

DEVOPS TOOLING

KUBERNETES TOOLING

GENAI & LLM INTEGRATIONS

LAKEHOUSE INTEGRATIONS





Help customers become a data and AI platform.  
Read more in our article in The Times.

THE  TIMES



EDB + Red Hat

# DATABASE THE CLOUD-NATIVE WAY



**Davide Tammaro**  
Senior Sales Engineer  
EDB



**Marco Betti**  
Sr. Account Solution Architect  
Red Hat





# EDB POSTGRES AI ON RED HAT OPENSIFT

**Modern, mission-critical applications demand a new data infrastructure that is open source, flexible, and cloud native.**

- Manage and scale PostgreSQL databases in complex hybrid and multi-cloud environments
- Ensure high availability, security and disaster recovery with minimal downtime
- Reduce total cost of ownership with streamlined deployment, automated operations and an open source secure software supply chain.

## Enterprise grade PostgreSQL for mission-critical databases in the OpenShift Container Platform



**EDB delivers:** Certified, robust PostgreSQL database solutions on OpenShift with faster onramp, built-in high availability, security, scalability, and multi-master replication, ensuring continuous data availability and performance across distributed environments.



**Red Hat delivers:** OpenShift (Kubernetes-based platform) that simplifies the deployment and management of containerized applications, providing the necessary infrastructure to run EDB's advanced database solutions efficiently and securely across any cloud or on-premises environment.



# BETTER TOGETHER: EDB & RED HAT OPENS SHIFT

Enterprise-grade PostgreSQL for mission critical databases in OpenShift Container Platform

With

## RED HAT OPENS SHIFT

HYBRID CLOUD

SELF-HEALING

CLOUD-NATIVE APPLICATIONS



## EDB POSTGRES

ENTERPRISE GRADE

ENHANCED SECURITY

ORACLE COMPATIBILITY

Customers Get

Enterprise-ready PostgreSQL to develop and deploy Cloud-Native applications with scalable databases on hybrid cloud

Key Use Cases

### App Modernization

Enterprise scale and availability

### Agility

DevOps and CI/CD pipeline

### Secure Open Source

PostgreSQL with secure software supply chain

### Cost Reduction

Migrate to Oracle alternative



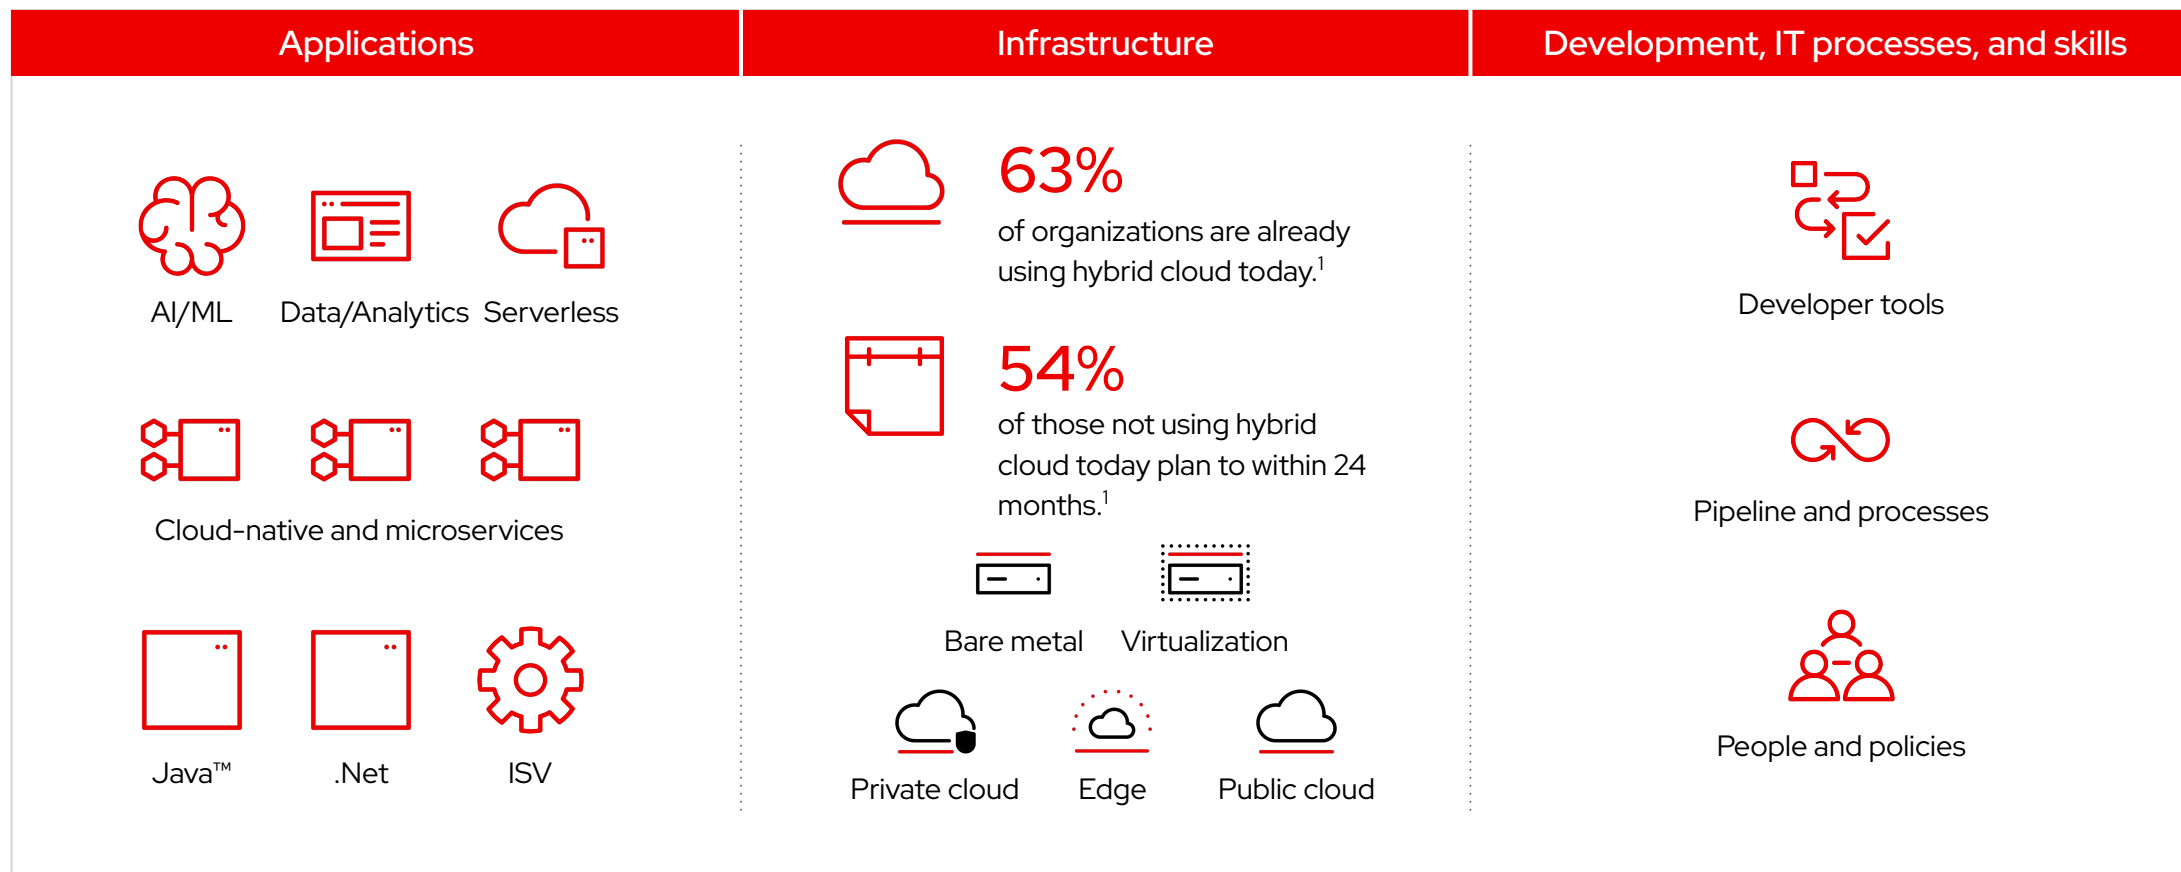




Why Red Hat  
OpenShift?

# Red Hat Open Hybrid Cloud

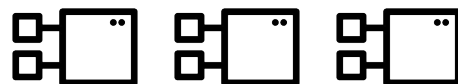
More than a strategy, reality of Enterprise IT



Traditional apps



Cloud-native apps



AI/ML, Functions



Communities of Innovation | Ecosystems of Solutions



Secure & Automated Infrastructure and Operations



Physical



Virtual



Private cloud



Public cloud



Edge



# You need an application platform that is...

## Trusted

Container engine

Reduce  
Risk



## Comprehensive

Application platform

Improve  
Productivity

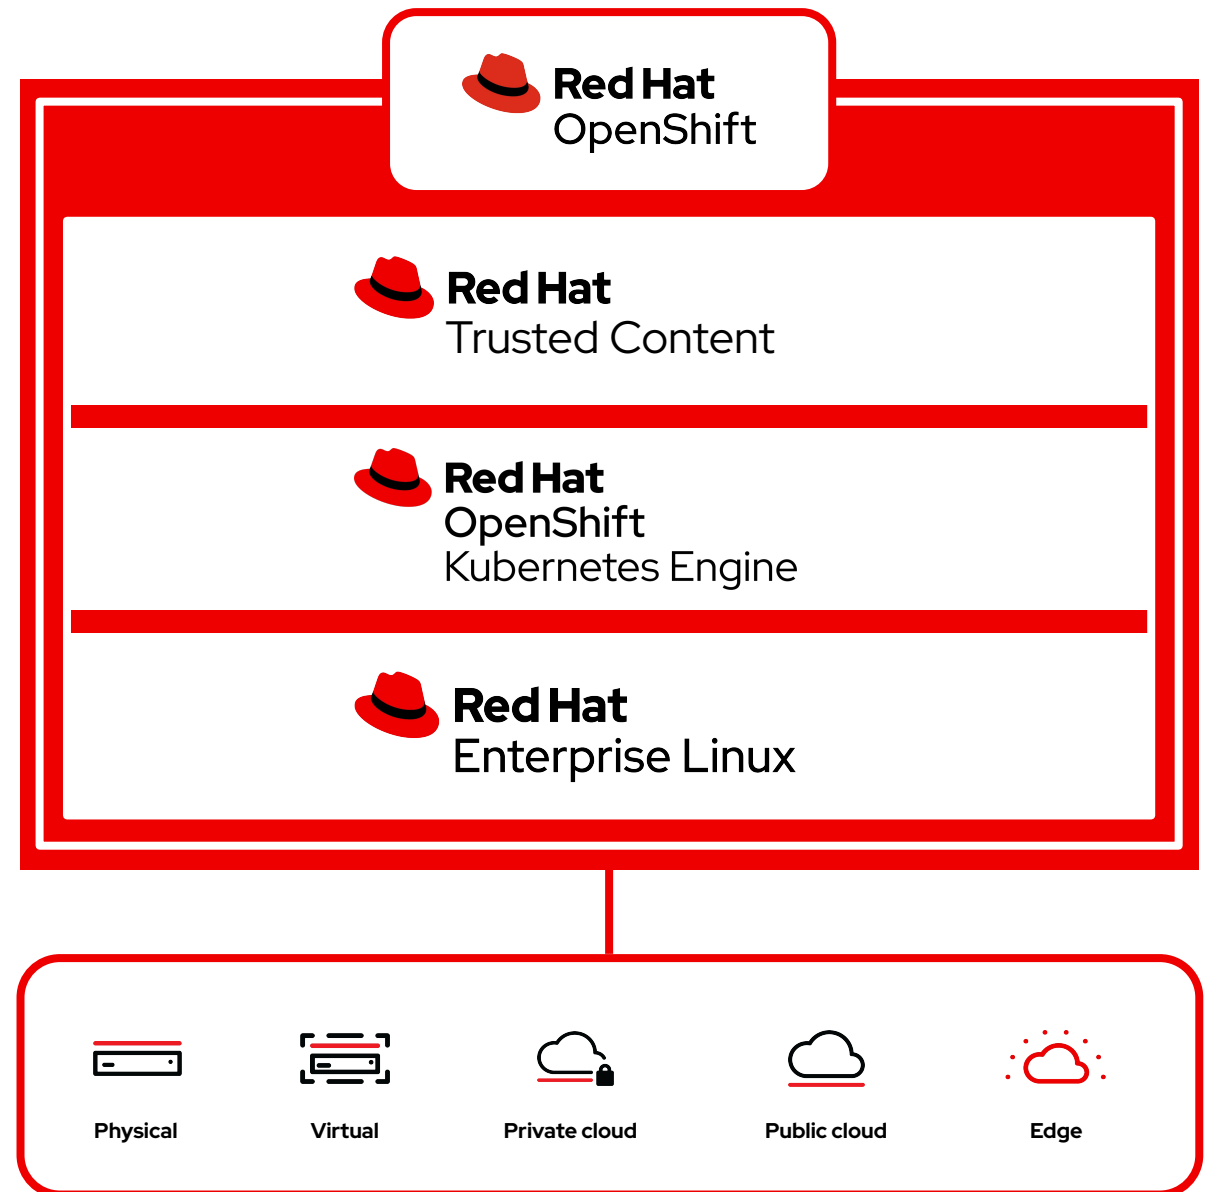


## Consistent

Across hybrid cloud

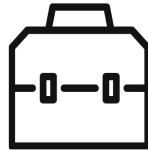
Increase  
Flexibility

# OpenShift built on a **Trusted** Container Engine

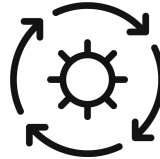


# Components of an Application Platform

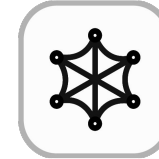
Application developers need more than Kubernetes to be productive



Developer Tools



CI/CD  
GitOps



Service mesh



Monitoring

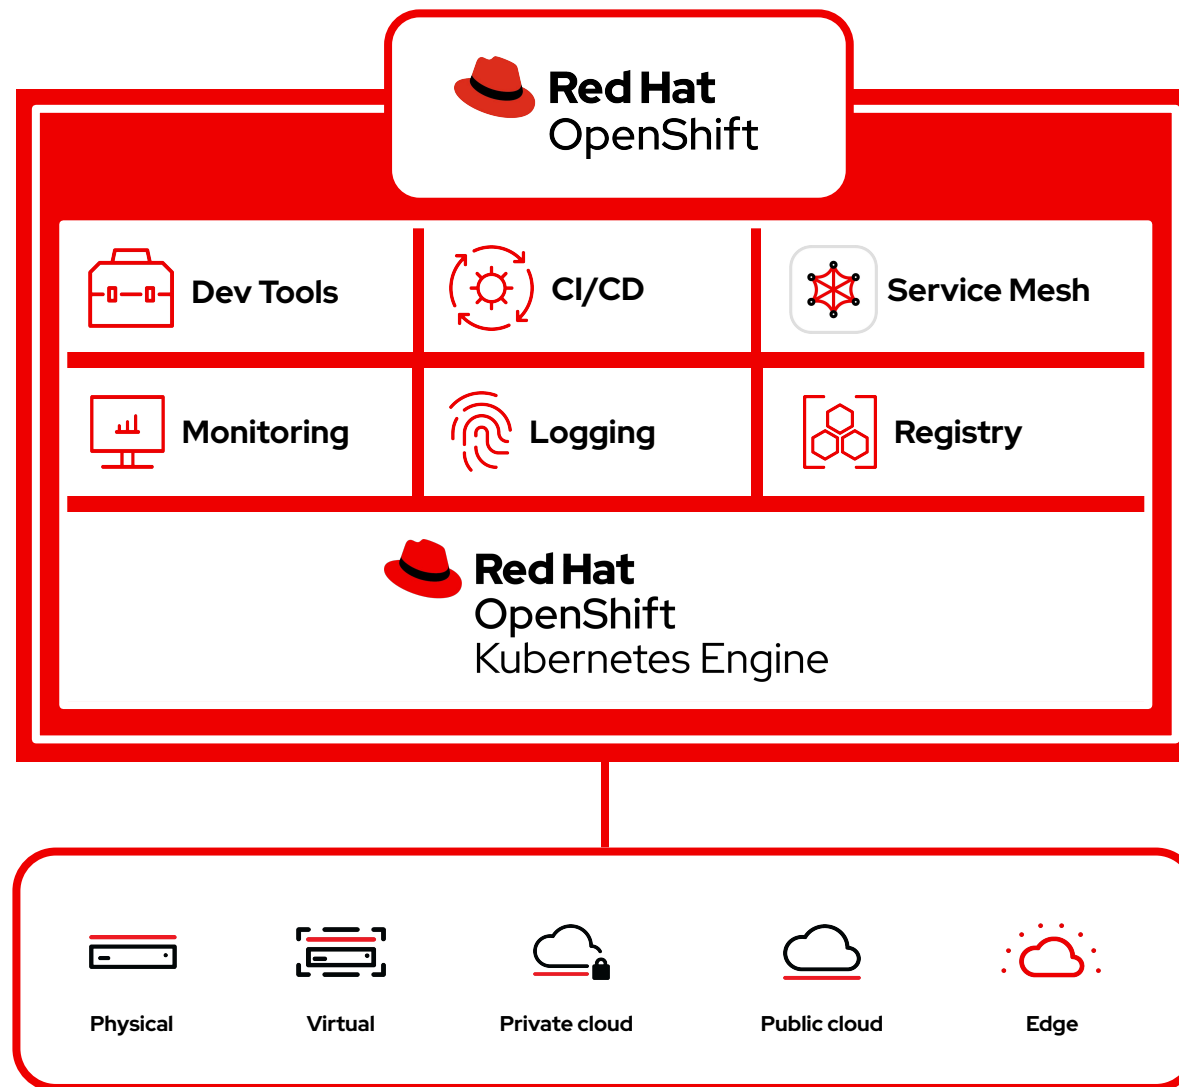


Logging



Container  
Registry

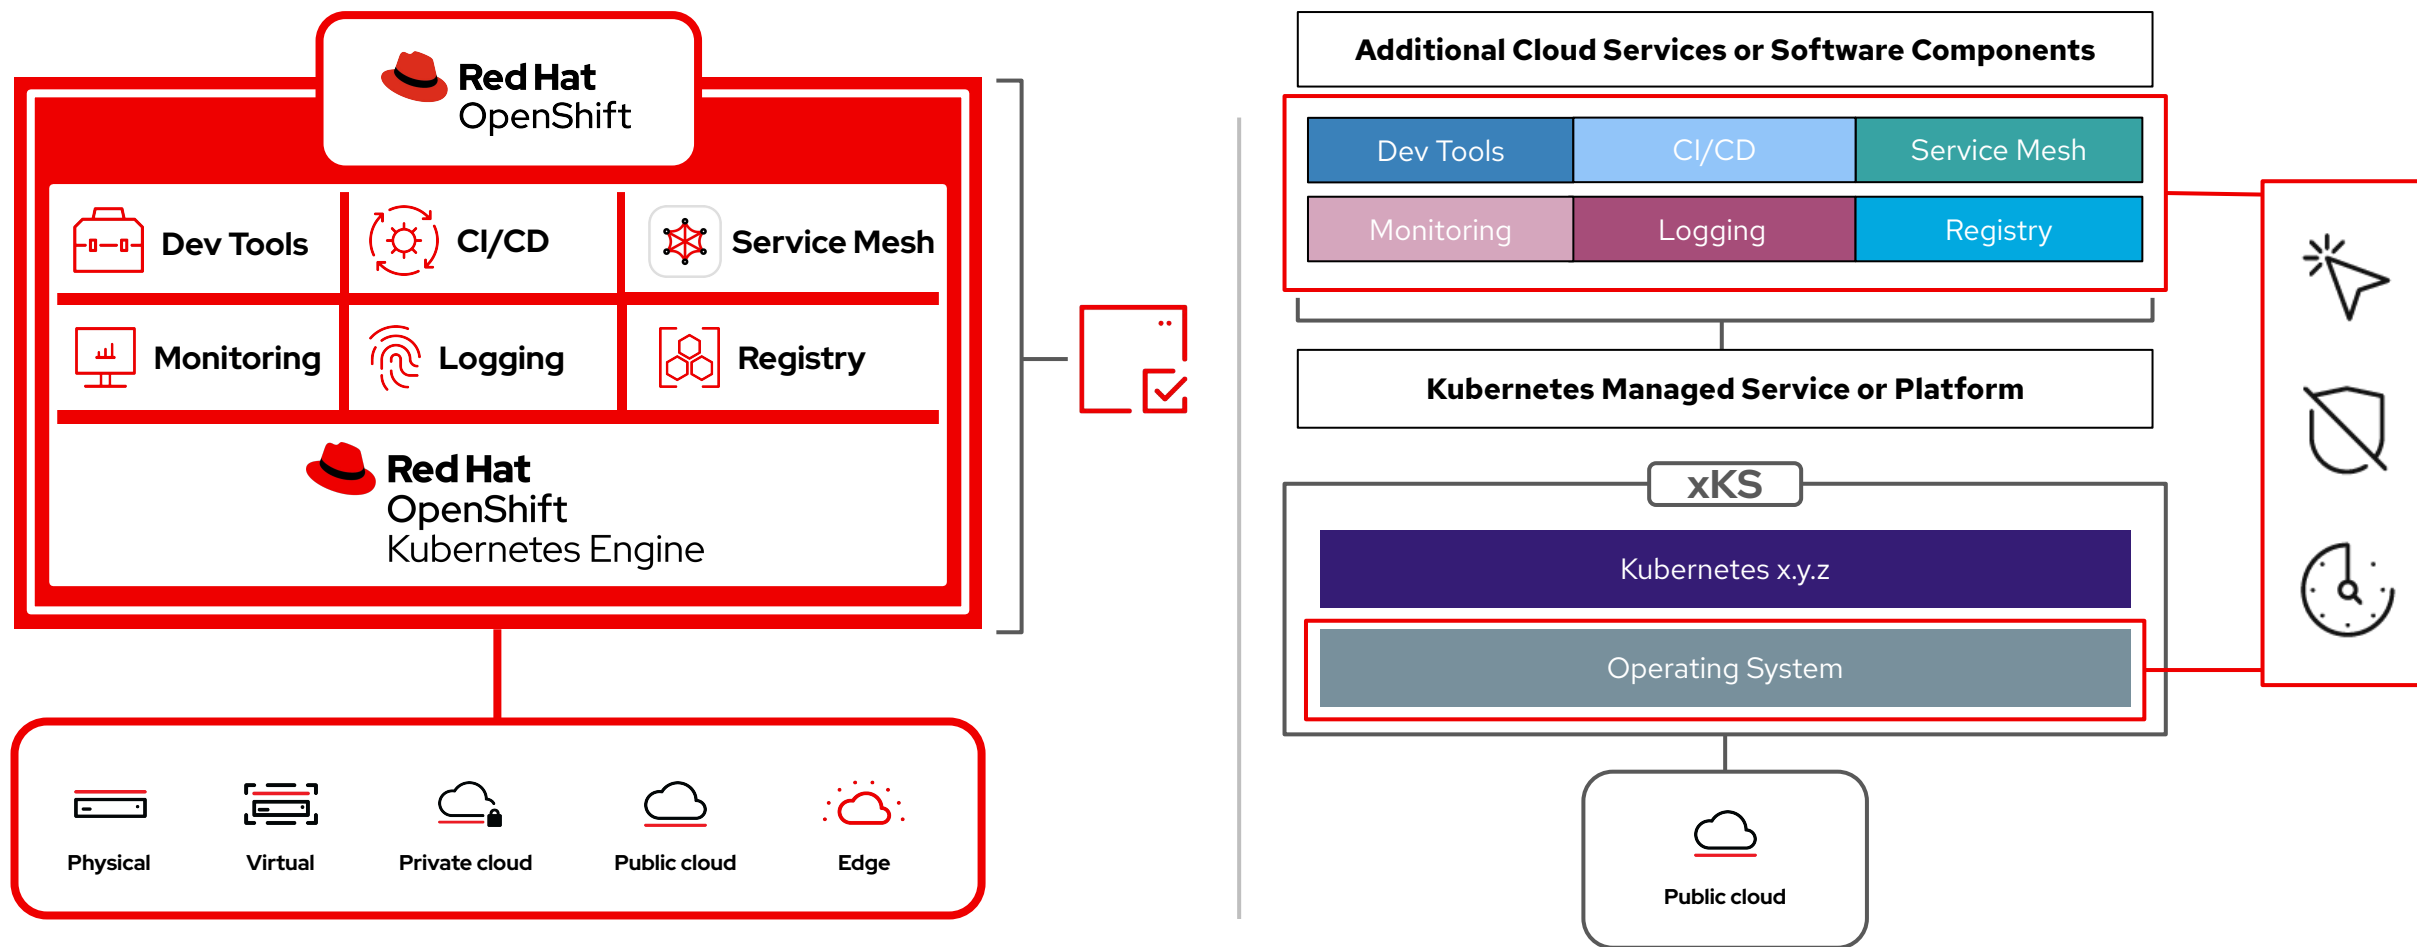
# OpenShift Delivers a **Comprehensive** Application Platform



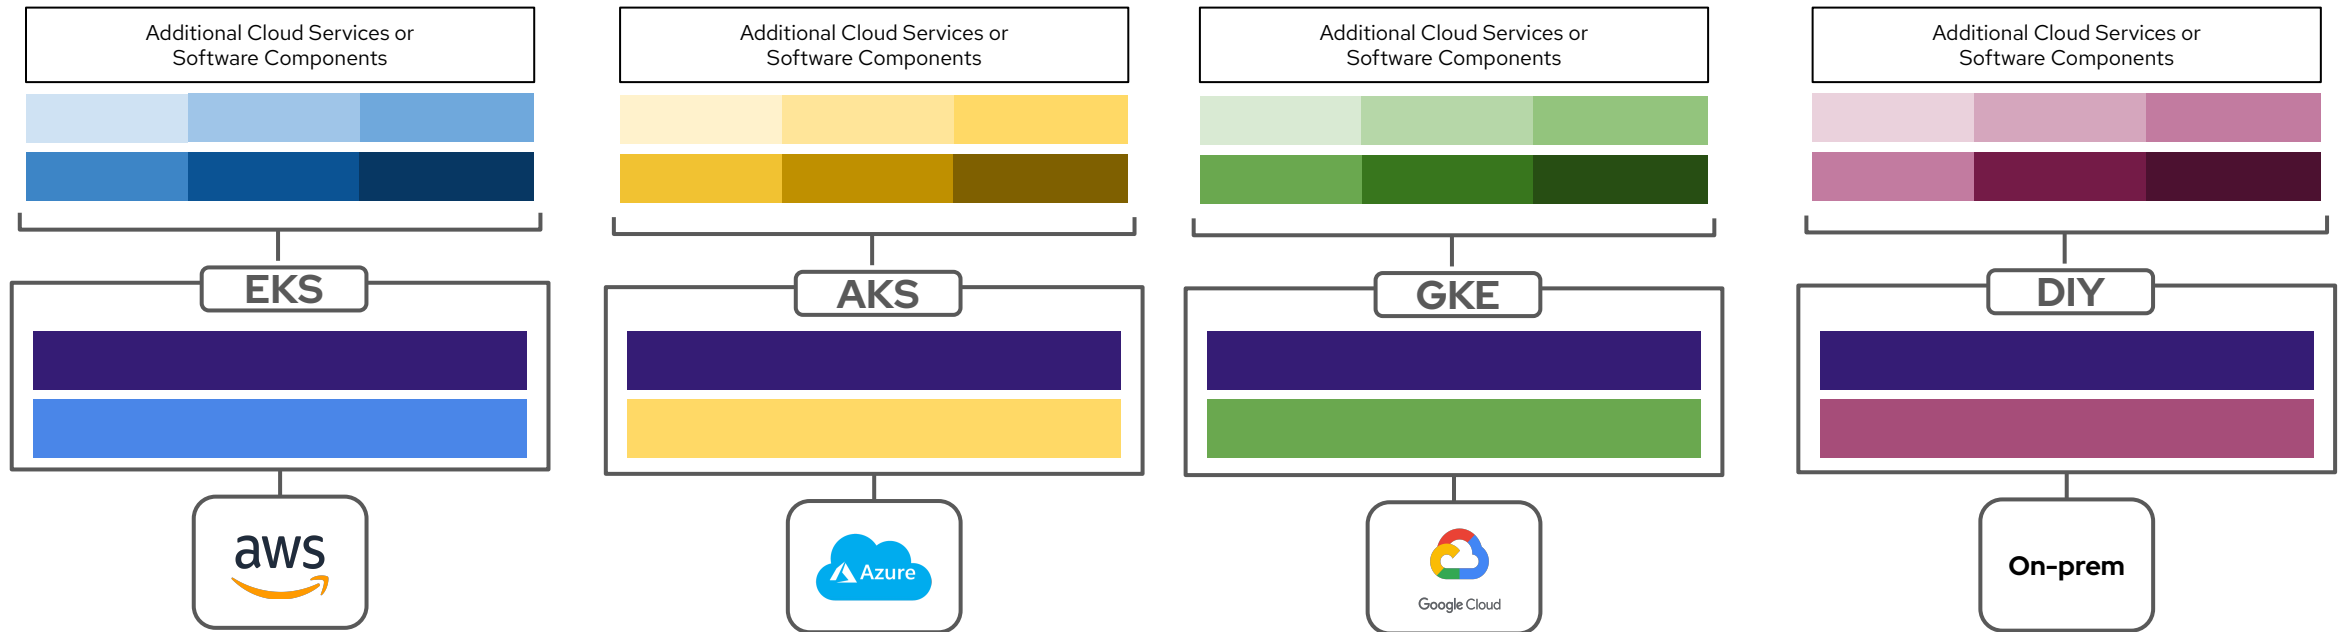


# OpenShift Trust & Consistency Across the Hybrid Cloud

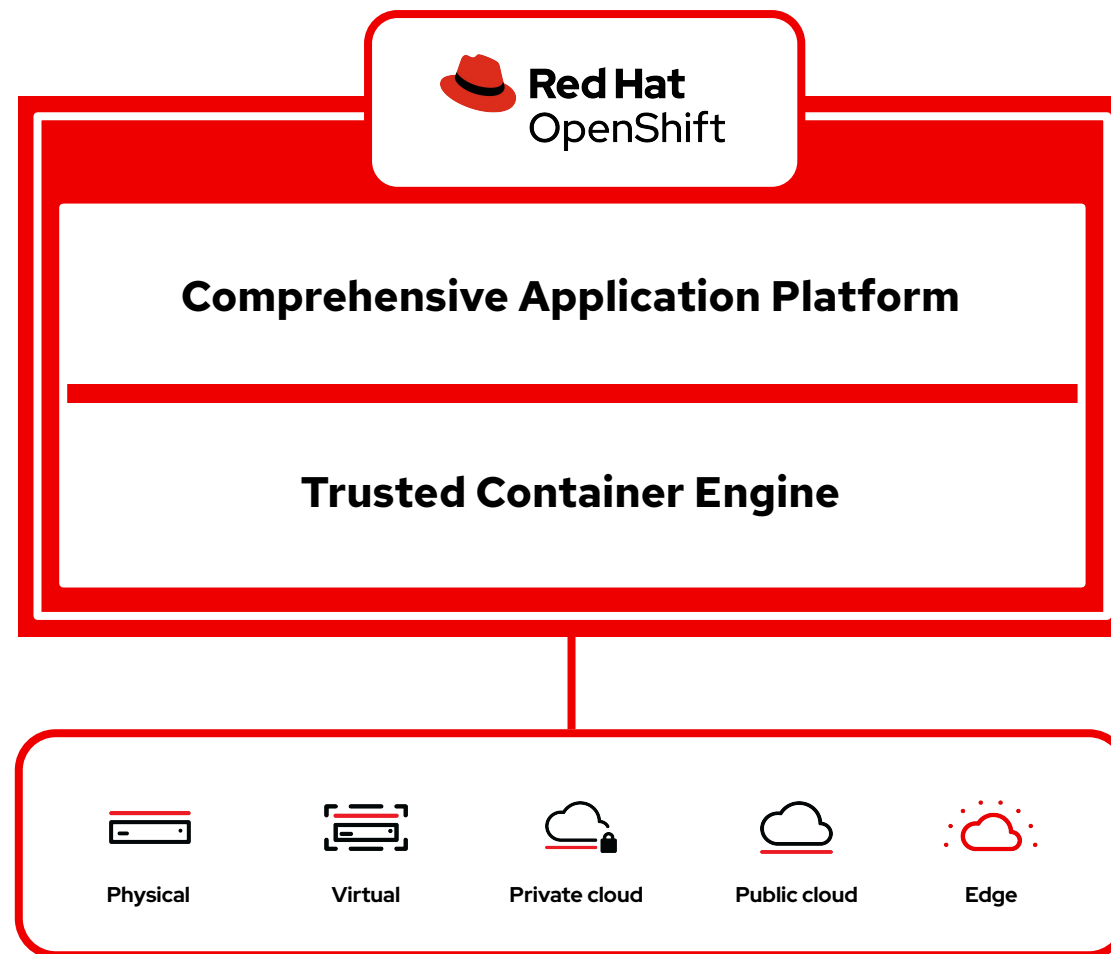
**Versus** managing different stacks across each environment



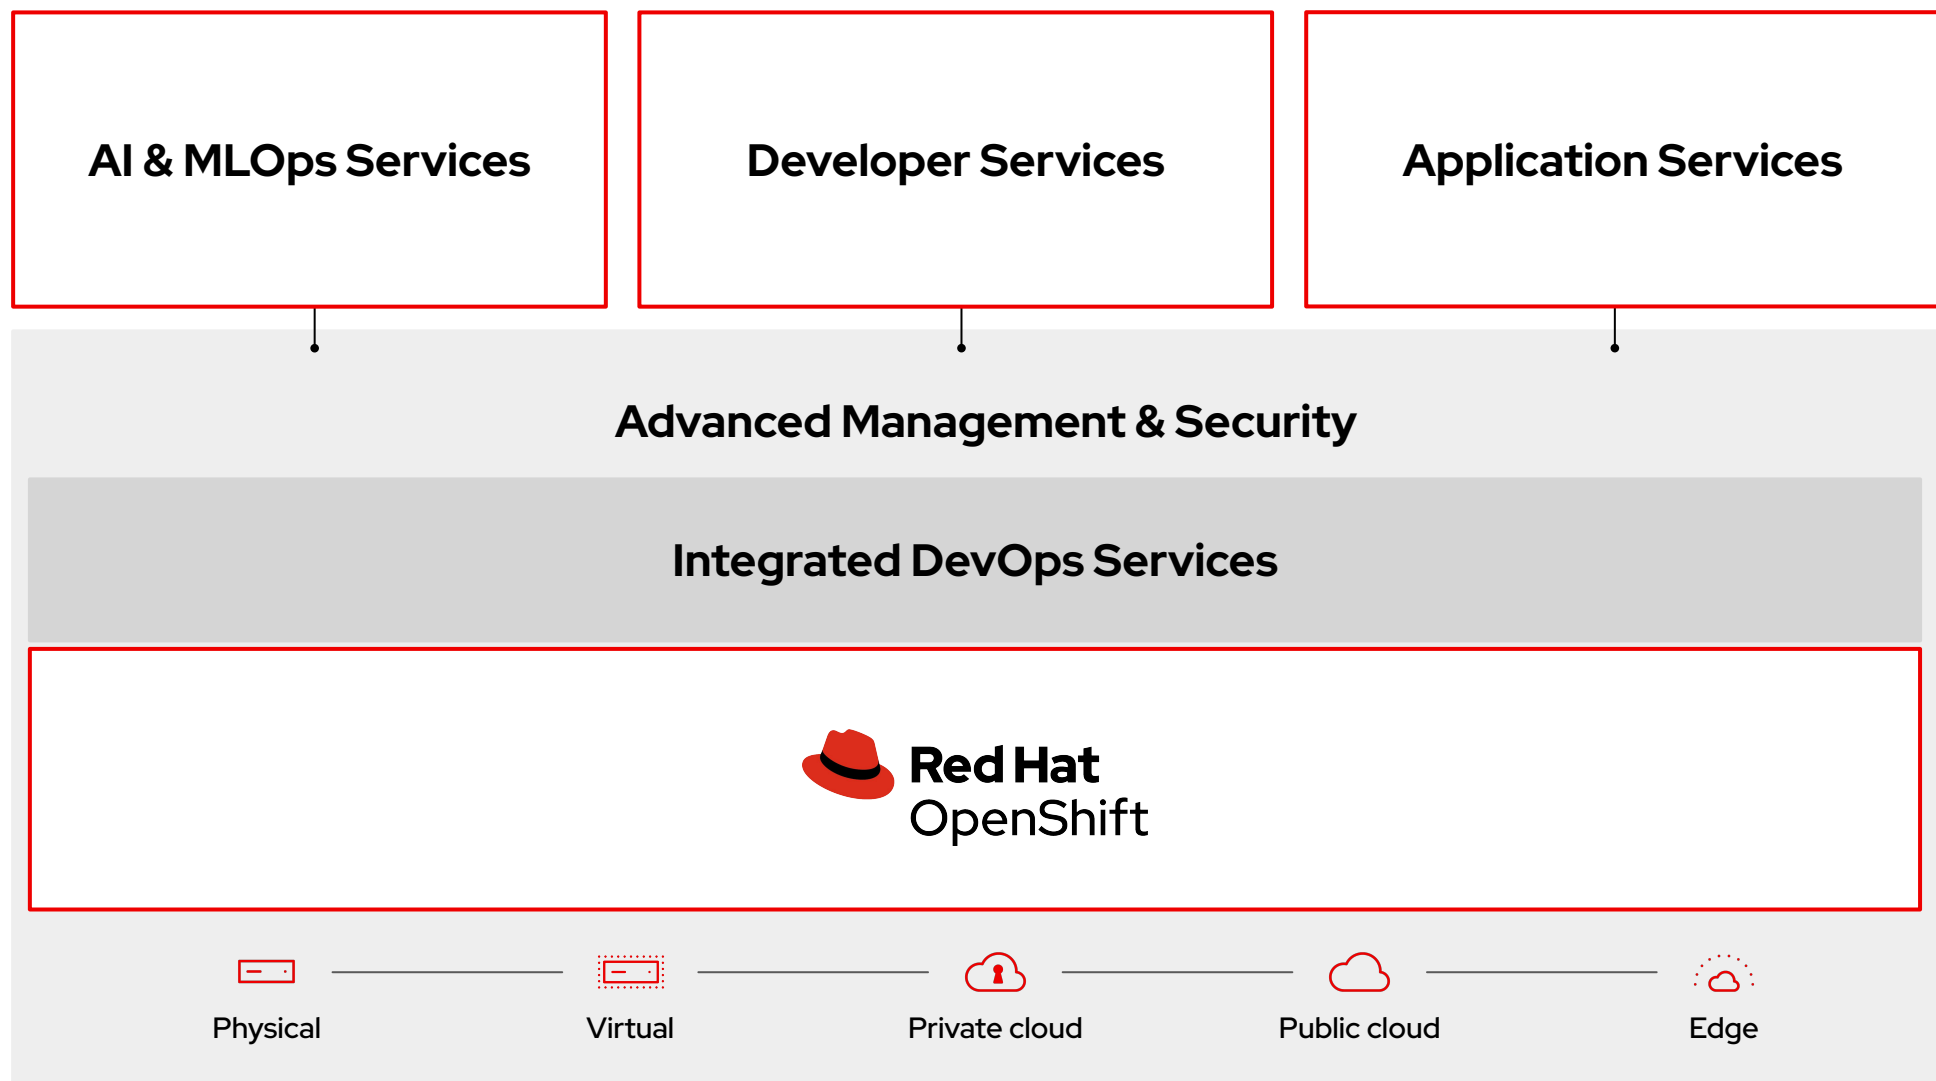
# And that's just for one Public cloud...



OpenShift is  
**Consistent**  
Across a Hybrid Cloud  
environment

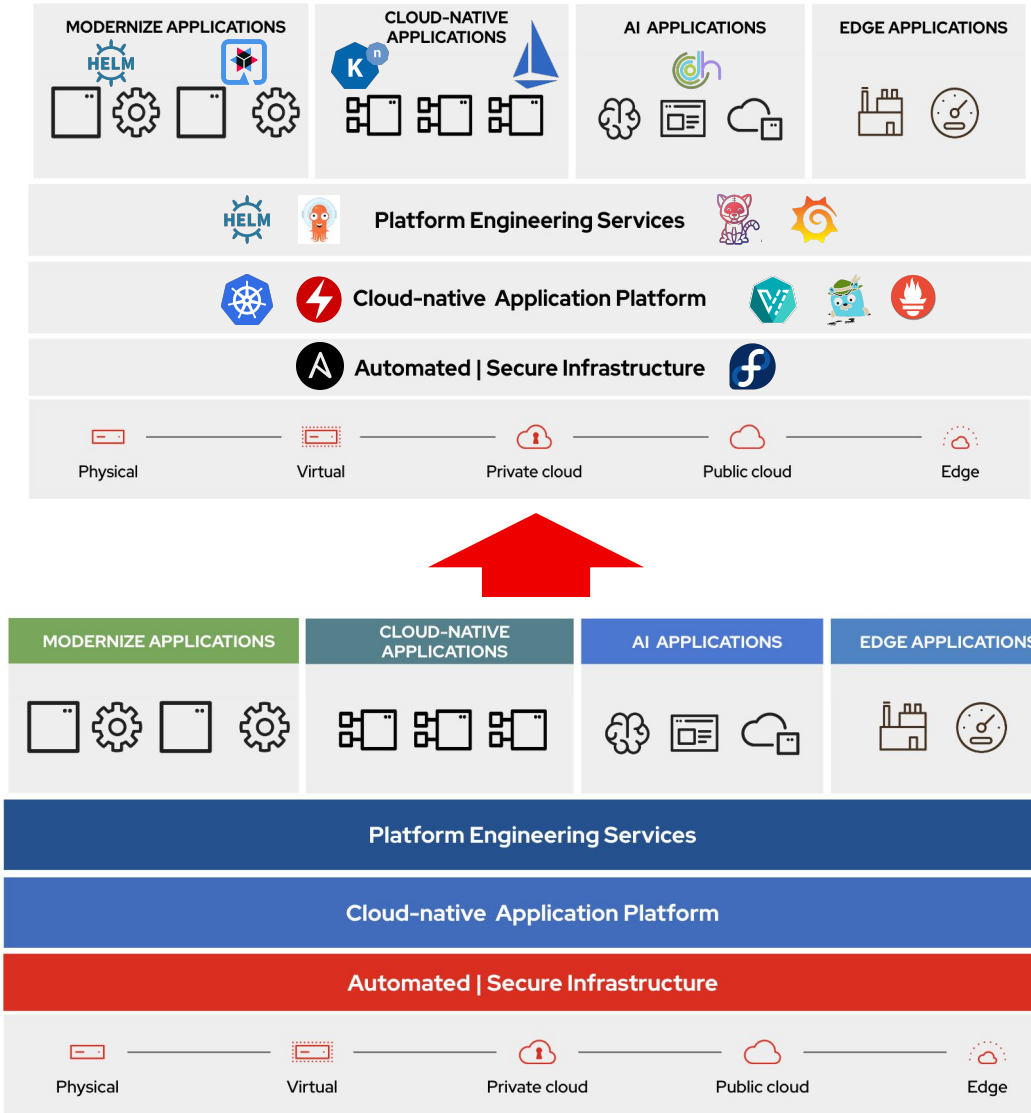


# Where to next?





# Expertise and Experience in CNCF Communities and Platforms



## Contributions to OSS / CNCF Projects\*

**Leading Contributor**



**#2 Contributor**



**#3 Contributor**



**Top 10 Contributor**

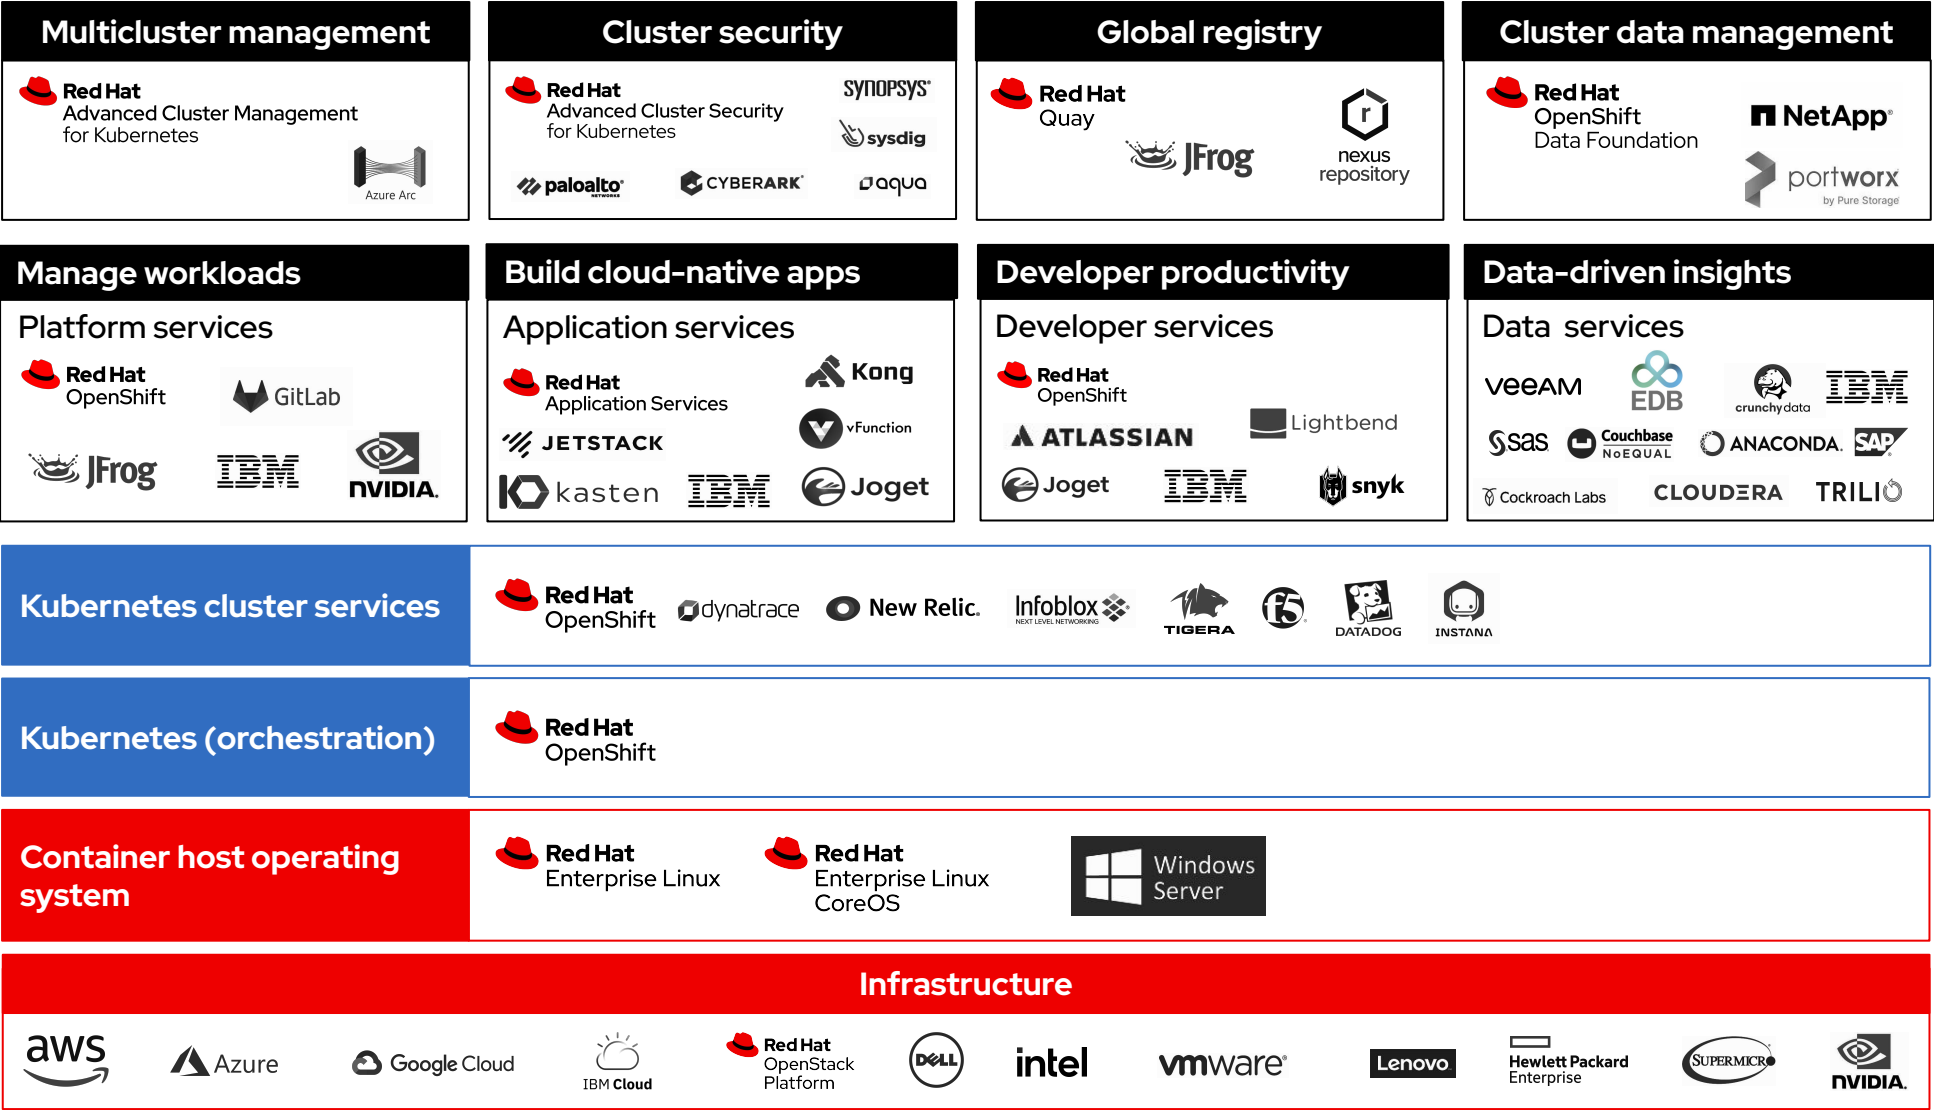


\*Updated Q4 2023

# Red Hat open hybrid cloud platform with ISV ecosystem



Simplify the trial, procurement, and deployment of ISV software on Red Hat OpenShift, anywhere with Red Hat Marketplace



# Why should you run a database within a Container Platform?

Besides objects, messages and unstructured data, Cloud Native Application will need to store structured data.



Running Postgres on Container platform provides:

- **Increased Uptime**
- **Operational Efficiency**
- **Faster Innovation**
- **Reduced TCO**

# EDB Postgres in Openshift





# 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 OPENSIFT OPERATOR FOR POSTGRES



Opensift 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



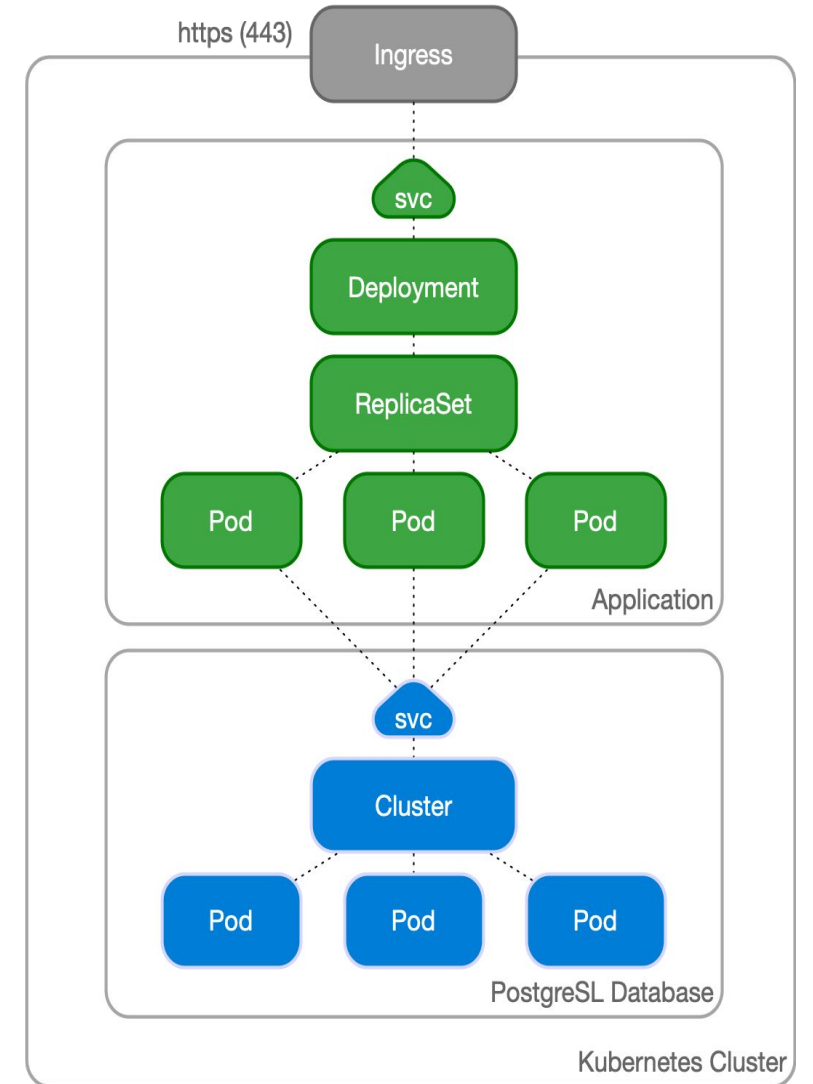
CNP tries to encode years of experience managing PostgreSQL clusters into an Operator which should automate all the known tasks a user could be willing to do

Our PostgreSQL operator must simulate the work of a DBA



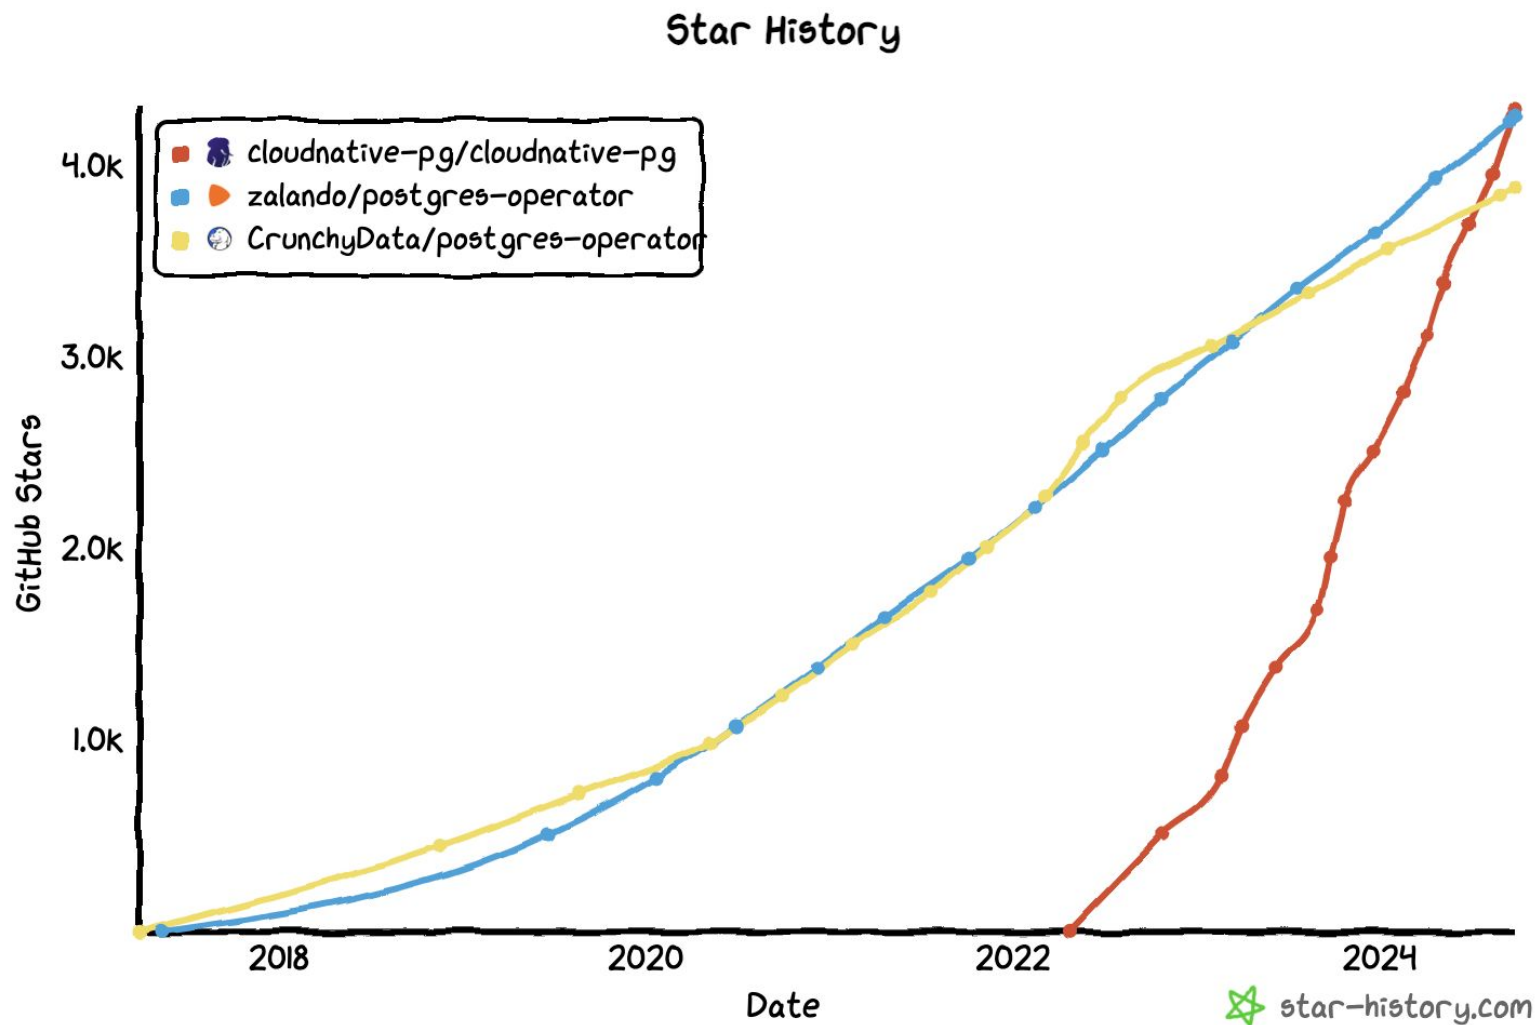
# APPLICATIONS AND DATABASES IN K8S

- Stateless application as a deployment
  - Rolling upgrades
  - ReplicaSet for scaling and HA
  - Custom application images (Go, Django, Java, Python, C, C++, ...)
- Stateful database using our operator
  - Embeds primary/standby logic
  - Service for RW and Read operations
  - Rolling upgrades, scaling, HA, ...
  - "Cluster" CRD



# GITHUB STARS

<https://github.com/cloudnative-pg/cloudnative-pg>



# FEATURES

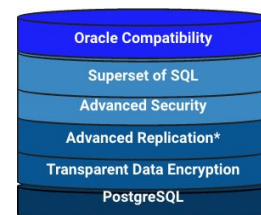
- Automated failover
- Services for RW and RO workloads
- Affinity control
- Backup and Recovery
- Rolling updates
- Scale up/down of read replicas
- Fencing and hibernation
- Native Prometheus exporters
- Log in JSON format to stdout
- OpenShift (and other K8S platforms) compatibility
- TDE (in EDB Postgres for Kubernetes)
- ... and much more



## CloudNativePG

CloudNativePG is a Kubernetes operator that covers the full lifecycle of a PostgreSQL database...

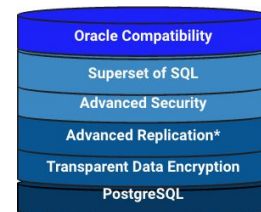
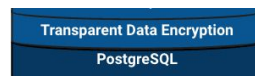
 Containerized application  Storage



## EDB Postgres for Kubernetes

PostgreSQL Operator for mission critical databases

 Containerized application  Storage



## EDB Postgres Distributed

EDB Postgres Distributed for Kubernetes is an operator designed to manage EDB Postgres Distributed...

 Containerized application  Storage

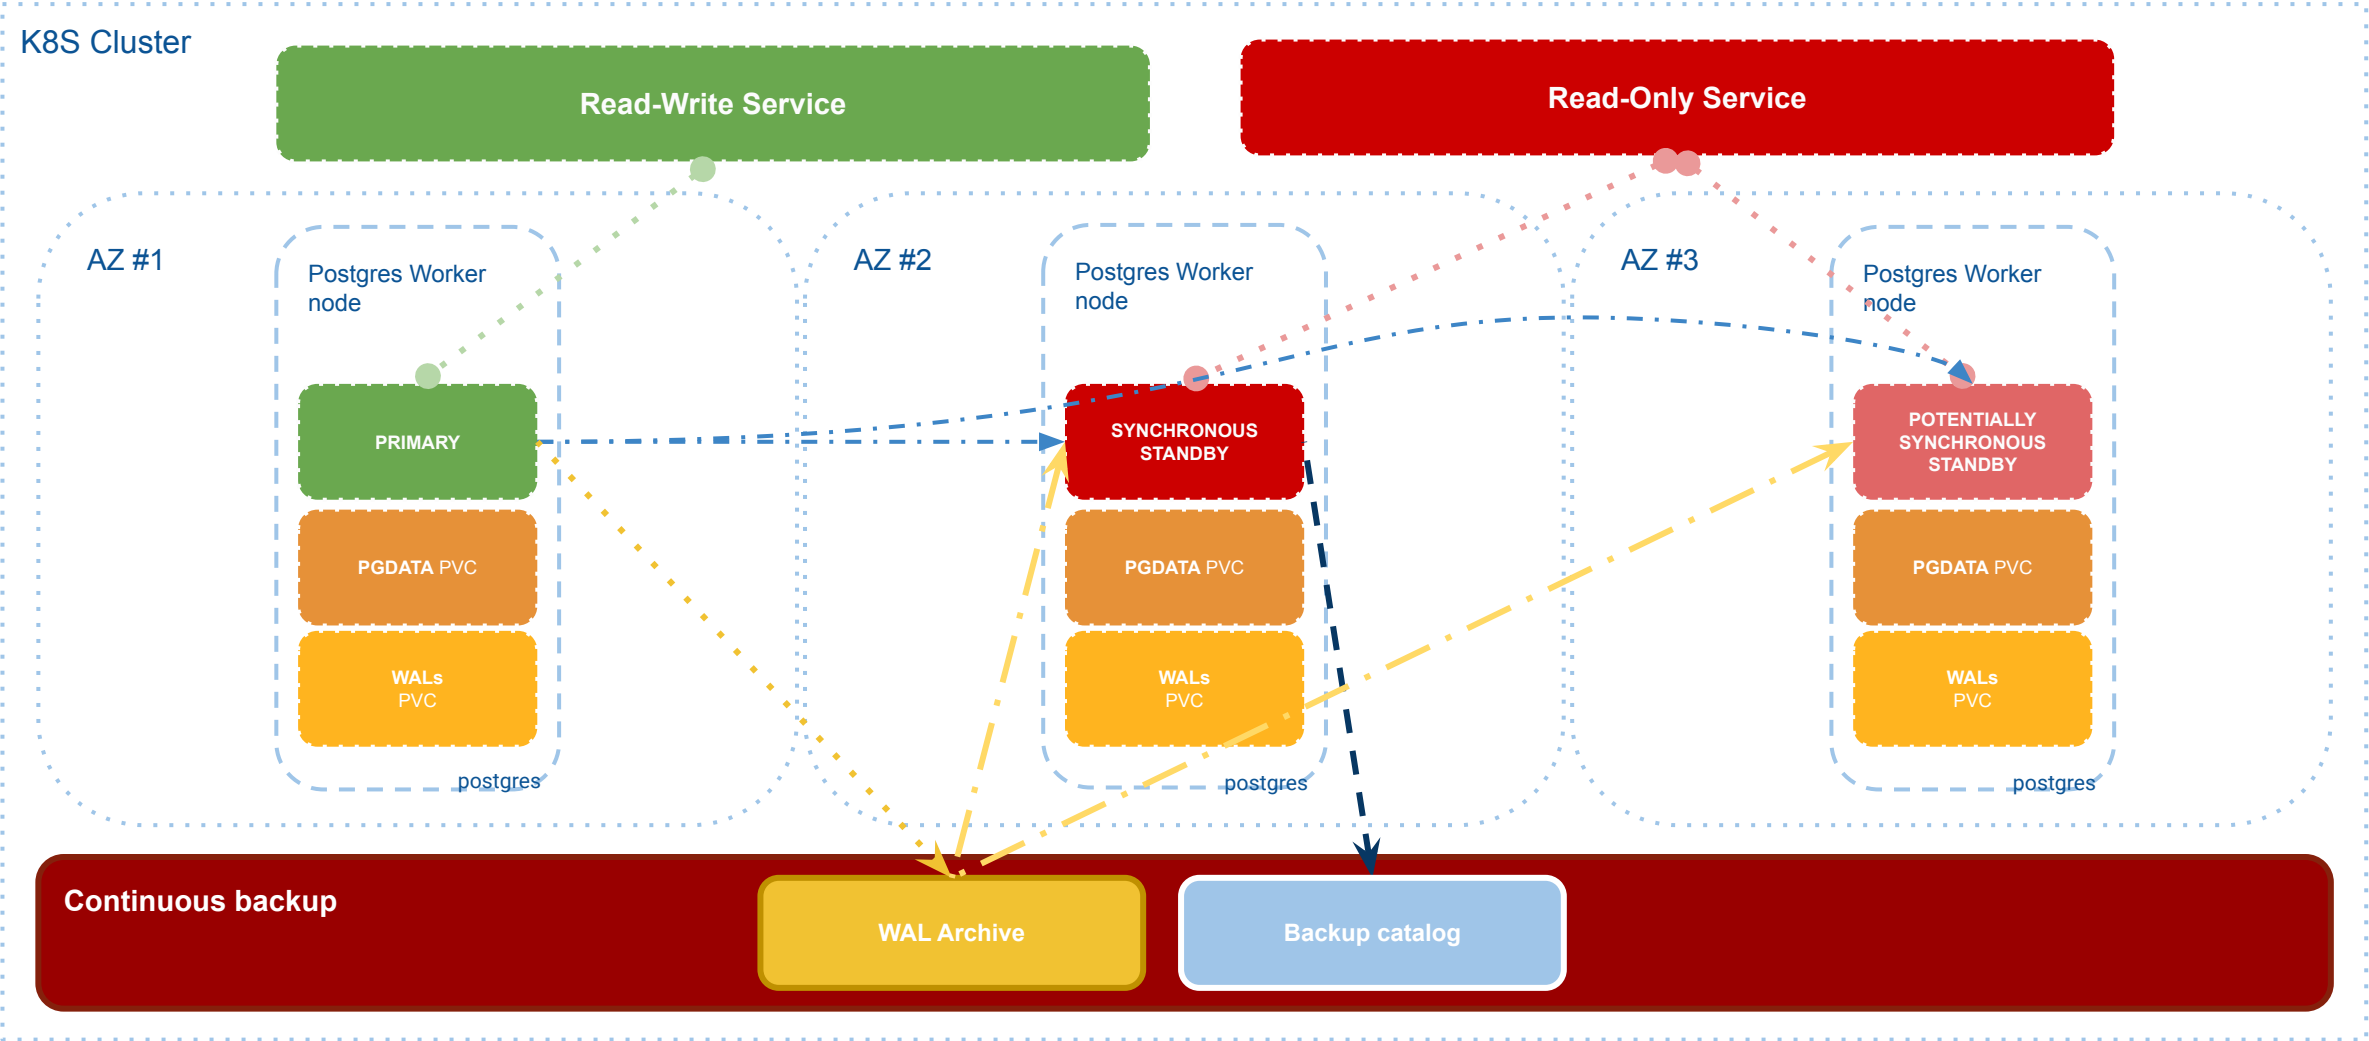




# Architectures



# Production Cluster with 3+ Availability Zones



# Configuration example

```
apiVersion: postgresql.cnpg.io/v1
kind: Cluster
metadata:
  name: cluster1
spec:
  instances: 3
  imageName: ghcr.io/cloudnative-pg/postgresql:14.3
```

Number of instances in streaming replica

Postgres version

```
bootstrap:
  initdb:
    database: app
    owner: app-own
    secret:
      name: app-secret
```

```
storage:
  size: 100Gi
```

```
monitoring:
  enablePodMonitor: true
```

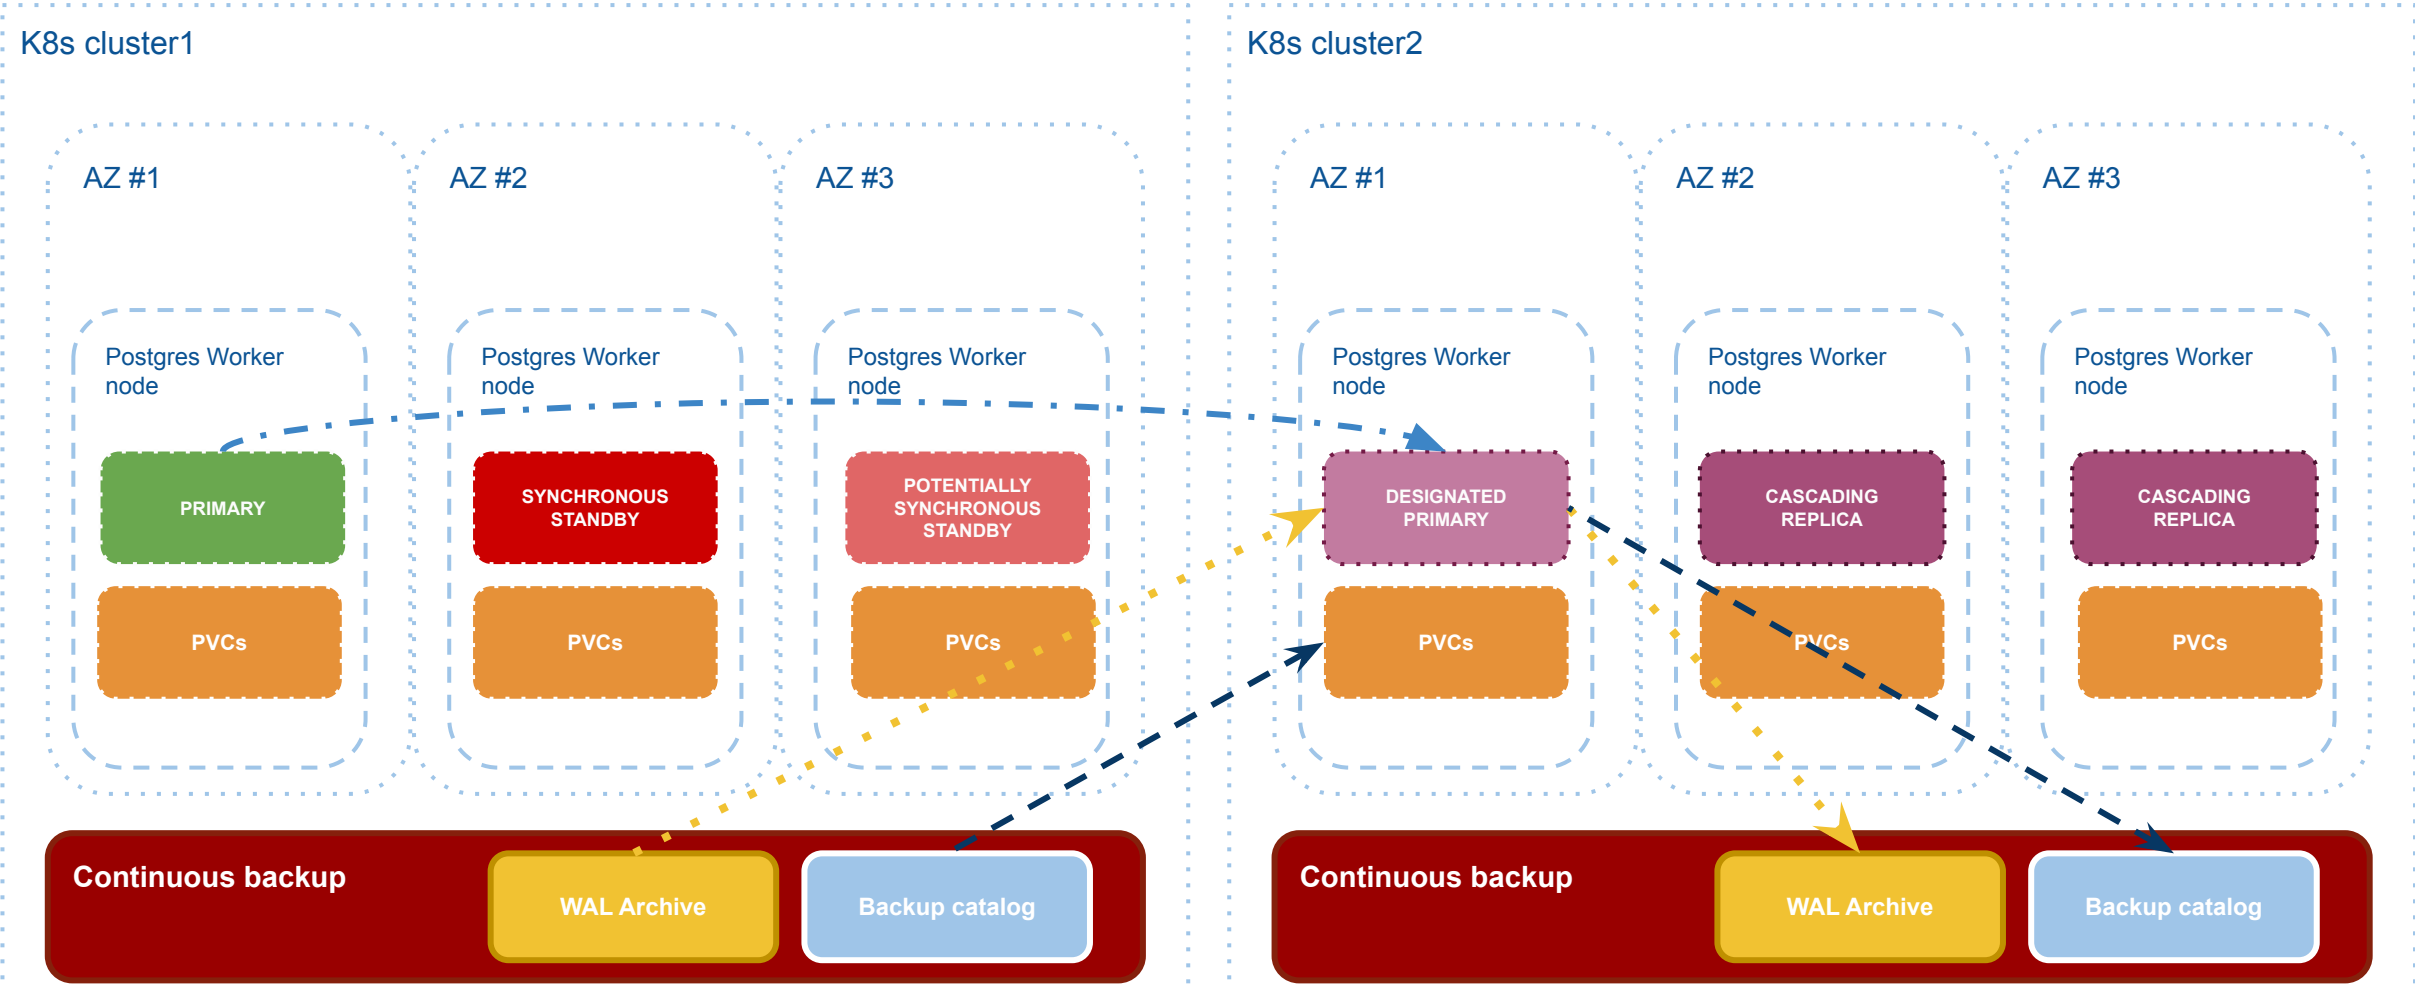
```
backup:
  barmanObjectStore:
    destinationPath: "s3://cloudnativepg/"
    endpointURL: "http://192.168.1.121:9000"
    s3Credentials:
      accessKeyId:
        name: minio-creds
        key: MINIO_ACCESS_KEY
      secretAccessKey:
        name: minio-creds
        key: MINIO_SECRET_KEY
    data:
      immediateCheckpoint: true
    retentionPolicy: "1w"
```

Initialize database (new)

Monitoring to prometheus

Barman backup repository

# Symmetric Architecture on 2 different K8S clusters



**Primary + DR, RPO=0, RTO=manual, 100% declarative configuration**



# Configuration example

```
apiVersion: postgresql.cnpg.io/v1
kind: Cluster
metadata:
  name: cluster1
spec:
  instances: 3
  imageName: ghcr.io/cloudnative-pg/postgresql:14.3
```

Number of instances in streaming replica

Postgres version

```
bootstrap:
  initdb:
    database: app
    owner: app-own
    secret:
      name: app-secret
  backup:
    name: backup-PIT
  replica:
    enabled: true
    source: cluster2
```

```
storage:
  size: 100Gi
```

```
monitoring:
  enablePodMonitor: true
```

```
backup:
  barmanObjectStore:
    destinationPath: "s3://cloudnativepg/"
    endpointURL: "http://192.168.1.121:9000"
    s3Credentials:
      accessKeyId:
        name: minio-creds
        key: MINIO_ACCESS_KEY
      secretAccessKey:
        name: minio-creds
        key: MINIO_SECRET_KEY
    data:
      immediateCheckpoint: true
      retentionPolicy: "1w"
```

Initialize database (as replica from)

Monitoring to prometheus

Barman repository

# Unlocking Insights for Growth

EDB Accelerators for Analytics and Gen AI





# EDB for Analytics & Gen AI agenda

- Business Drivers and Architecture Overview
- PG Generative AI Demo and Highlights
- PG Analytics Demo and Highlights



# What is the GAP

Drivers, Requirements, challenges, pains, future state

## GAP Identification

**Current State**

- Digital Transformation
- Data Strategy
- Open Source Adoption
- Move to the Cloud/Hybrid Cloud
- CFCN & DevOps
- Migrate from legacy DB (e.g. cost optimization)

- **Analytics, Gen AI**
- Always-On / HA (Business Critical Apps)
- Performance and scalability
- Security Compliance
- Backup & Recovery

**Desired Future State**

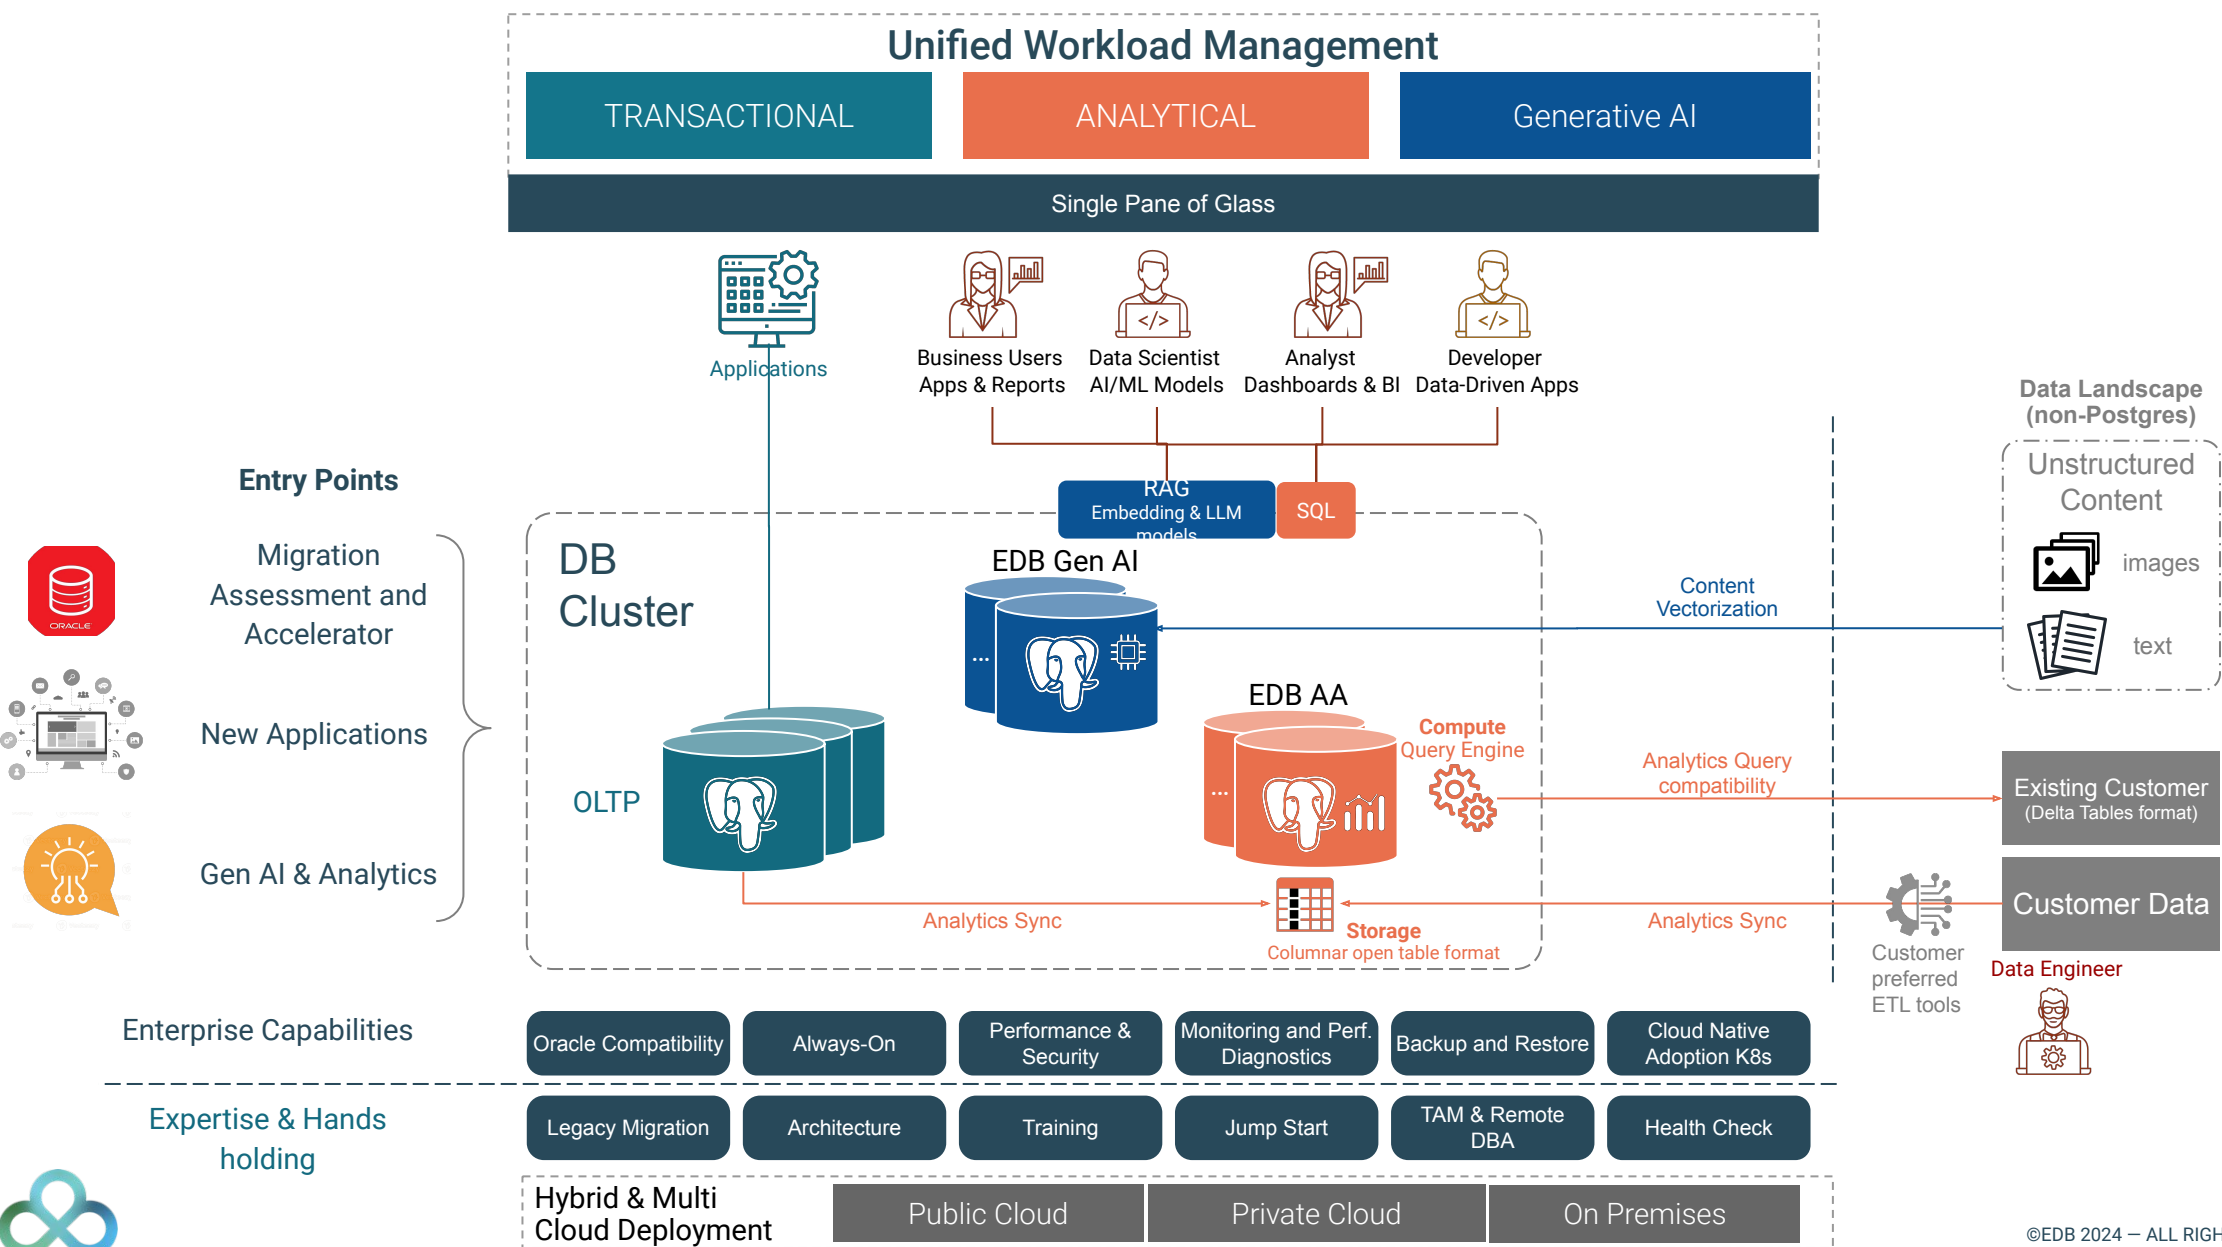
This is how we can connect EDB Postgres AI platform capabilities + Expertise to bridge the GAP

**EDB Postgres AI  
Platform capabilities**

**EDB Professional  
services Expertise**



# Architecture Overview - The Big Picture



# PG Generative AI



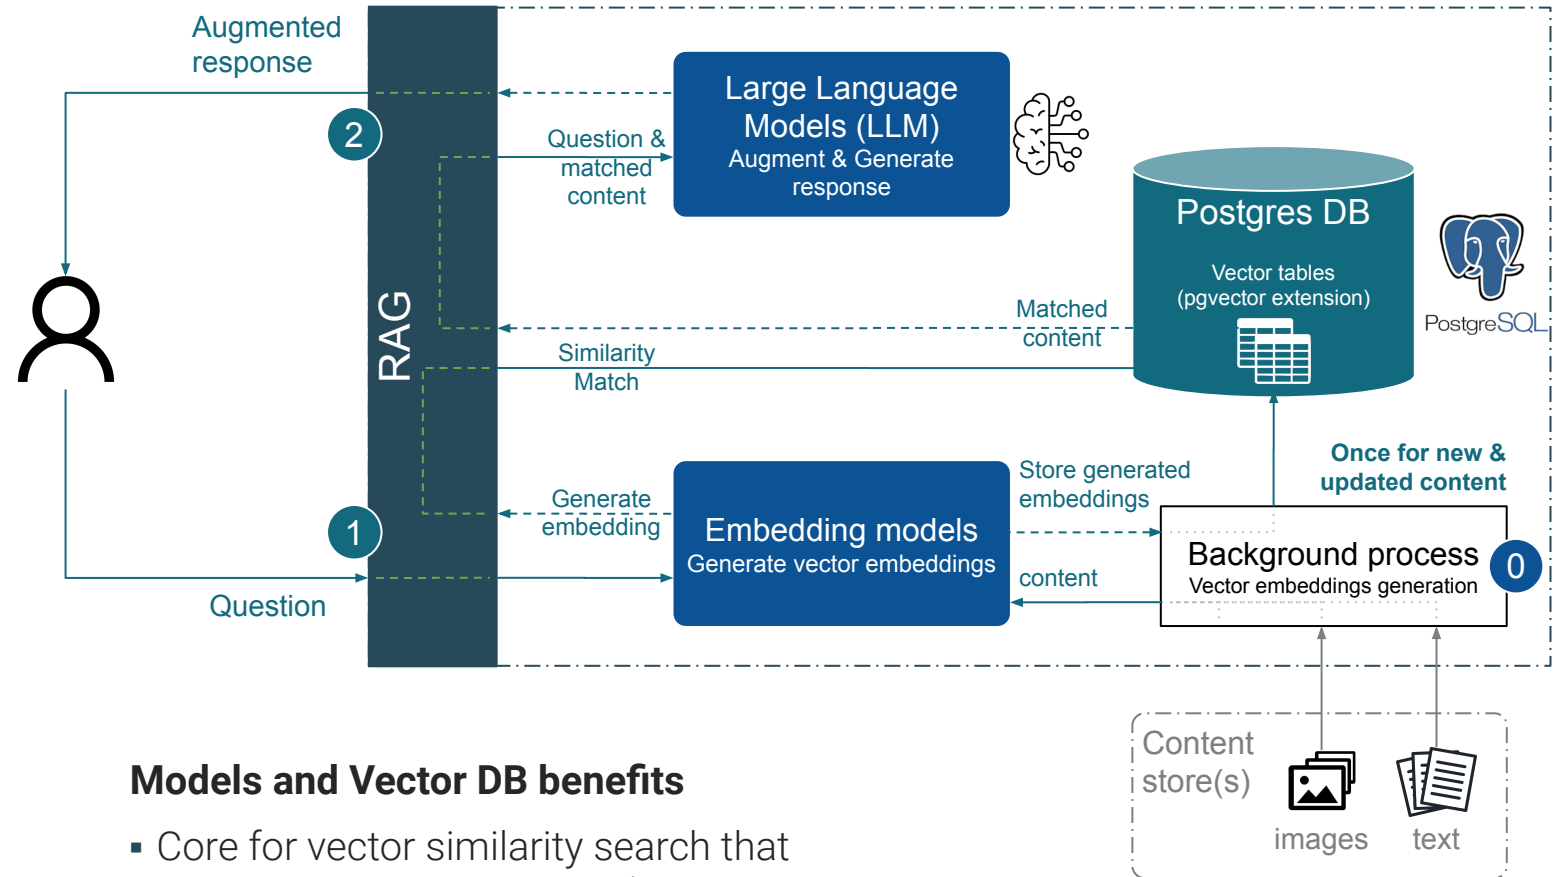
# RAG (**R**etrieval **A**ugmented **G**eneration) flow

## Use Cases

- Conversational Assistance
- Chatbot
- Semantic search
- Visual and Semantic search for images
- Real-time and personalized search experience
- ...

## Content examples

- Documents (contracts, user guides, product/technical documentation, policies, CVs...)
- Intranet documents (e.g. sharepoint, ...)
- ....

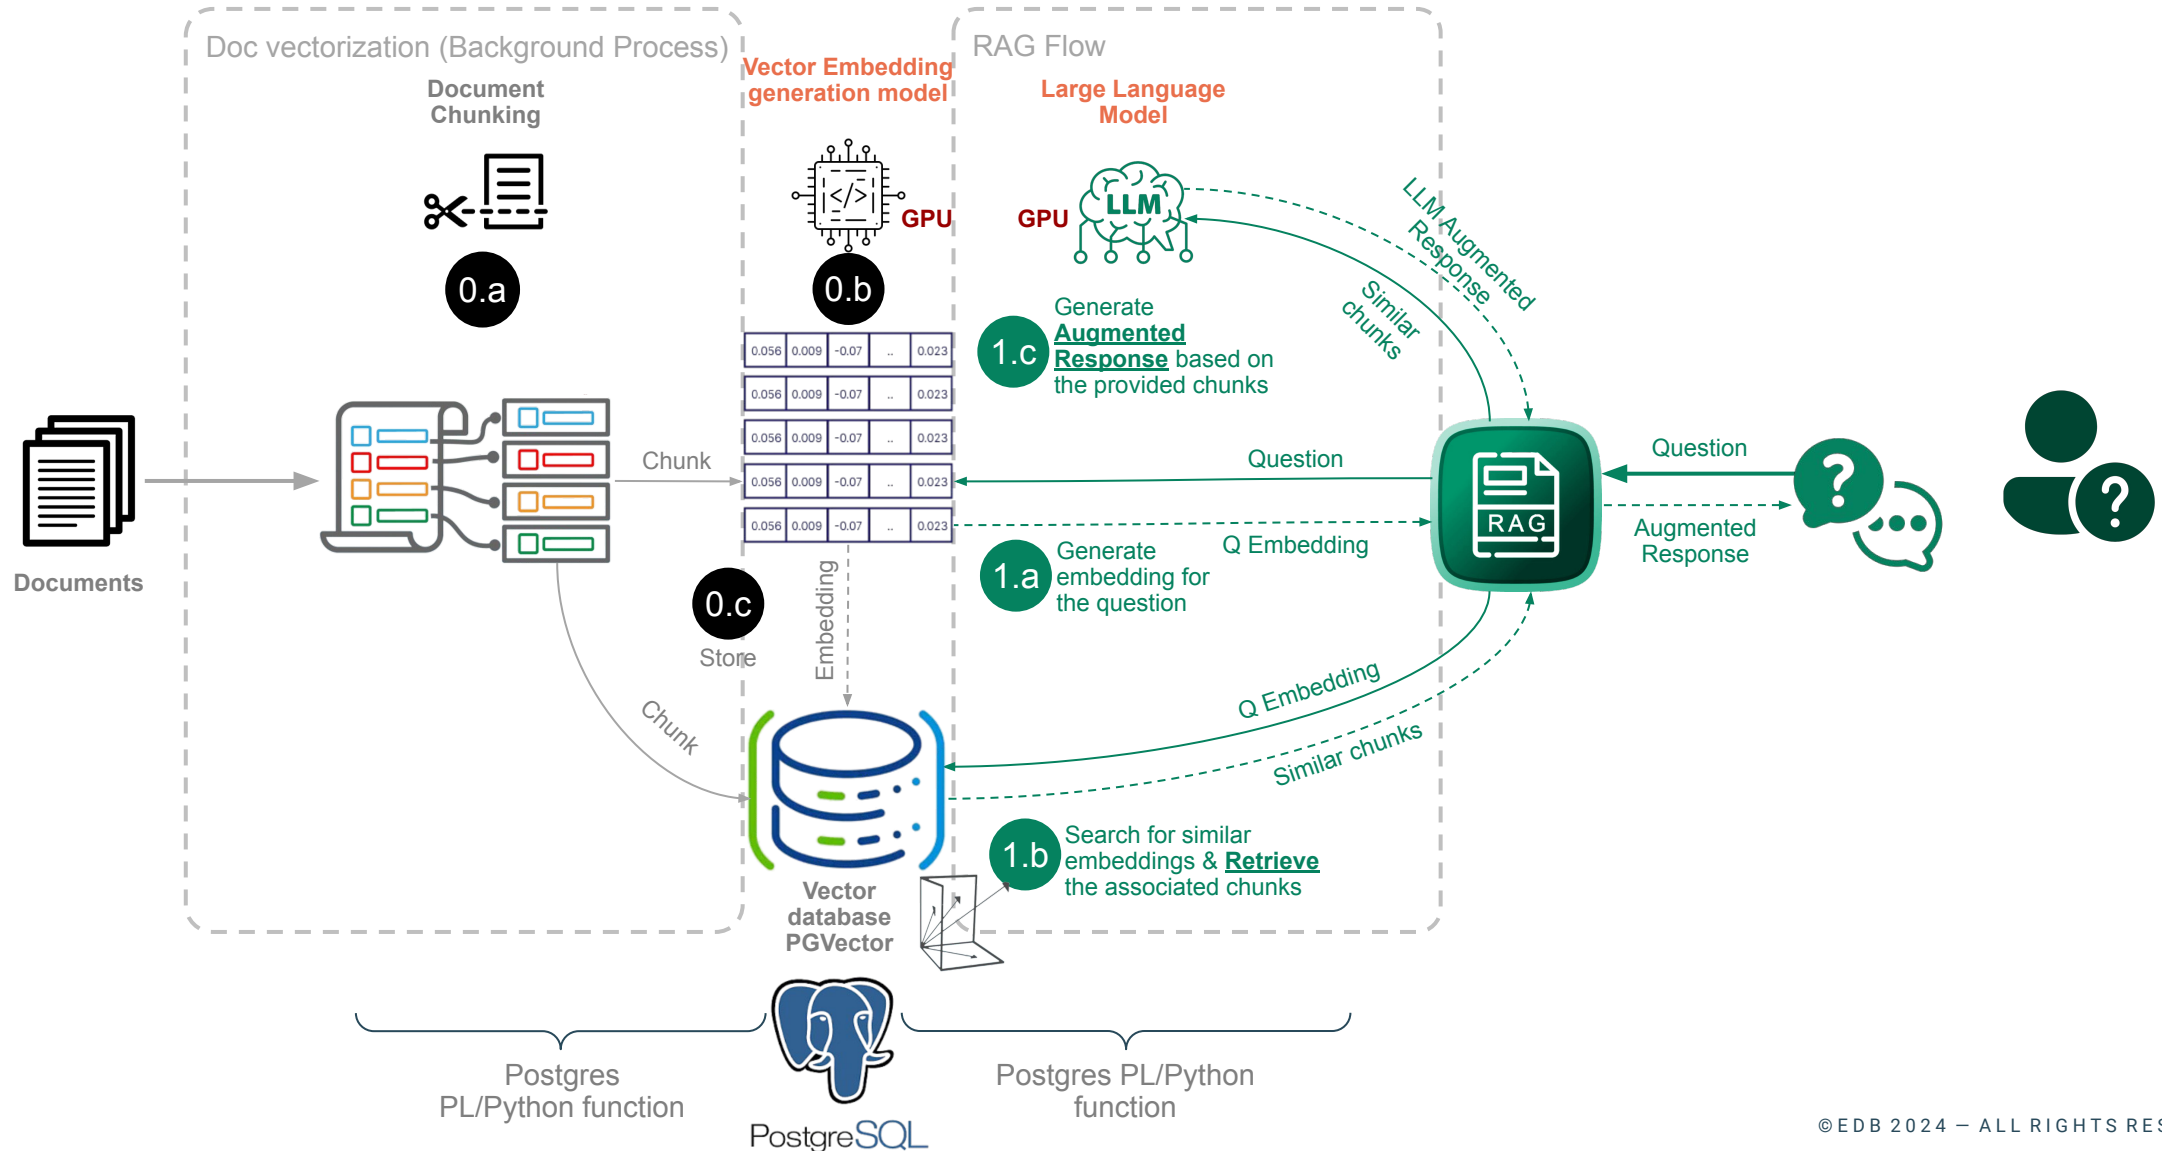


## Models and Vector DB benefits

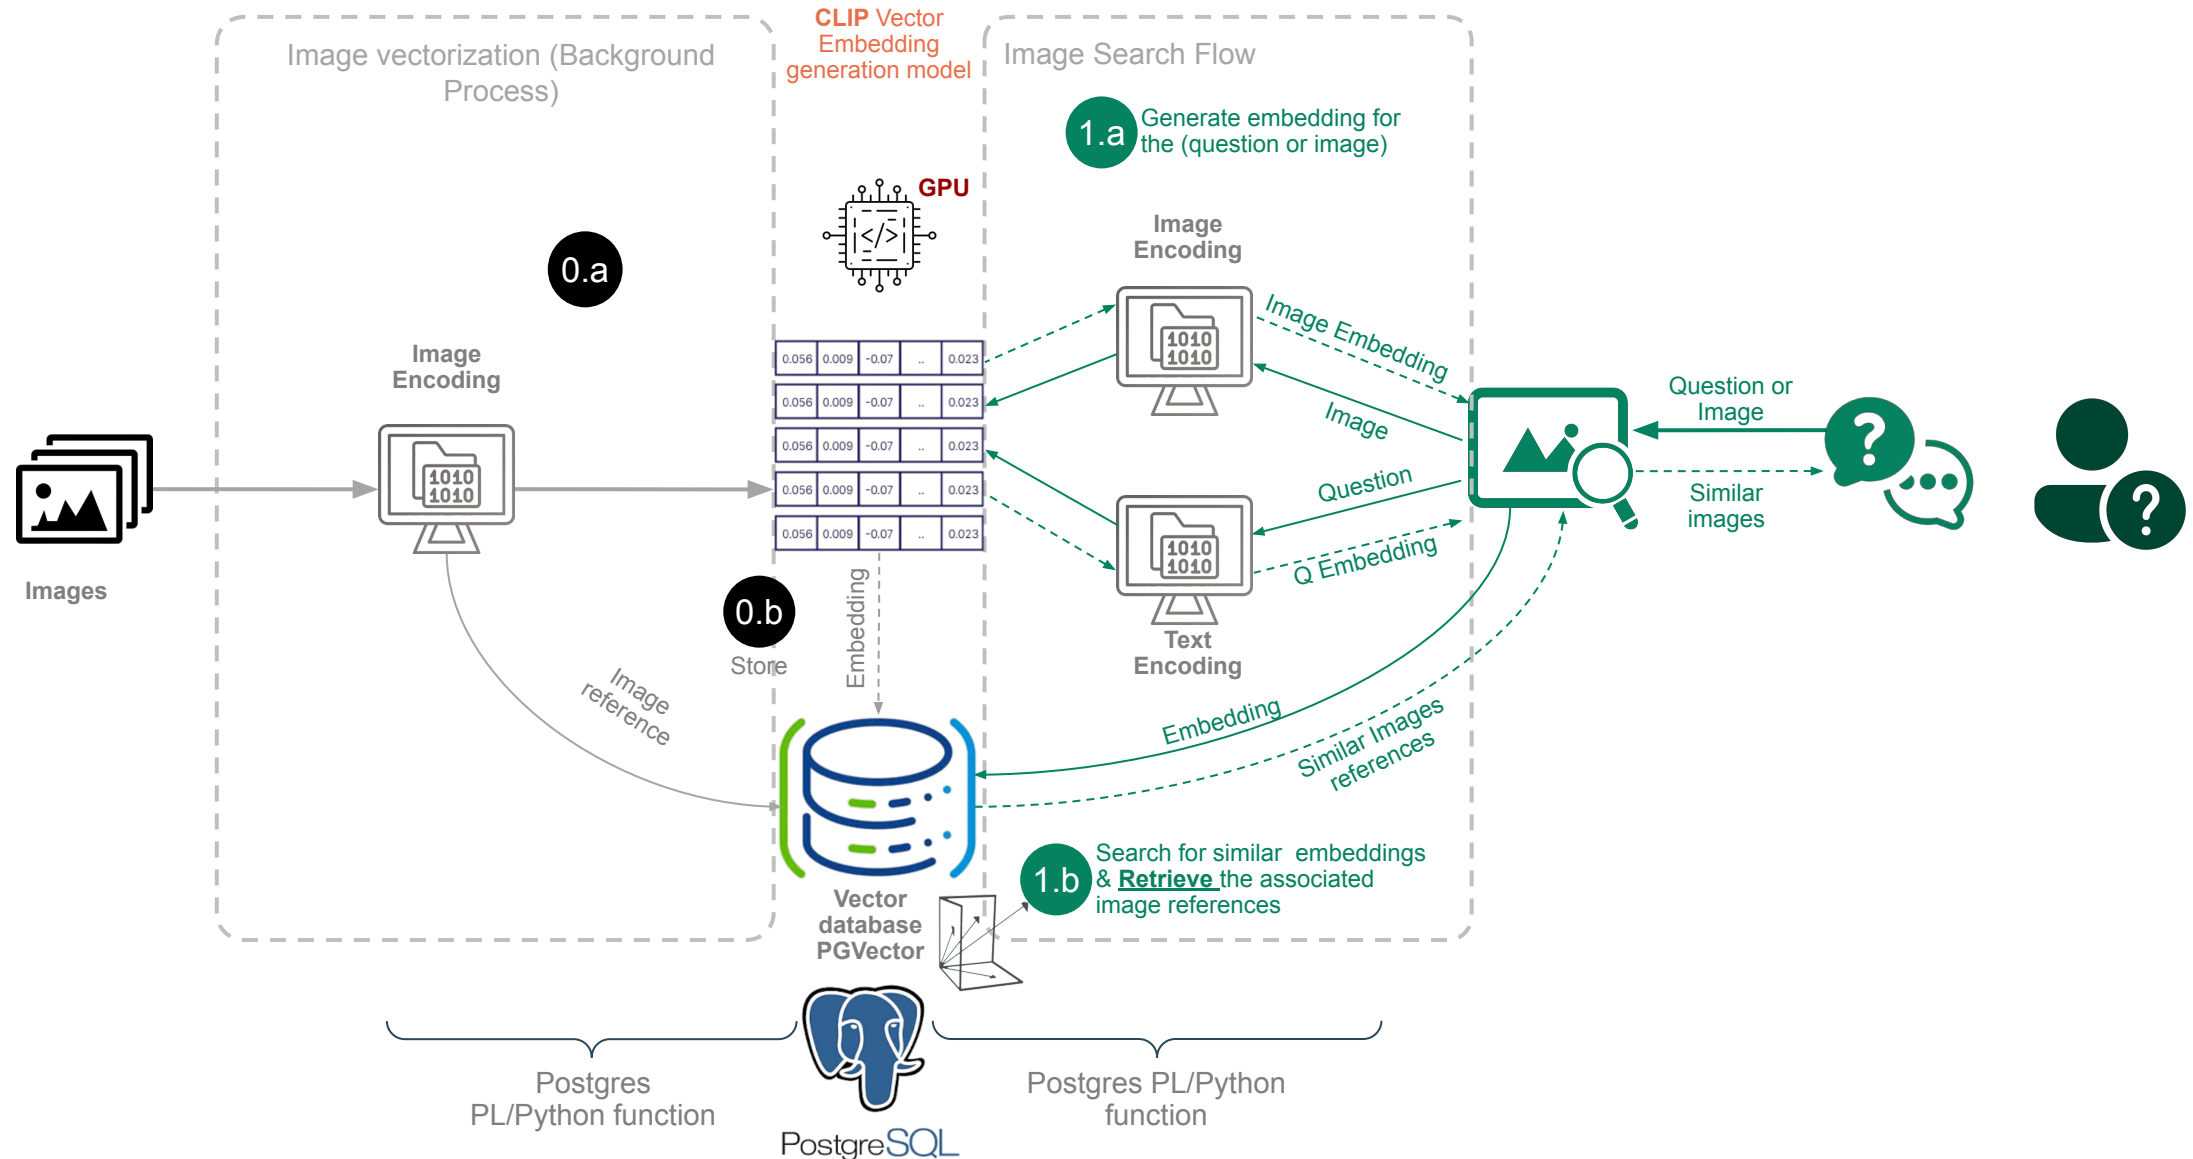
- Core for vector similarity search that provides fast and scalable (with improved accuracy) experience



# RAG Interactions - detailed flow



# Image Search Interactions - detailed flow





# Ideation Workshop

High Value Low Complexity

Use Case / MVP



# PG Analytics Accelerator



# Key Business Challenges Addressed by PGAA

- Separates storage from compute, and includes support for S3 compatible storage, local FS
- Makes open Tables format queryable just like any Postgres table
- Analytical Query Engine for up to 30x faster queries compared to transactional Postgres

## PGAA + PGD (Real-time analytics)

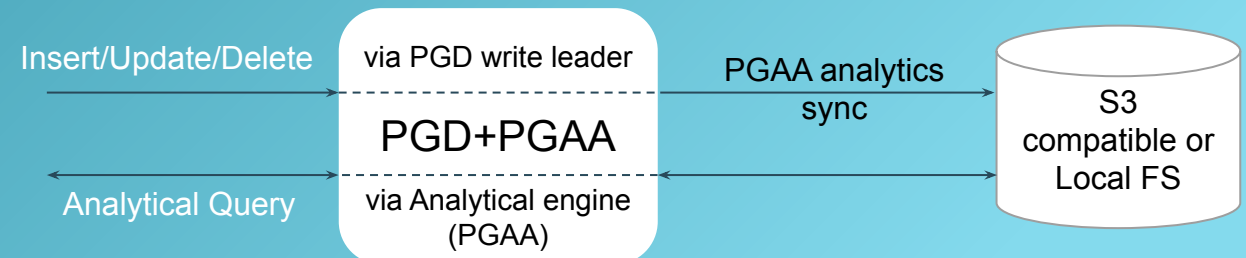
- Replicate transactional tables directly to open tables format (PGD)
- Query tables with the analytical engine (PGD)

## High-Performance Analytics:

Extremely Fast Queries

**Cost-Effective Solution:** Lower  
Ownerships Costs & Ease of Use

**Real-Time Data Insights:** Efficient  
Data Processing & Open Data Formats



Reduced total cost of ownership by decoupling compute and storage.

**30X**  
faster

on *average* for analytical queries compared to Postgres

**5X**  
smaller

on disk open tables format vs. Postgres tables and indexes

**18X**  
cost-effective

Object storage vs. solid state drives (SSDs)





**EDB**

Postgres® for the AI Generation

# Customer Testimonial







# Da Oracle a PostgreSQL

COME AFFRONTARE E SUPERARE SFIDE COMPLESSE

Marco Console - Global Head of DBAs



# Superare le Sfide: da Oracle a PostgreSQL con Successo

"Il successo nasce dalla collaborazione con i partner giusti, un team affiatato e una visione condivisa."

*Marco Console*

Il nostro obiettivo, estremamente sfidante per il team, è stato migrare con successo una complessa infrastruttura da Oracle a PostgreSQL per la più grande PayTech Europea.

- Reazioni iniziali del team: un mix di shock e motivazione.
- Analisi della complessità dell'infrastruttura esistente.







# Analisi dell'AS-IS

Continuità operativa e alta affidabilità

L'infrastruttura originaria basata su Oracle includeva diverse componenti critiche, come la gestione dei dati tramite cluster di database distribuiti tra più data center, con una necessità di garantire continuità operativa e alte prestazioni

# Il Potere del Team Giusto

Il valore di avere il personale tecnico adeguato.

Il team con competenza e dedizione ha affrontato sfide legate alla configurazione delle risorse, come l'allocazione delle VM su base ambiente ovvero (PROD, PERF, UAT, SIT), garantendo che ogni ambiente fosse dimensionato correttamente in termini di CPU, RAM e spazio disco per una migrazione fluida e sicura,



# I Partner Giusti fanno la Differenza

Un support a 360°

La collaborazione con partner specializzati, come EDB, ha reso possibile un'analisi dettagliata delle prestazioni PostgreSQL, inclusa la configurazione di strumenti come PgBouncer per migliorare la gestione delle connessioni e ottimizzare le prestazioni del sistema sotto carichi intensi. Inoltre, devo riconoscere che il team di Broway, società svizzera composta da professionisti altamente qualificati, ha rappresentato un fattore determinante nel successo del progetto.





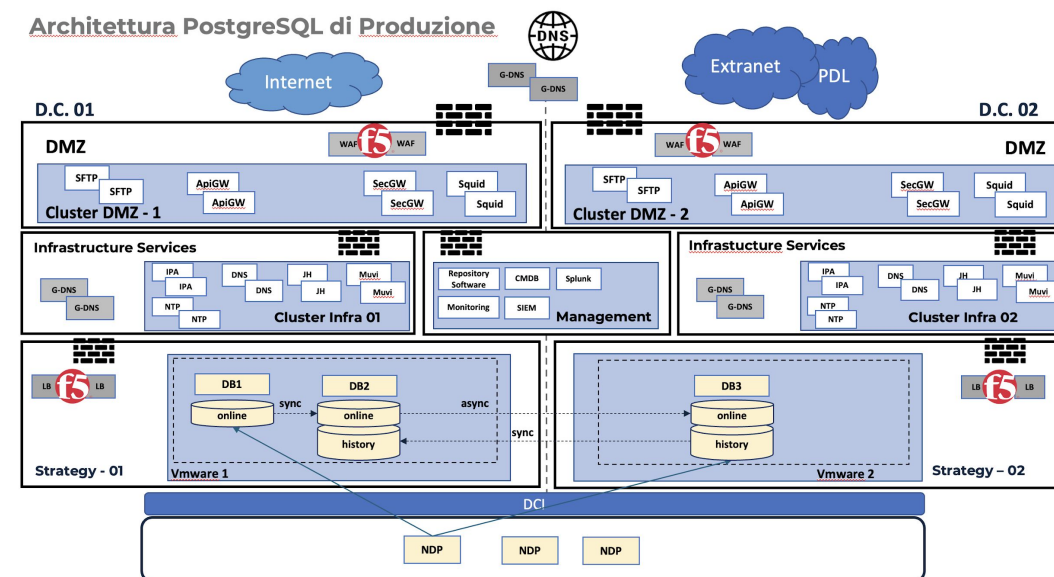


# La Migrazione: Un Progetto altamente Sfidante

La transizione da Oracle a PostgreSQL non è semplice.

Durante la migrazione, uno dei principali ostacoli era l'inefficace pruning delle partizioni dovuto alla struttura SQL generata dalle applicazioni. Questo problema è stato risolto ottimizzando le query con l'inserimento di colonne di partizione per tutte le tabelle coinvolte, migliorando drasticamente la performance

## Architettura PostgreSQL di Produzione



# Esigenza di Migrazione

Migrare da Oracle a PostgreSQL un'importante e critica piattaforma di pagamenti.



Requisiti da soddisfare:

01. **Aderenza alle compliance:** Necessità di aderire agli standard PCI-DSS per garantire la sicurezza delle transazioni.
02. **Performance e Alta Affidabilità:** Garantire prestazioni elevate per gestire un carico intensivo di transazioni.
03. **Disaster Recovery con RPO=zero:** Implementare un sistema di disaster recovery che assicuri la continuità dei dati senza perdita (Recovery Point Objective = 0).

# Problemi Iniziali e Superamento delle Difficoltà

Approccio collaborativo per risolvere ogni problema.



## Problematiche Specifiche

01. Tuning delle applicazioni
02. Compatibilità del codice
03. Innumerevoli test di validazione

L'uso di `plan_cache_mode` e la gestione delle sessioni con **PgBouncer** sono stati fondamentali per risolvere i problemi di performance. Sono stati effettuati numerosi test per validare le configurazioni e ridurre i colli di bottiglia legati al caching dei piani di esecuzione



# Le Sfide Incontrate

Abbracciare il futuro, guidare il cambiamento.



## Esperienza e Evoluzione

01. **Competenze interne:** Team con forte esperienza su Oracle, ma con competenze in PostgreSQL da sviluppare ulteriormente
02. **Fiducia nel Cambiamento:** Superare le esitazioni iniziali nel lasciare una tecnologia consolidata come Oracle per approdare a PostgreSQL.



# Come abbiamo Vinto le Sfide

Dedizione, Esperienza e Nuove Competenze

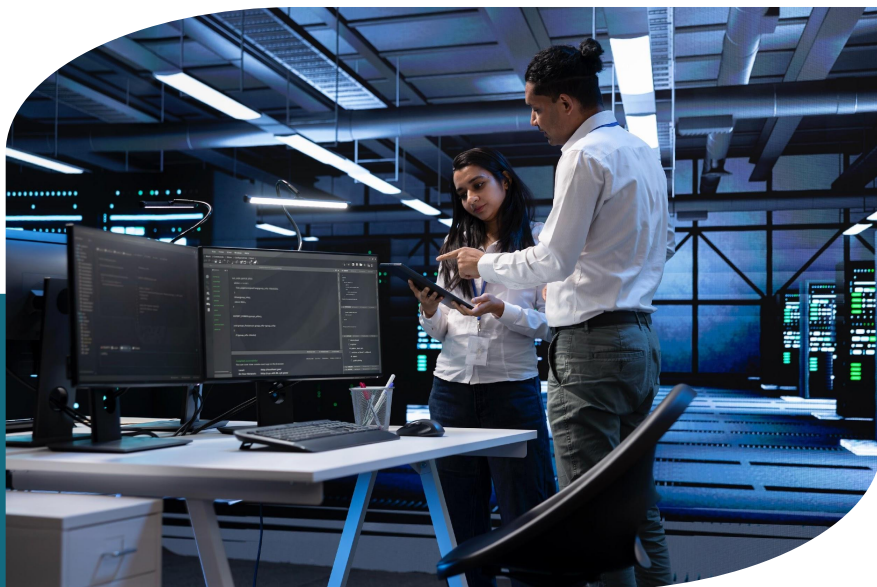


## Impegno e Innovazione

01. **Disponibilità del Team EDB:** Supporto continuo e dedizione del team EDB.
02. **Competenze delle Persone Coinvolte:** Analisi dei requisiti e design dell'architettura basati su esperienze pregresse.
03. **Voglia di Imparare:** Partecipazione del team al corso EDB per Oracle DBAs, con la volontà di acquisire nuove competenze su PostgreSQL.
04. **Maturità di EDB Postgres Advanced Server:** La compatibilità avanzata con Oracle di EDB Postgres ha facilitato in modo davvero unico la transizione e soprattutto senza interruzioni.

# Il Ruolo di EDB nel Progetto

Il pilastro tecnologico per una migrazione sicura e performante



## Pilastro Tecnologico

01. **Ruolo Tecnologico:** Utilizzo di **EDB Postgres Advanced Server** per garantire un'infrastruttura affidabile e performante.
02. **Ruolo Consulenziale:** Supporto nell'**Architectural Design** per Alta Affidabilità (HA), Disaster Recovery (DR) e ottimizzazione delle performance.
03. **Ruolo di Partner:** Supporto 24/7 per garantire il successo della migrazione e delle operazioni post-migrazione.

# Spirito di Collaborazione come Chiave del Successo

La fiducia e la comunicazione aperta tra partner e team interno.

La collaborazione tra team di sviluppo, amministratori di database e partner esterni ha permesso di raggiungere una sintonia perfetta nell'ottimizzazione delle query, con un particolare focus sulla riduzione della latenza e l'aumento delle transazioni per secondo (TPS)



Fiducia



Comunicazione



Collaborazione

# Risultato Finale: Migrazione con Successo

Infrastruttura migrata, performance migliorate

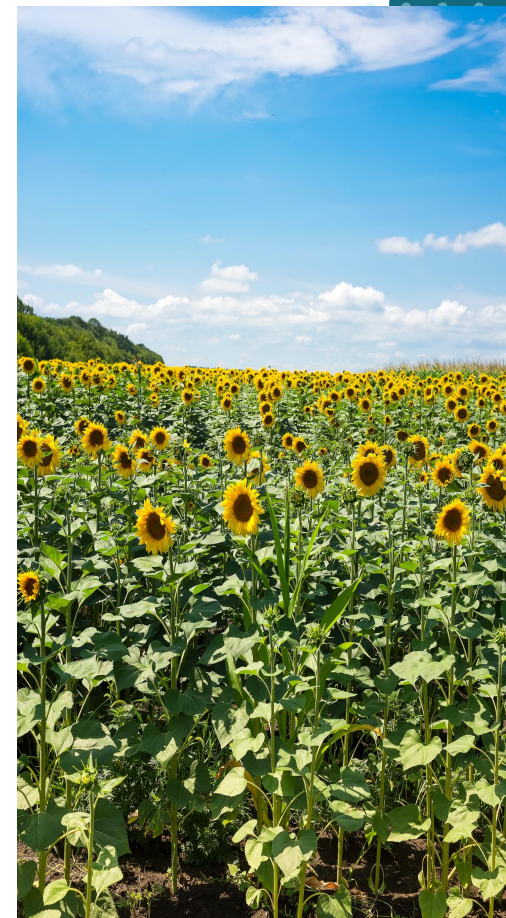
I benchmark eseguiti successivamente alle attività di tuning post-migrazione hanno permesso di misurare prestazioni mai inferiori a **5104 TPS** e una latenza di **0,979 ms**, grazie alle ottimizzazioni effettuate sulle query SQL e all'uso di PgBouncer.



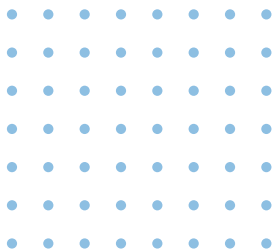
# Benefici a Lungo Termine di PostgreSQL

I vantaggi in termini di efficienza, costi ridotti e flessibilità.

PostgreSQL offre un'infrastruttura scalabile e flessibile, con una gestione più semplice delle partizioni e una migliore ottimizzazione delle risorse, garantendo **costi operativi ridotti** e un'**efficienza superiore** rispetto all'ambiente Oracle precedente







### DETTAGLI DELL' EVENTO

Email: [marco.cesare@enterprisedb.com](mailto:marco.cesare@enterprisedb.com)

Phone: +39 335 6519644

Website: [www.enterprisedb.com](http://www.enterprisedb.com)

# Ringraziamenti

Un grazie al team, ai partner e ai partecipanti.

Il successo della migrazione è stato reso possibile grazie all'adozione di un approccio iterativo, alla continua collaborazione e all'uso di strumenti avanzati per il monitoraggio e l'ottimizzazione delle prestazioni.





**EDB**

Postgres® for the AI Generation

# Customer Testimonial





# Agenda

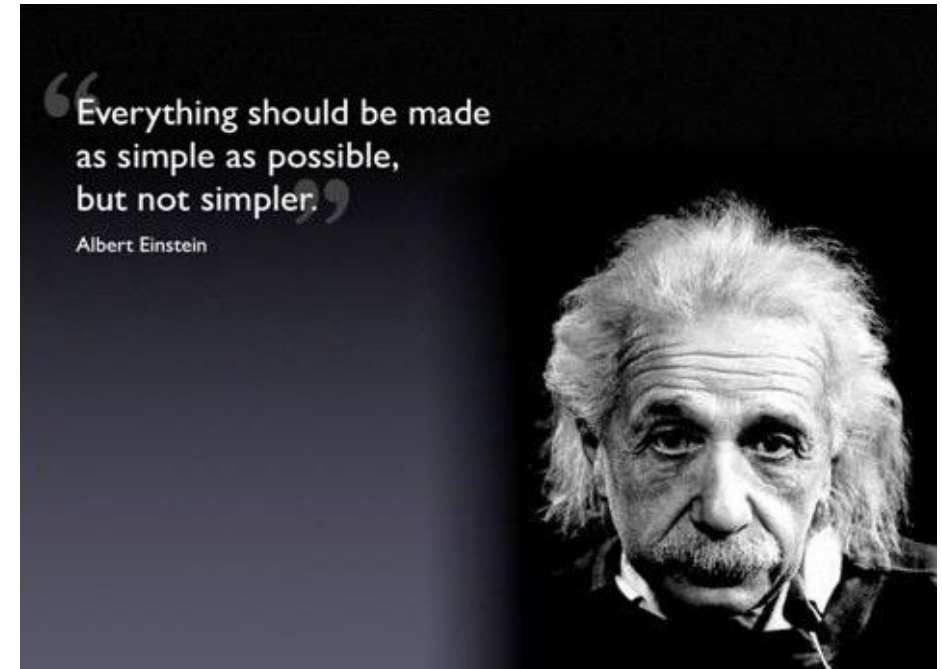


- ✓ **Dove eravamo?**
- ✓ **«Res Gestae» : dove siamo?**
- ✓ **Simplicity is the ultimate sophistication**





# Postgres AI Day, Milan



## Dialogare con il presente per (cercare di) interpretarlo

Paolo Sarchini

IT Solutions Architect - Uno Informatica

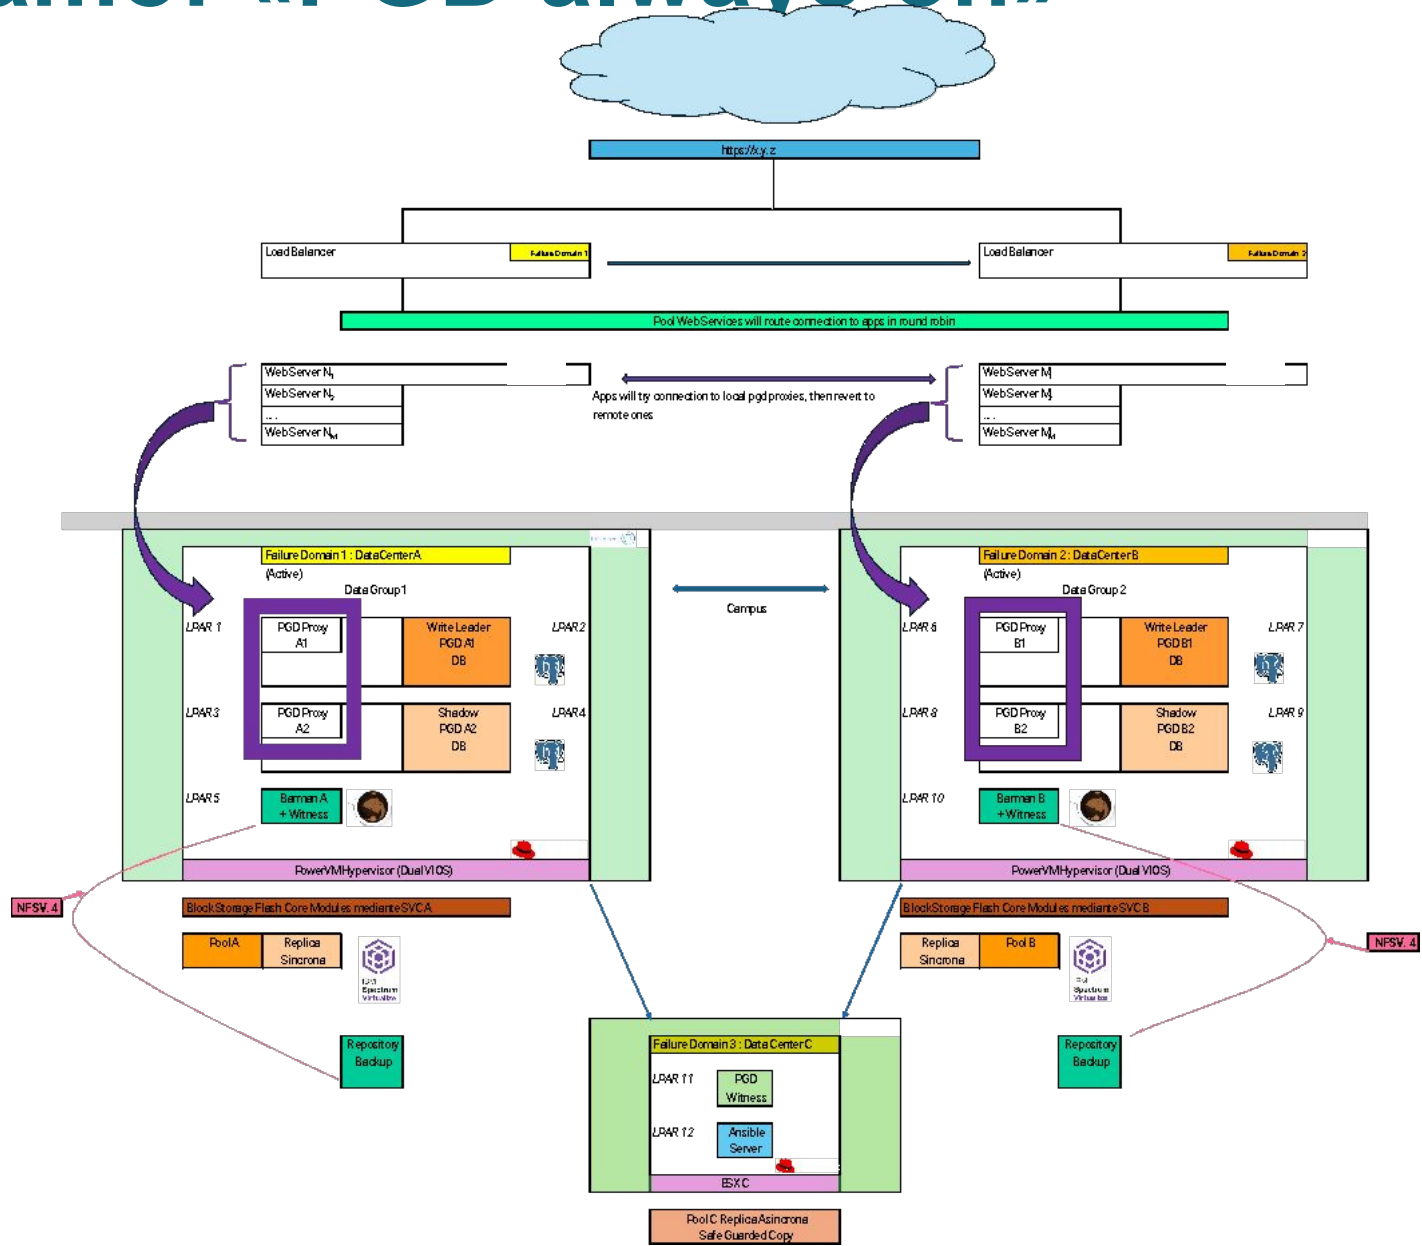


# Ma per andare dove dobbiamo andare... per dove dobbiamo andare?

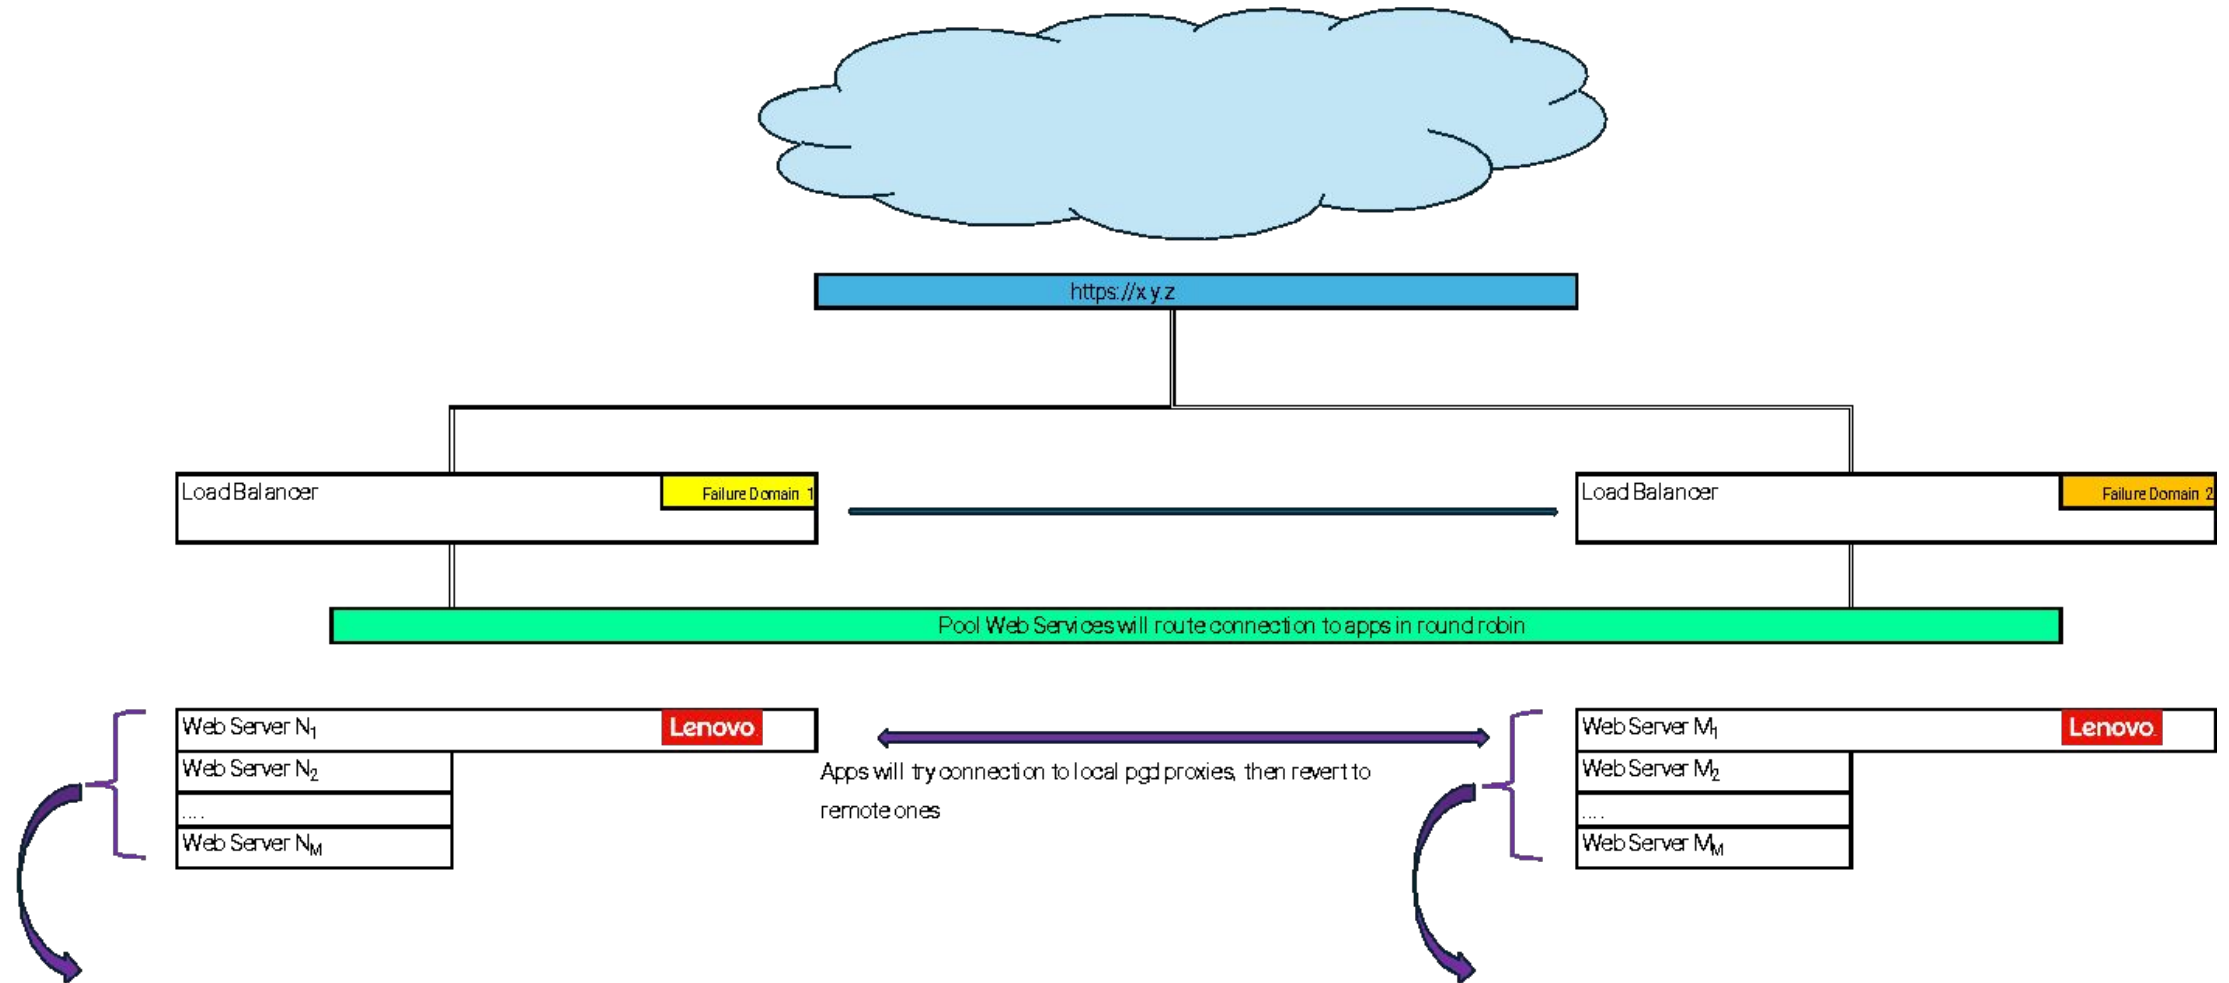
- Dove eravamo? Esigenze di Partenza:
  - Minimizzare i fermi (rolling Upgrades in caso di Major Release)
  - Distribuire le sessioni;
  - Eseguire il Backup a caldo dei sistemi
  - Distribuire i carichi (tra letture e scritture) sui vari nodi
  - Massimizzare l'efficienza operativa
  - (ottimizzazione risorse IT, utilizzando molteplici OS sulla medesima piattaforma



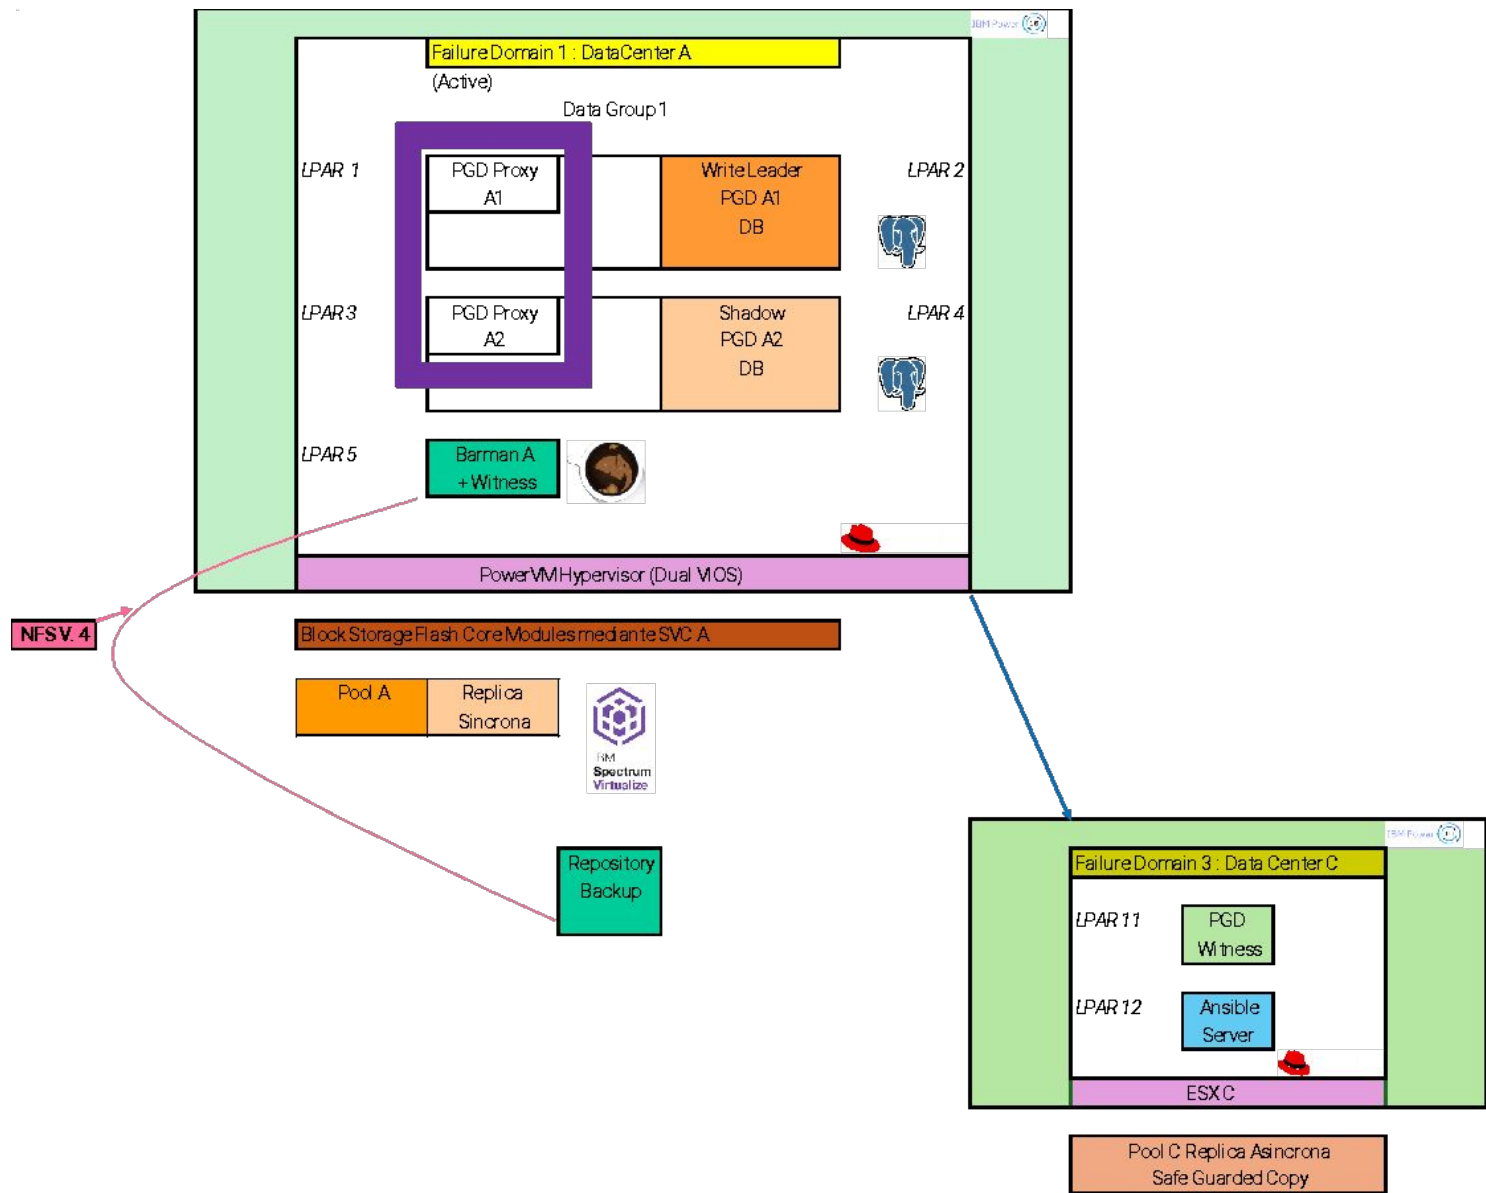
# Dove siamo: «PGD always on»



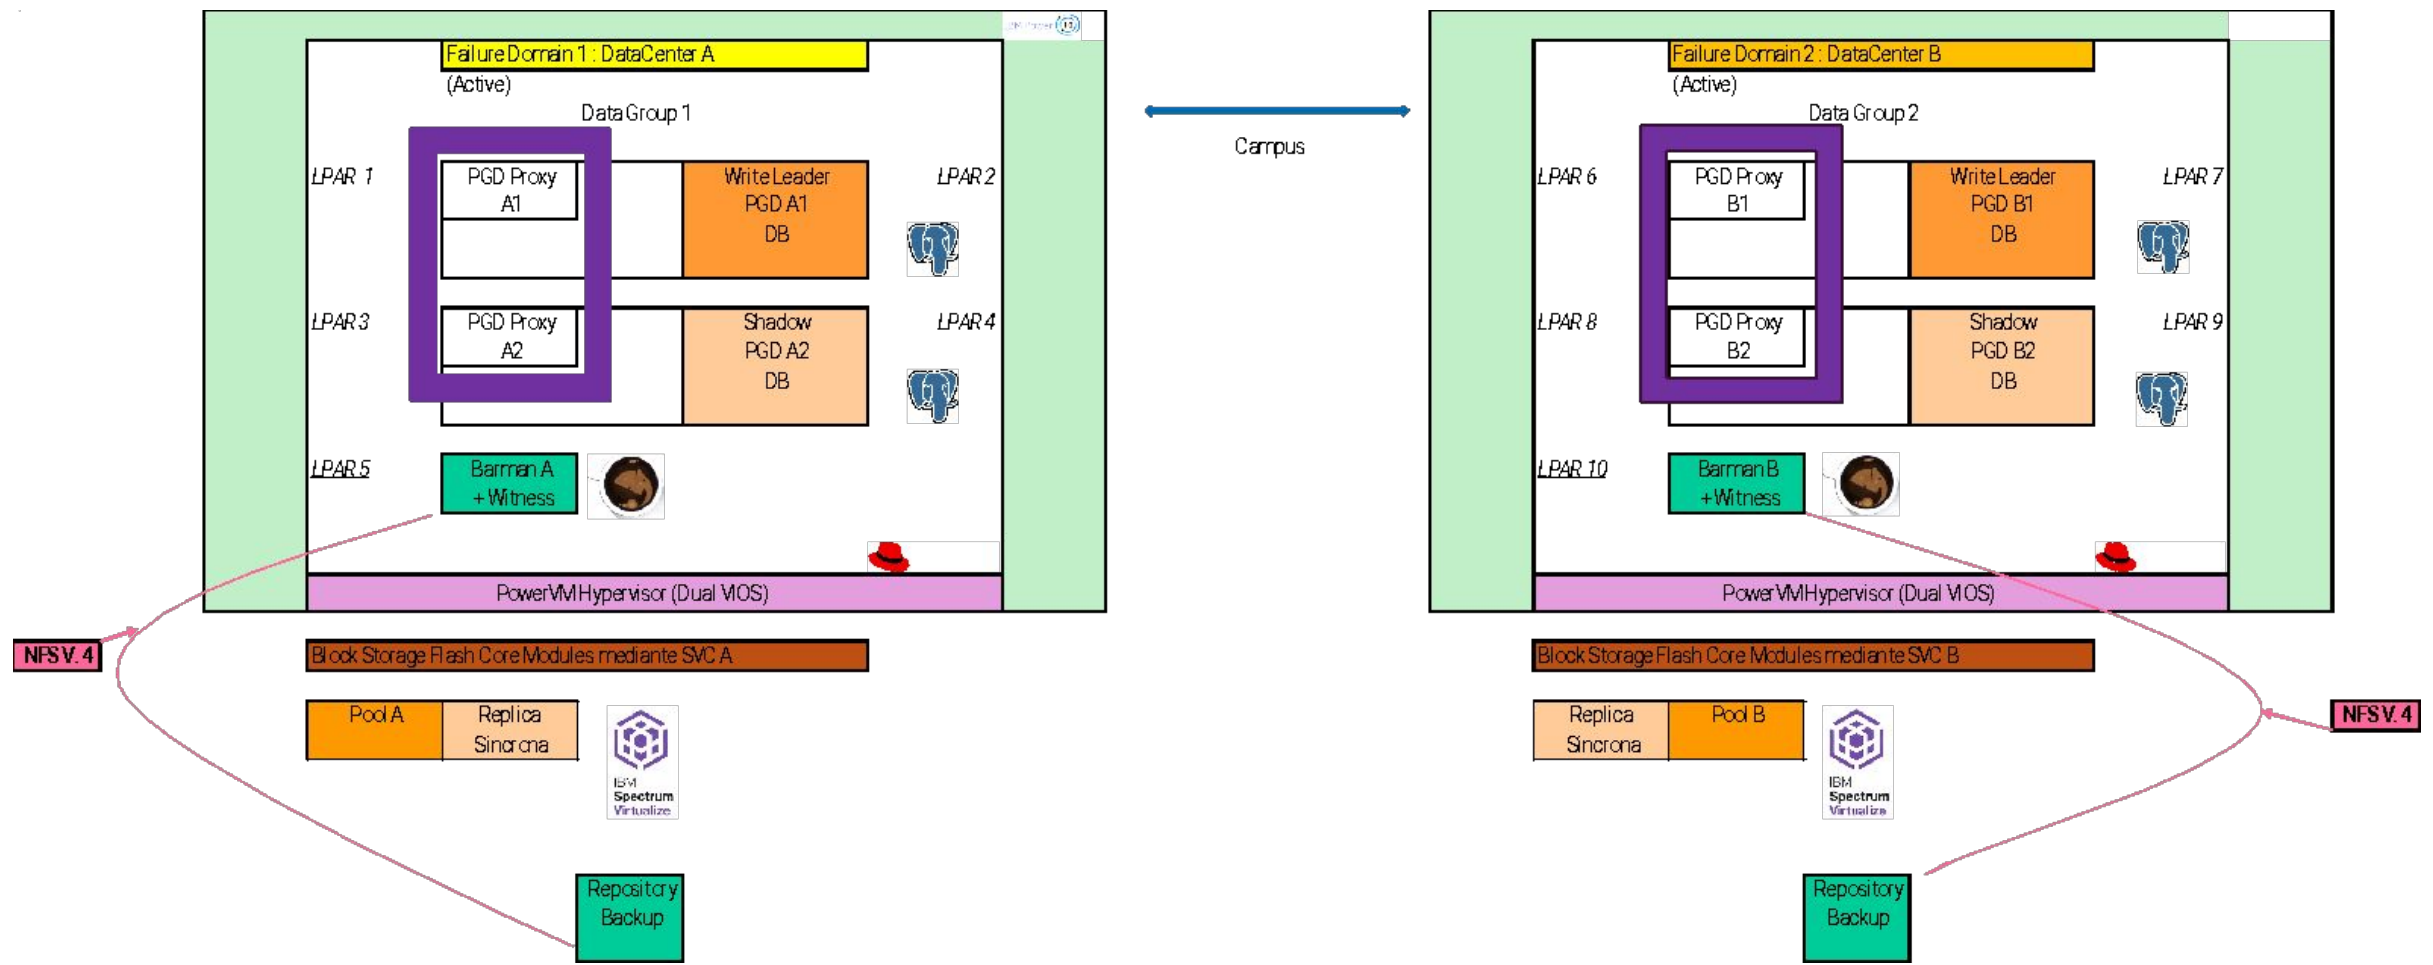
# Dove siamo: «PGD always on»



# Dove siamo: «PGD always on»



# Dove siamo: «PGD always on»





# Simplicity is the ultimate Sophistication : un'Architettura progettata per il futuro



✓ Come altre tecniche occulte di divinazione, il mondo ICT ha un gergo deliberatamente inventato per rendere oscuri i suoi metodi ai non addetti.

- Affidabilità : tasso di disponibilità molto elevato (sondaggio ITIC)  
(<https://www.ibm.com/account/reg/it-it/signup?formid=urx-39584>)
- Performance : gestione di 8 thread in maniera simultanea (SMT8)
- Sicurezza : hardware memory Encryption nativa e coprocessori dedicati per la crittografia
- Tecnologia avanzata per l'I/O : PCIe Gen5 ed Open Memory Interface, SVC, SGC
- Power Enterprise Pool : pay per Use and not pay per System e possibilità di far convivere molteplici OS
- Motore di inferenza ottimizzato: 4 Matrix Math Accelerator (MMA) per ogni core Power 10 consentono l'esecuzione di carichi di lavoro ML ed AI senza risorse aggiuntive (GPU)
- **Sul versante PGD: distribuire i carichi (tra letture e scritture) sui vari nodi**



# Questions?

***Thank you for your attention !***



© 2003 United Feature Syndicate, Inc.

# Grazie per l'attenzione

**Uno Informatica**

Via Puccini, 150 – Arezzo 52100

Tel. 0575 98321

[info@unoinformatica.it](mailto:info@unoinformatica.it)

[www.unoinformatica.it](http://www.unoinformatica.it)



# Beyond the Cloud: EDB's Formula for Maximizing Postgres Value

**Mario Vetrano**

EDB, Cloud Specialist EMEA



# Why? A chat between Mario and the AI



**Mario:**

How many postgres databases are deployed in cloud (any cloud provider) Vs onprem?

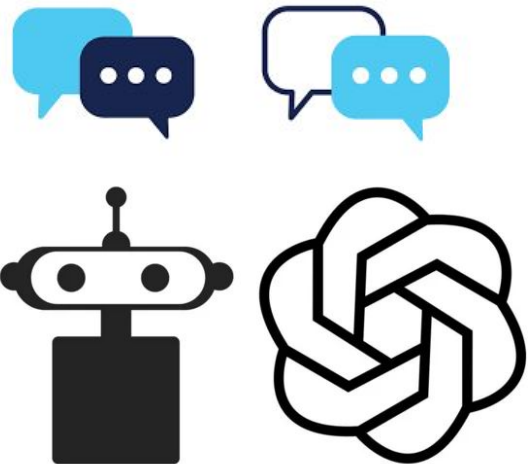
**ChatGPT:**

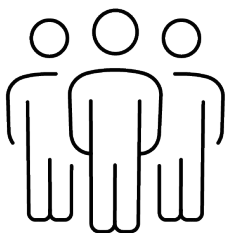
**As of recent reports, PostgreSQL deployments are increasingly moving to the cloud**, but a significant portion still remains on-premises.

Currently, public cloud deployments for PostgreSQL, including DBaaS, KaaS and IaaS, are used by 34.8% of organizations, showing an increase of 15% over the past year.

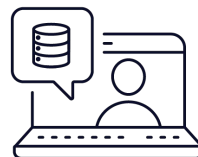
This demonstrates a broader trend towards cloud adoption, especially for managed database services that offer benefits like automatic scaling, high availability, and reduced maintenance burden compared to on-premises deployments.

**In summary, the distribution of PostgreSQL deployments is now roughly 35.6% on-premises, 24.4% hybrid cloud, and 34.8% fully cloud-based.**





EDB **builds Postgres**,  
for and with the  
community.



EDB **supports Postgres**,  
giving customers access  
to the people making it.



EDB **enhances Postgres**,  
with enterprise-grade  
features/tools/services.

**We are the specialists.**



# EDB Portfolio

## 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)

### PLATFORM TOOLS AND SERVICES

MIGRATION  
PORTAL

CONTINUOUS HIGH  
AVAILABILITY

BACKUP AND  
RECOVERY

### EXTENSIBILITY

CSP INTEGRATIONS

DEVOPS TOOLING

KUBERNETES TOOLING

GENAI & LLM INTEGRATIONS

LAKEHOUSE INTEGRATIONS



**Support 24\*7 for Postgres and SLO on Remedy**



**Professional services across the Postgres lifecycle**



**Technical account management**



**Remote DBA**

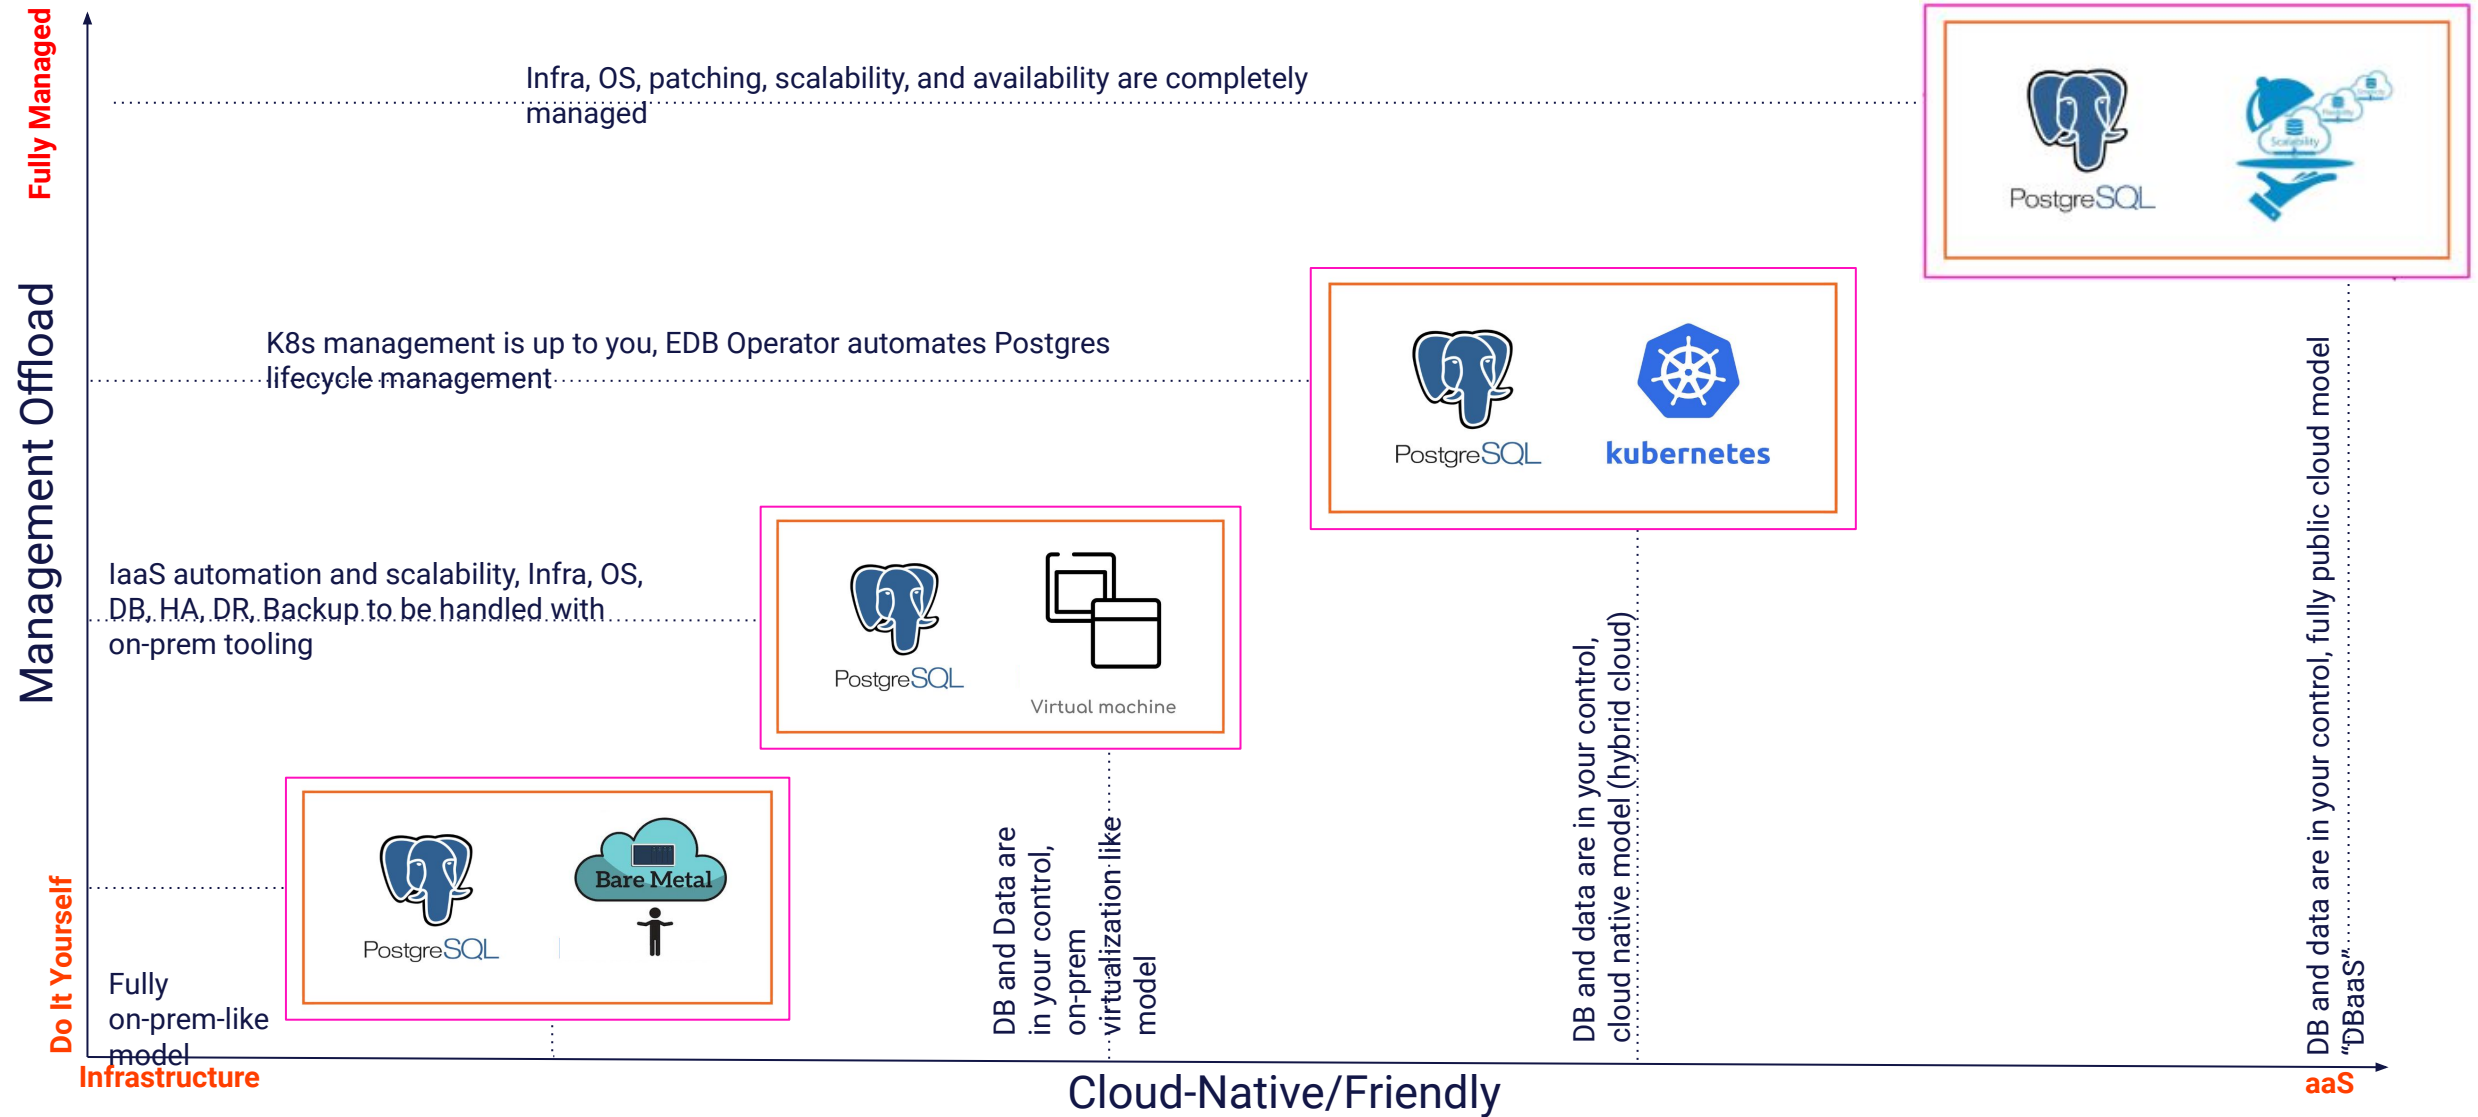


**Training:** instructor-led, on-demand, and certification





# Meet Customers Where They Are



# From Big Animal to Big Data and AI: a year in review

## Q1 2024

Cluster Pause-Resume  
WAL on dedicated volumes  
DataDog & New Relic  
TDE and BYOK for AWS

## Q3 2024

Volume Snapshot backup  
TDE for Azure and GCP  
PgBouncer and PGD Proxy  
co-location with PG nodes  
GCP Marketplace

## Q4 2023

AWS Marketplace  
PostGis support  
Built-in Monitoring

## Q2 2024

Big Animal becomes Database CS as part  
of **EDB Postgres AI Platform Launch**

-Analytics with Postgres Lakehouse  
clusters

-AI/ML tech preview: EDB Postgres AI  
database (aiddb) for vectors

-Hybrid data estate management: visibility  
into self-managed Postgres and RDS

Q4 2024 and onward

...



# Closes the gap

between what enterprises need and what DBaaSes offer today

## What is it?

It is EDB's PostgreSQL delivered **on any CSP with any deployment model.**



All EDB offerings are available via classic order form or via cloud providers' marketplaces.

Get the best from your cloud spending, via **EDB Commit Draw Down.**

Get new subscriptions, renewal, expansion, and services.



**For each cloud provider, EDB has you covered with a seamless experience.**

Same Postgres, same pricing, same management, same licensing, same support, same tooling.

## How to get it?



EDB Order Form

- EDB Invoice for SW
- CSP Invoice for Infrastructure



Cloud Marketplaces  
'Private Offer'

- Pay via Cloud Committed Spend
- Unified Invoice from CSP

New Subscriptions,  
Renewals,  
Expansions

License Portability  
Professional Services

All EDB Offerings, Cloud  
and On-Prem

# EDB cloud values



OLTP, OLAP, AI via Postgres	No Platform Lock-In	More Control and Optimization	Insurance for Postgres Workloads	Cheap and Fast Oracle Exit	Granular Database and Storage Encryption	Cross-Estate Postgres Observability	Max Uptime and Geo- Distribution
✓	✓	✓	✓	✓	✓	✓	✓

NEW

Unified  
Workload  
Management

Same Postgres  
Everywhere

Tunable 97%  
SuperUser

Support with Postgres  
Focus and **SLO on  
Remedy**

Oracle  
Compatibility

Transparent Data  
Encryption  
(TDE)

NEW

Postgres  
Observability

Multi-master  
99.995% SLA



✗	✗	✗	✗	✗	✗	✗	✗
---	---	---	---	---	---	---	---

Multiple  
Services with  
Their Interfaces

CSP Only  
and No  
Exit

Limited  
Optimization

Support on  
the API, not  
on Postgres

Refactoring  
and  
Migration

Encryption on  
Storage Volume  
Only

Own Service  
Limited  
Observability

Only Master  
/ Slave  
Model



Aurora  
and RDS



DB for Postgres  
and Citus



Cloud SQL  
and AlloyDB



# Closes the gap

between what enterprises need and what DBaaSes offer today

## Where to deploy?

In Your Cloud Account



In EDB's Cloud Account

## Deploy in your cloud tenant or use EDB's.

Same flexible choice of provider, regions, and availability zones.

Different consideration for infrastructure pricing and internal governance.

## The EDB DBaaS provides maximum flexibility in DBaaS deployment options:

Pick your preferred tenancy option by cluster / workload.

Pick your preferred combination of cloud provider, region, and infrastructure sizing.



### Where to deploy

Please reference [this page](#) for more information on account types.



### Provider & Region

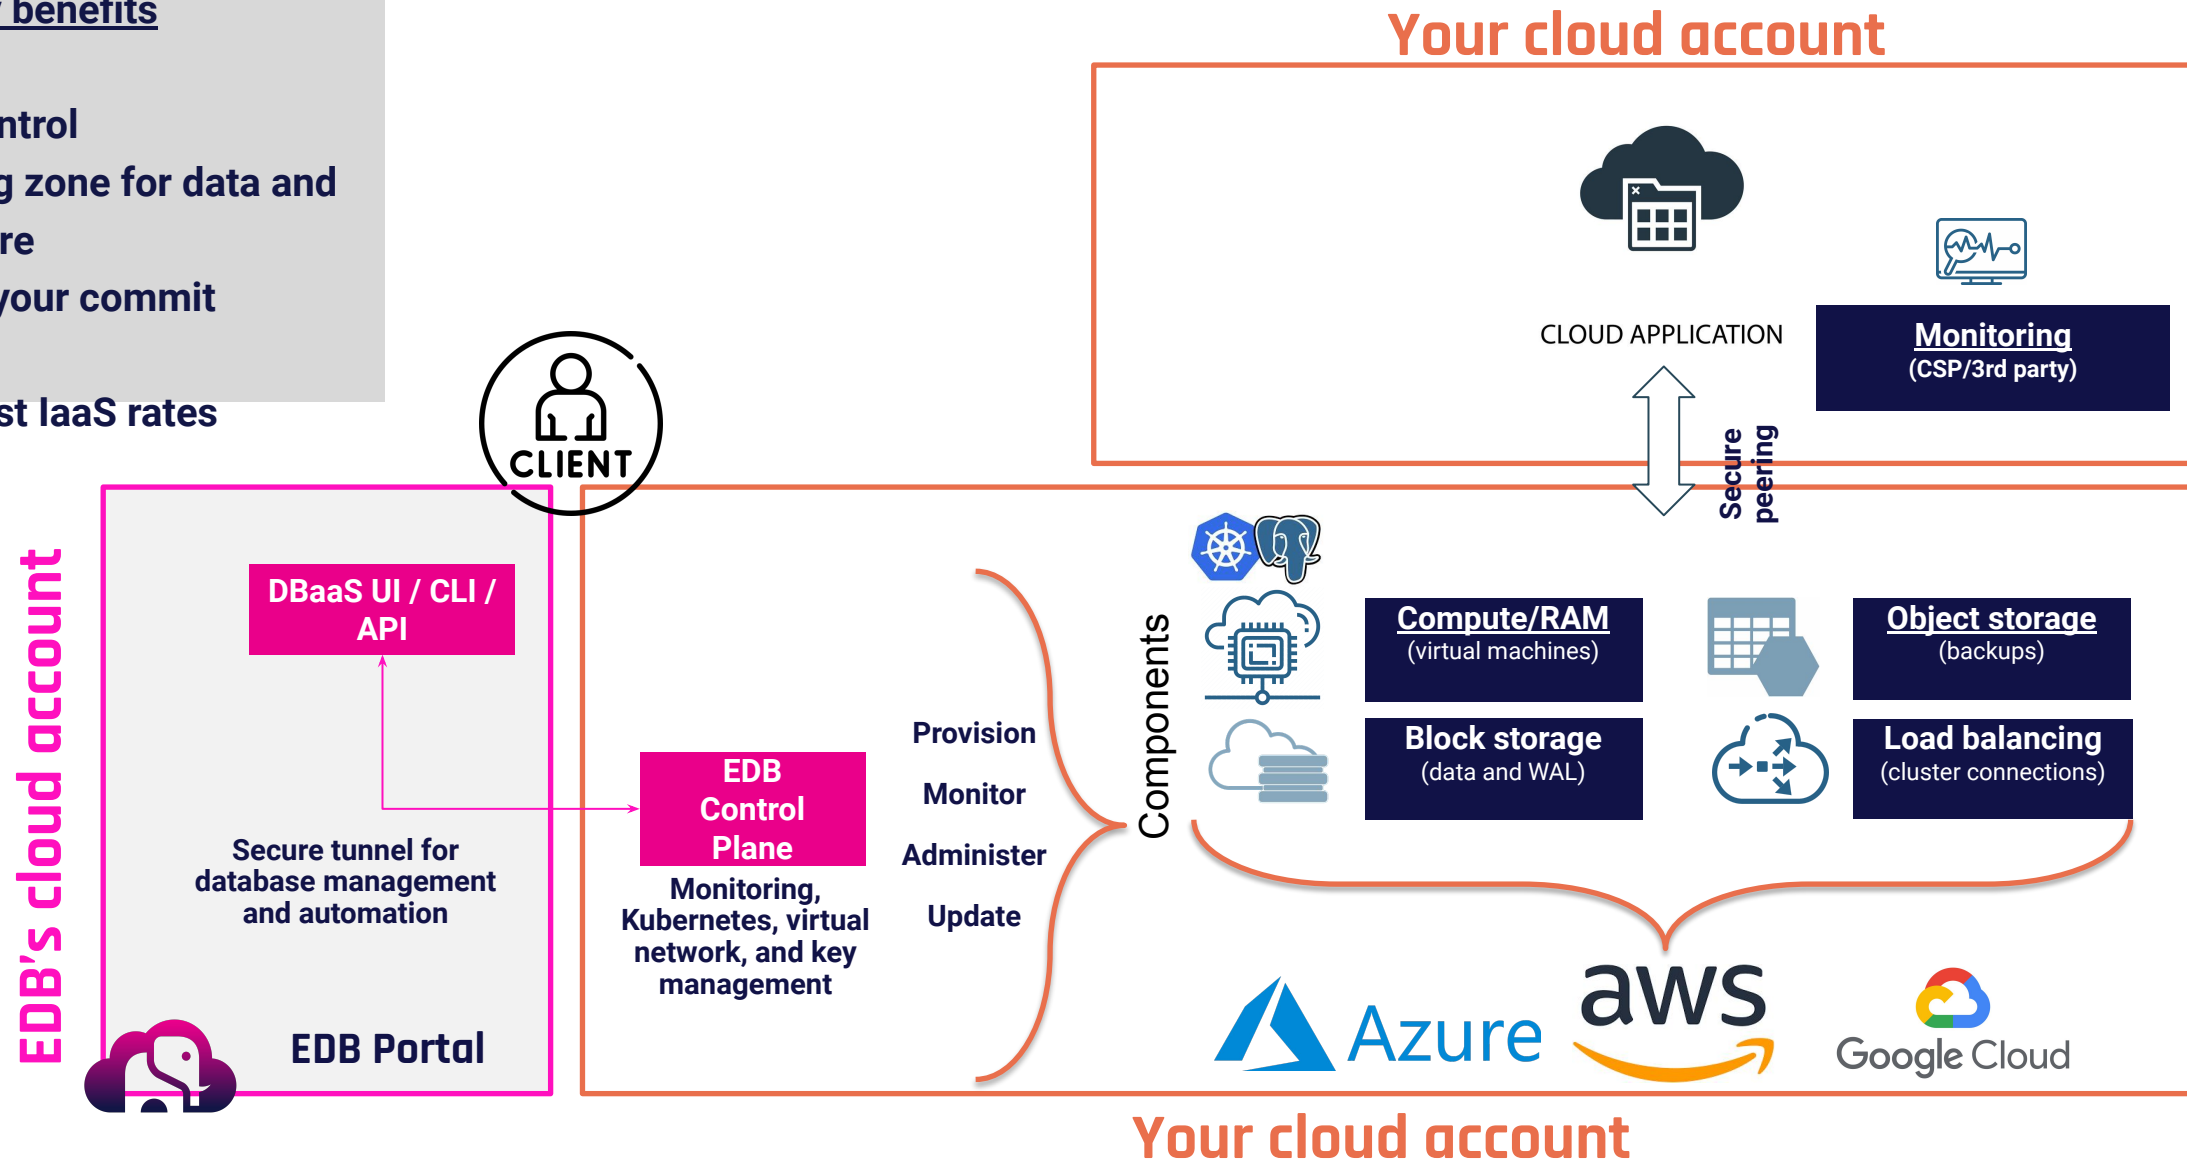
Please reference [this page](#) for more information on Regions.



# DBaaS running in *your* cloud account

## Key benefits

- Granular control
- Your landing zone for data and infrastructure
- Burn down your commit contract
- Get your best IaaS rates

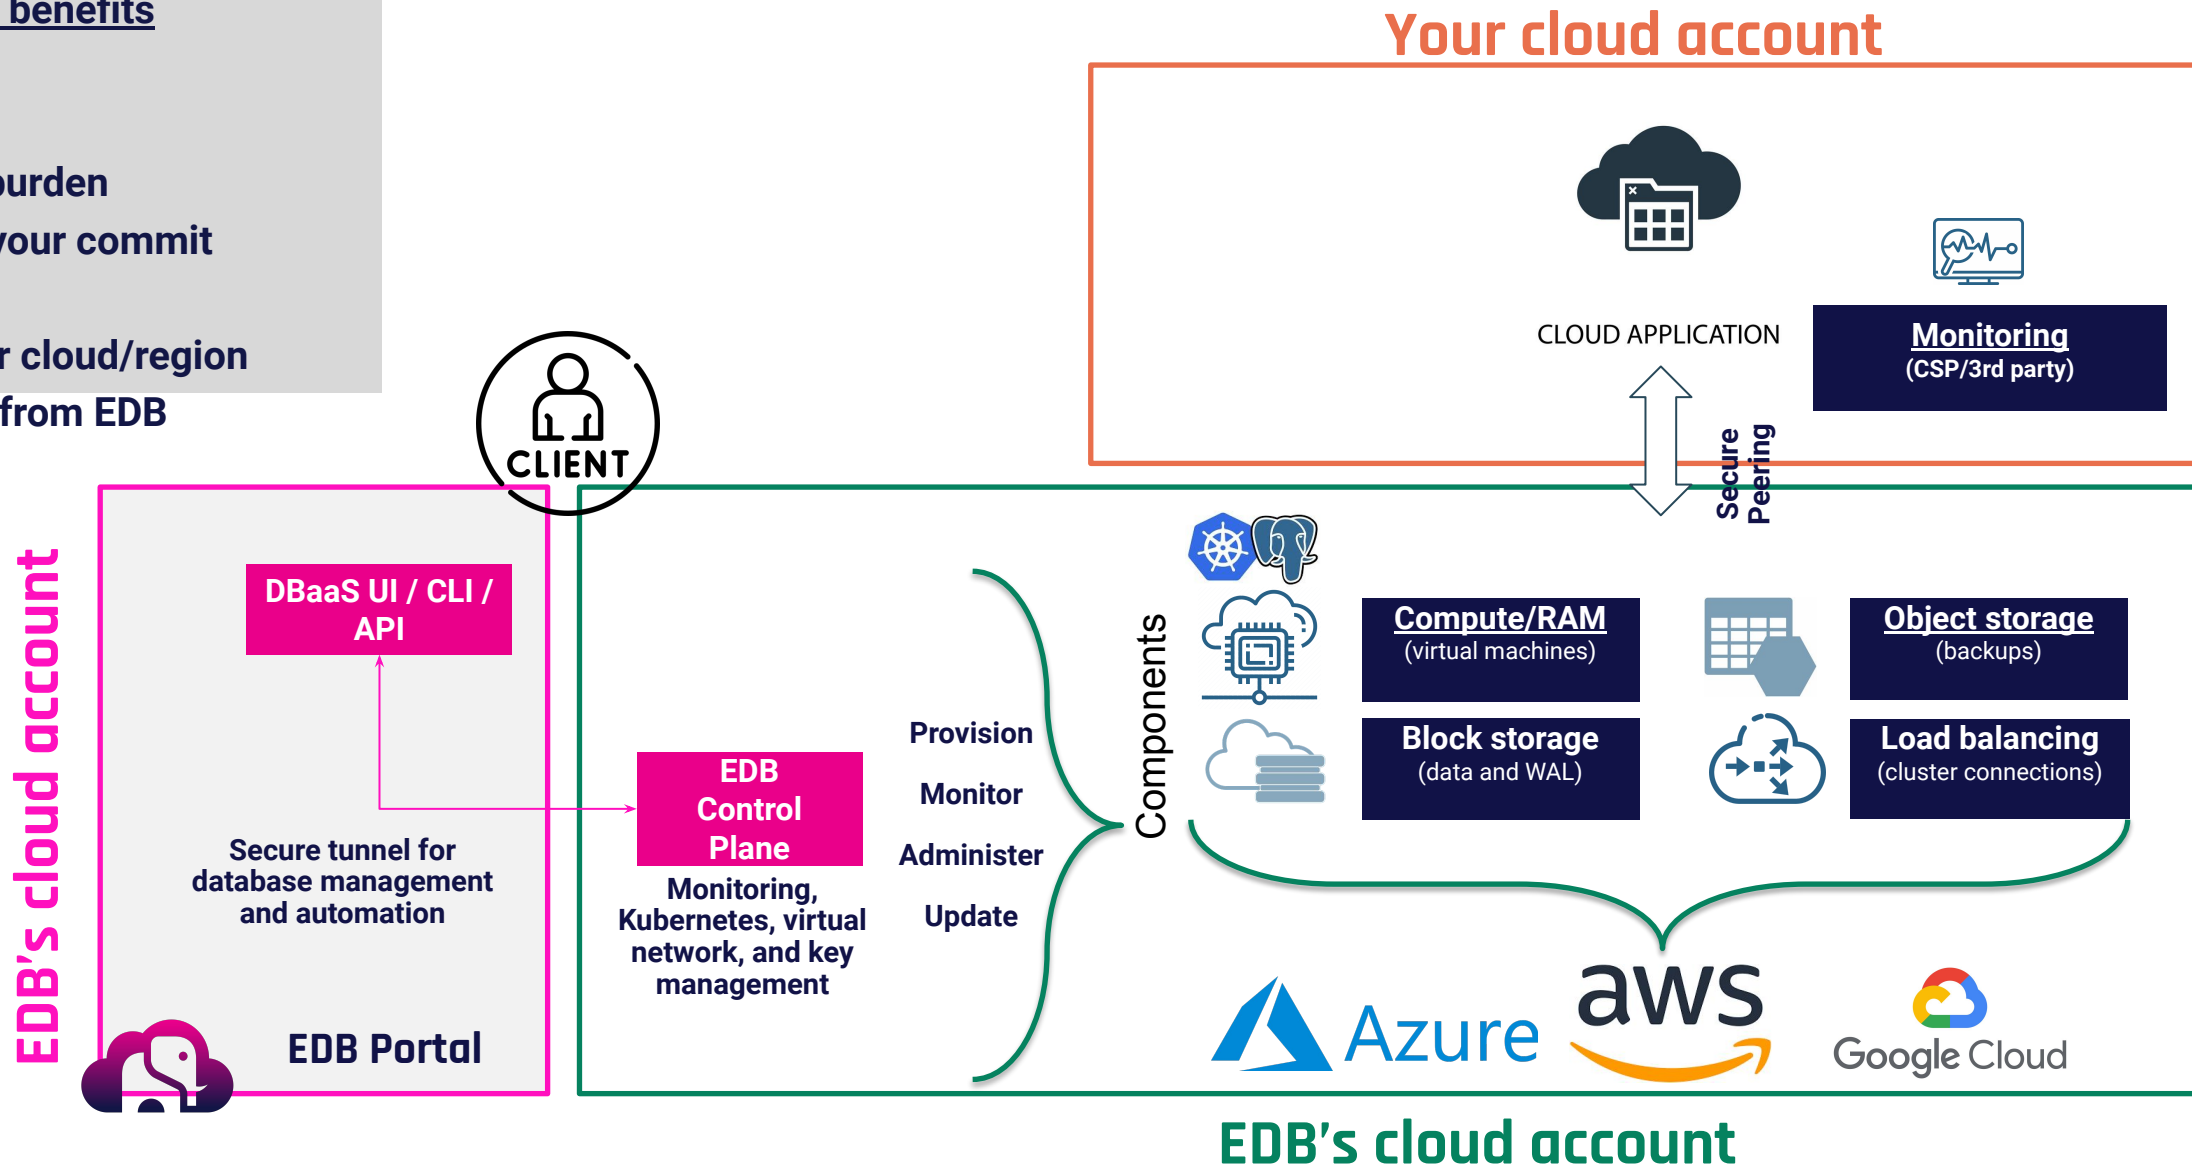




# DBaaS running in *EDB's* cloud account

## Key benefits

- Fully aaS
- Minimized burden
- Burn down your commit contract
- Choose your cloud/region
- Get one bill from EDB



# Packaging

- Different plans, with related features
- vCpu/hr metric, per plan
- Committed contracts based on volume
- Support included



**PostgreSQL**

Open source PostgreSQL

Fully managed service  
24/7 technical support  
Run on Azure, AWS, GCP  
Open source tools  
Automatic backups  
HA and DR options



**EDB Postgres  
Extended Server**

PostgreSQL

- + Transparent data encryption
- + EDB tools
- + EDB Postgres extensions



**EDB Postgres Advanced  
Server**

EDB Postgres Advanced Server

- + Oracle compatibility
- + Admin, security, tuning, & audit enhancements
- + EDB tools
- + EDB Postgres extensions
- + Transparent data encryption



**EDB Postgres Distributed  
(Extended or Advanced  
Server)**

PostgreSQL or Advanced

- + Distributed high availability (99.995%)
- + Transparent data encryption
- + Oracle compatibility (Advanced)
- + Admin, security, tuning, & audit enhancements



# Service Level Driven

We supports different cluster types and replications.

- Which Service Level does your workload require?
- Which RPO/RTO?

You have flexibility to mix the following:

- Type of cluster and replication
- Single / Multi Az-Region

to match your requirements.

## Choosing your cluster type and configuration

Consideration	Single node	Primary/standby HA	HA + standby replica	Distributed HA single region	Distributed HA multi-region
Data replication	None	Physical	Physical	Logical	Logical
Region	Single	Single	Multi	Single	Multi
VM failure tolerance	✓	✓	✓	✓	✓
AZ failure tolerance	TBD	✓	✓	✓	✓
Region failure tolerance	TBD	TBD	✓	TBD	✓
Recovery time objective	varies	35s-60s	varies	0	0
Recovery point objective	<5 min	0	<5 min	0	30s (configurable)
Service level agreement	99.5%	99.99%	99.99%	99.99%	99.995%



# Types of clusters and replicas

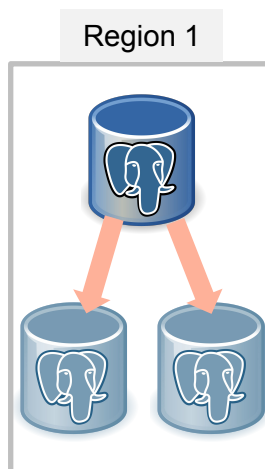
## Single node

- Single zone
- Generally for non-production environments that can tolerate zonal outages



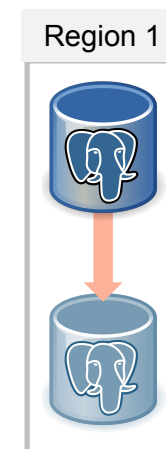
## Master + 2 replicas

- 2 secondary replicas within a region across 3 zones
- Replication is sync to one standby and async to the other
- Can modify sync behavior on a per-transaction, per-session, per-user, or per-database level



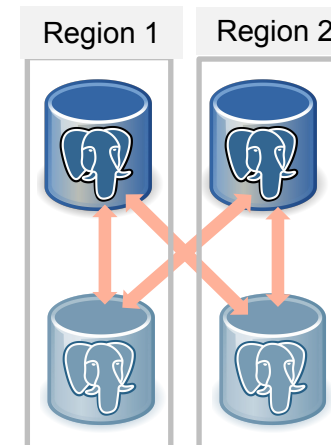
## Master + 1 replica

- 1 secondary replica within a region across two zones
- Automatically disables sync replication during maintenance operations to ensure write availability
- Can change from sync to async on per-session/per-transaction basis



## Extreme high availability

- “Active/Active” mesh replication topologies cross region/cross zonal
- No-downtime major version upgrades
- Self-service superuser access on clusters
- Fast recovery and higher performance workloads

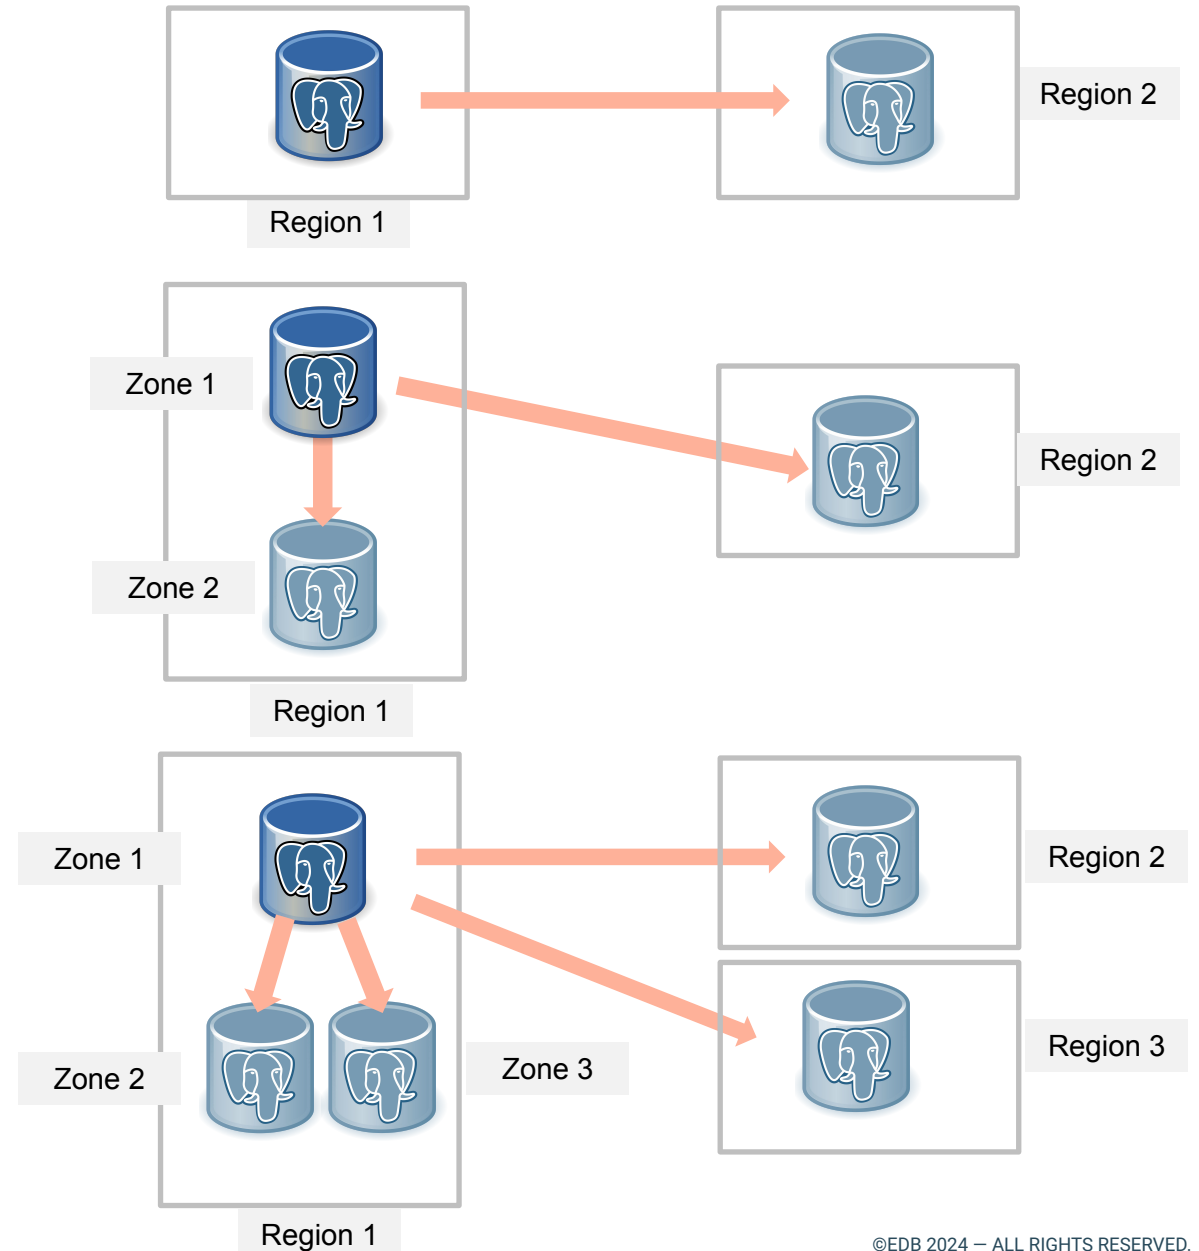


## + 'Faraway Replicas'

Allows across region async replicas

- Read-only workloads on replicas in another region
- Allows independent database backups in another region
- Can have a mix of faraway replicas in a mix of regions
- Can be different sized compared to active nodes
  - Different sized compute/RAM
  - Different storage subsystem

### 1 region to many replicas



# Who does what with EDB DBaaS?

---

## **EDB** ensures smooth running operations:

- 24x7 monitoring
- Infrastructure availability and performance
- High availability and backups
- Security and compliance
- OS, database patching – including major version upgrades
- Network configuration and security

## **Customer** optimizes the database for their enterprise needs:

- Selecting appropriate compute resources (VM and disk) and replicas to fit workload
- Data modeling
- Query performance
- Securely managing database credentials
- Postgres Parameter tuning



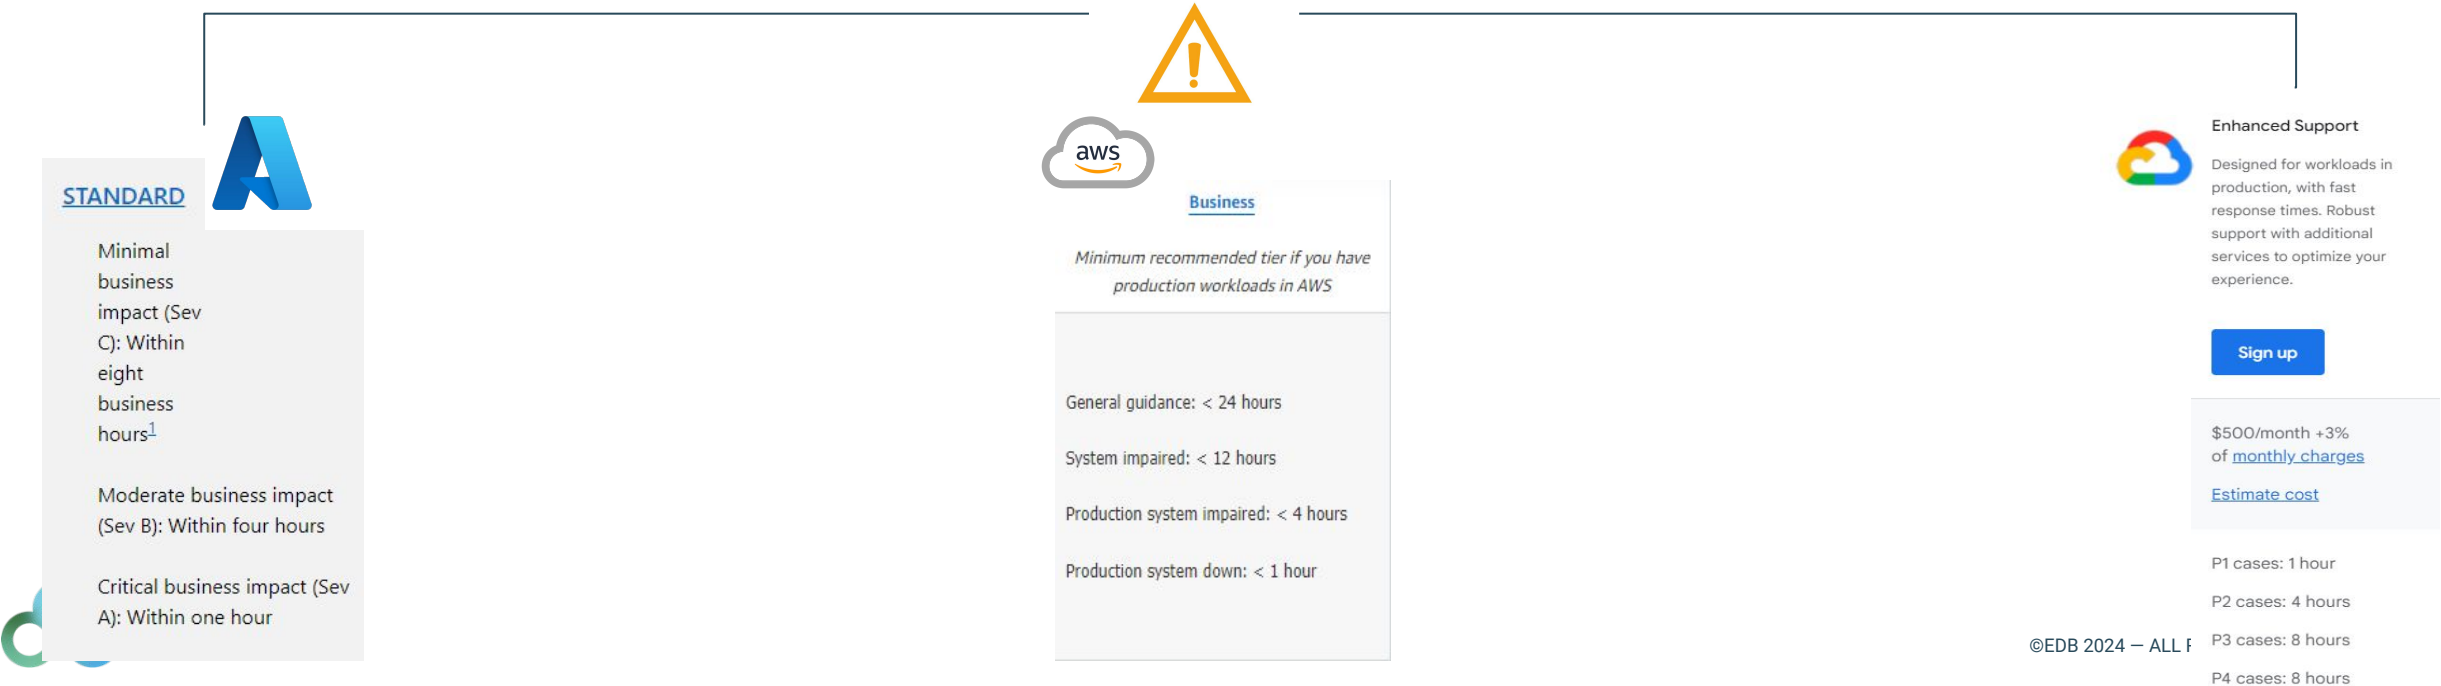


# Support

- Support on the Postgres engine
- Included in the DBaaS price
- SLO for response and remedy

Availability:		24 hours x 7 days
Initial Response	Severity-1	30 minutes
	Severity-2	30 minutes
	Severity-3	60 minutes
Service Level Objective	Severity-4	1 business day
Remedy	Severity-1	24 hours
	Severity-2	48 hours
	Severity-3	15 business days
Service Level Objective	Severity-4	45 business days

## Cloud providers' only commitment is on response time



# How to interact and integrate with our DBaaS?

## WebUI

### Cluster Type

Please reference [this page](#) for more information on cluster types.

#### Single Node

Ideal for non-production workloads; creates a single primary with no standby replicas. [Learn More](#)

#### Primary/Standby High Availability

Creates a cluster with one primary and up to two standby replicas in different availability zones. [Learn More](#)

#### Distributed High Availability

Creates a cluster, powered by EDB Postgres Distributed (PGD), with 3 nodes and up to 2 Data Groups spread across three availability zones. [Learn More](#)

### Number of Standby Replicas

Adding standby replicas will increase your number of CPUs, as well as your cost.

2 replicas

### Where to deploy

Please reference [this page](#) for more information on account types.

### Provider & Region

Please reference [this page](#) for more information on Regions.

### Cluster Summary

Cluster Type	Primary/Standby High Availability
Cluster Nodes	3 nodes
Deployment	BigAnimal
Provider	AWS (BigAnimal)
Region	US East 1
Cluster Name	Cluster0
Postgres Type	PostgreSQL
Postgres Version	15
Instance Size	m5.large, 2vCPU, 8GB RAM
Volume Type	General Purpose SSD (gp3)
Volume Properties	4 Gi, 3000 IOPS, 125 MB/s Disk Throughput
Networking	Public
Backups Retention	30 Days
Maintenance Window	Not Enabled
Connections	Read/write
PyBouncer	Not Enabled

## APIs

Search...

Billing >

Cloud Providers >

Cluster Architectures >

Clusters >

getClustersOverview

getClusters

createCluster

deleteCluster

getCluster

patchCluster

updateCluster

getClusterConnection

getClusterFarawayReplicas

getClusterGroup

getClusterGroupConnection

### Clusters

getClustersOverview

Returns clusters overview

AUTHORIZATIONS > bearerAuth

PATH PARAMETERS

projectId string  
required Example: prj\_abcd1234

QUERY PARAMETERS

q string  
Example: q=John  
Fulltext search

sort string  
Examples:

- sort=firstName - sort by "firstName" property ASC
- sort=state.city - sort by "city" property of the "state" property ASC
- sort=-firstName - sort by "firstName" property DESC
- sort=-state.city - sort by "city" property of the "state" property DESC

GET /projects/{projectId}/clusters

Response samples

200	400	401	403
404	409	412	429
500	501	503	504

Content type  
application/json

Copy Expand all Collapse all

```
{  - "data": [    + {      }  ]}
```

## CLI

Use the Command Line Interface (CLI) for BigAnimal management activities such as cluster provisioning and getting cluster status from your terminal.

The CLI is an efficient way to integrate with BigAnimal and enables system administrators and developers to script and automate the BigAnimal administrative operations.

[Documentation: Using the BigAnimal CLI](#)

```
EDB $ biganimal create-cluster
Cluster Name: my-biganimal-cluster
Password: *****
Postgresql Type: EDB Postgres Advanced Server
Use the server type for installation: 1
Postgresql Version?
  13
  12
  11
```

### Install latest version

```
curl -LO "https://cli.biganimal.com/download/${uname -s}/${uname -m}/latest/biganimal"
```

### Download latest binary

Linux (x86_64)	SHA256: 7756a736cd7072ebcd8302b278cd3b994e85f3b8993482084945723d	Download
Linux (arm64)	SHA256: 6ae7f7f0a877723380085736e49965276569642410db47a020ab2c0bdf	Download
Darwin (x86_64)	SHA256: 7b547b798b40b3bdebb55c49726298298012bb108e599b7f3eb5205d3	Download
Darwin (arm64)	SHA256: e69727b4d79b9858d8c855472e8fc7bd7225297b2cb54f2f6e4d409a0	Download
Windows (x86_64)	SHA256: 0d3512664576cbad79d4d7f8f48f53e0272530c9b3d36a7c33f5c505bc	Download

## Terraform

biganimal

Overview

BIGANIMAL DOCUMENTATION

Filter

[biganimal provider](#)

> Resources

> Data Sources

### BigAnimal Provider

A Terraform Provider to manage your workloads on EDB BigAnimal by interacting with the BigAnimal API. The provider is licensed under the [MPL v2](#).

Use the navigation to the left to read about the available provider resources and data sources.

### Requirements

- A BigAnimal account with an organization set up. If you don't already have a BigAnimal account, see [Getting started with the BigAnimal free trial](#).
- Terraform version 0.13.x or later.
- A BigAnimal API token (<https://www.enterprisedb.com/docs/biganimal/latest/reference/api/#using-the-get-token-script>)

©EDB 2024 — ALL RIGHTS RESERVED.



Thank You

Mario Vetrano,  
Cloud Specialist EMEA  
[mario.vetrano@enterprisedb.com](mailto:mario.vetrano@enterprisedb.com)







**EDB**

Postgres for the AI Generation

Demystify complexity through reference  
architectures and automation

Davide Tammaro - Senior Sales Engineer

# Architecture requirements

It's not a matter of what Postgres can do for you. It's a matter of what you want to achieve.

**Target availability:** Business accepts service unavailability in case of **Minor Issues**

SLA level: 99.99% allowed downtime

- Daily: 8.6s
- Weekly: 1m 0.48s
- Monthly: 4m 21s
- Quarterly: 13m 2.4s
- Yearly: 52m 9.8s

SLA level: 99.999% allowed downtime

- Daily: 0.86s
- Weekly: 6s
- Monthly: 26s
- Quarterly: 1m 18s
- Yearly: 5m 13s

Questions to answer:

- How long does it take a resource switchover that implies application reconnection?
- How long does it take a major release upgrade that implies a service maintenance window?



# Architecture requirements

It's not a matter of what Postgres can do for you. It's a matter of what you want to achieve.

**Target availability:** Business accepts service unavailability in case of **Minor Issues**

Architecture Principle:

**synchronous data copies and self-healing automations:**

- DDL + DML synchronous replica
- Active/Passive or Active/Active redundancy architecture

Questions to answer:

- How long does it take a resource switchover that implies application reconnection?
- How long does it take a major release upgrade that implies a service maintenance window?





# Architecture requirements

It's not a matter of what Postgres can do for you. It's a matter of what you want to achieve.

**Disaster recovery metrics:** Business accepts service unavailability in case of **Major Issues**

**Recovery Time Objective:**

To set the maximum business process outage  
(future transactions)

**Recovery Point Objective:**

To set the maximum **unrecoverable** business lost  
acceptable (transactions committed in the past)

Questions to answer:

- RTO: How long does it take to bring the service up again?
- RPO: How old is my last recoverable transaction, or, what is the maximum acceptable interval during which transactional data is lost?



# Architecture requirements

It's not a matter of what Postgres can do for you. It's a matter of what you want to achieve.

**Disaster recovery metrics:** Business accepts service unavailability in case of **Major Issues**

## Architecture Principle:

**Multi Region Redundancy and Availability Objective:** How many locations are storing the data and providing the service according to the RTO and RPO metrics

Questions to answer:

- RTO: How long does it take to bring the service up again?
- RPO: How old is my last recoverable transaction, or, what is the maximum acceptable interval during which transactional data is lost?



# Reference Architectures

Templates for concrete Postgres architectures

Architecture	Data Redundancy	Automatic Failover	Multi node commit	Multi Region Data redundancy	Multi Region availability	Recommended Use
Single Node	No	No	No	no	no	Development/testing
Primary/Replica MultiNode	Yes	Yes	Yes	Yes	Primary/Replica	Production, with no transaction loss acceptable. Target Availability 99.99%
Always-on Single Location	Yes	Yes	Yes	No	No	Production, with no transaction loss acceptable. <b>Target availability 99.999%</b>
Always-on Multi Location	yes	Yes	multiple commit scenarios	YES	Active/Active - distributed	Extreme high availability and/or Geo-distributed applications. <b>Target availability 99.999% and RTO=0</b>

**Multi-Region Data redundancy:** data and service is redundant on more than one geography

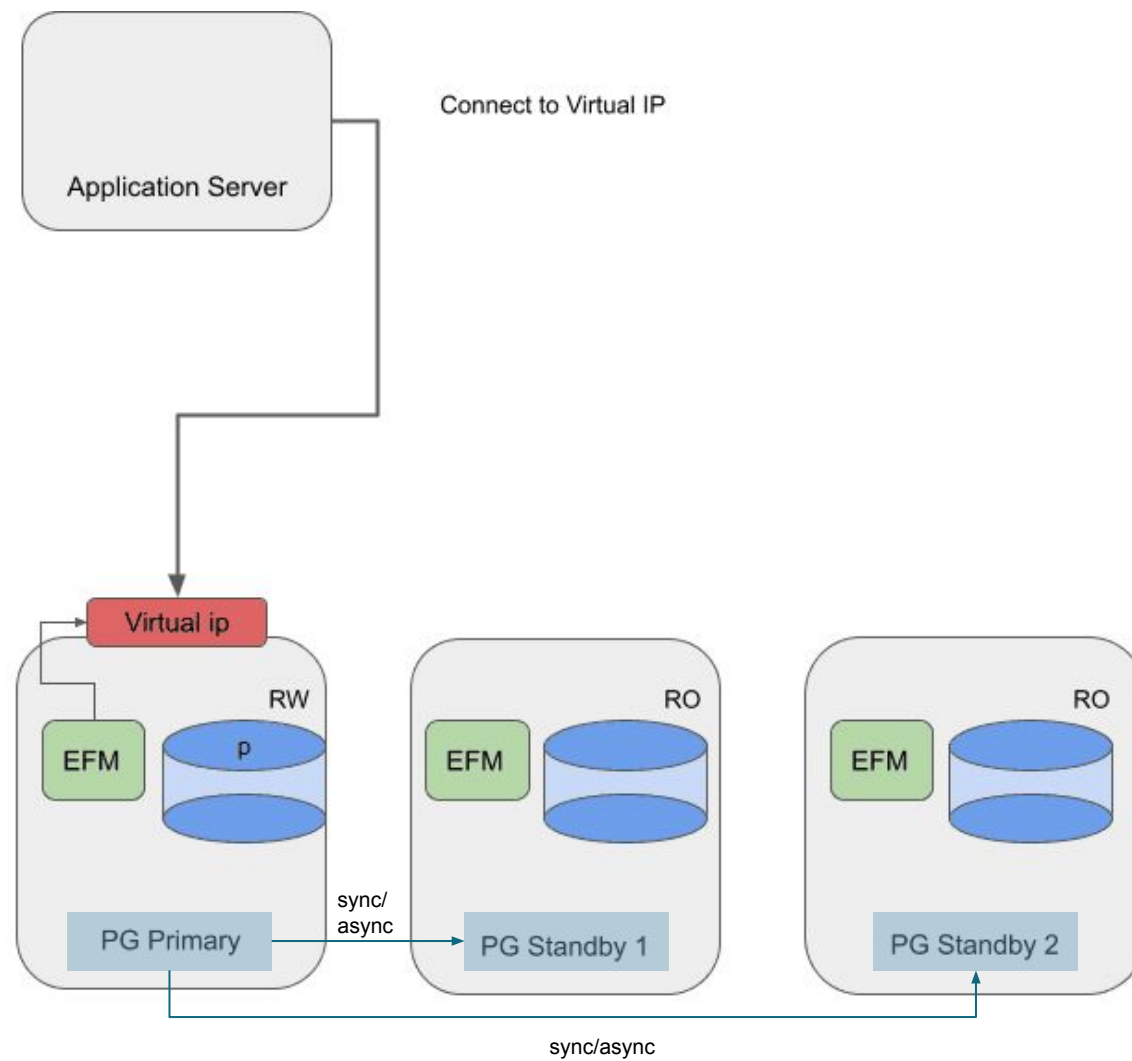
**Multi-Region availability:** data and service is also available for users in more than one geography



# Multi-node

Production, with no transaction loss acceptable. Target Availability: 99.99%

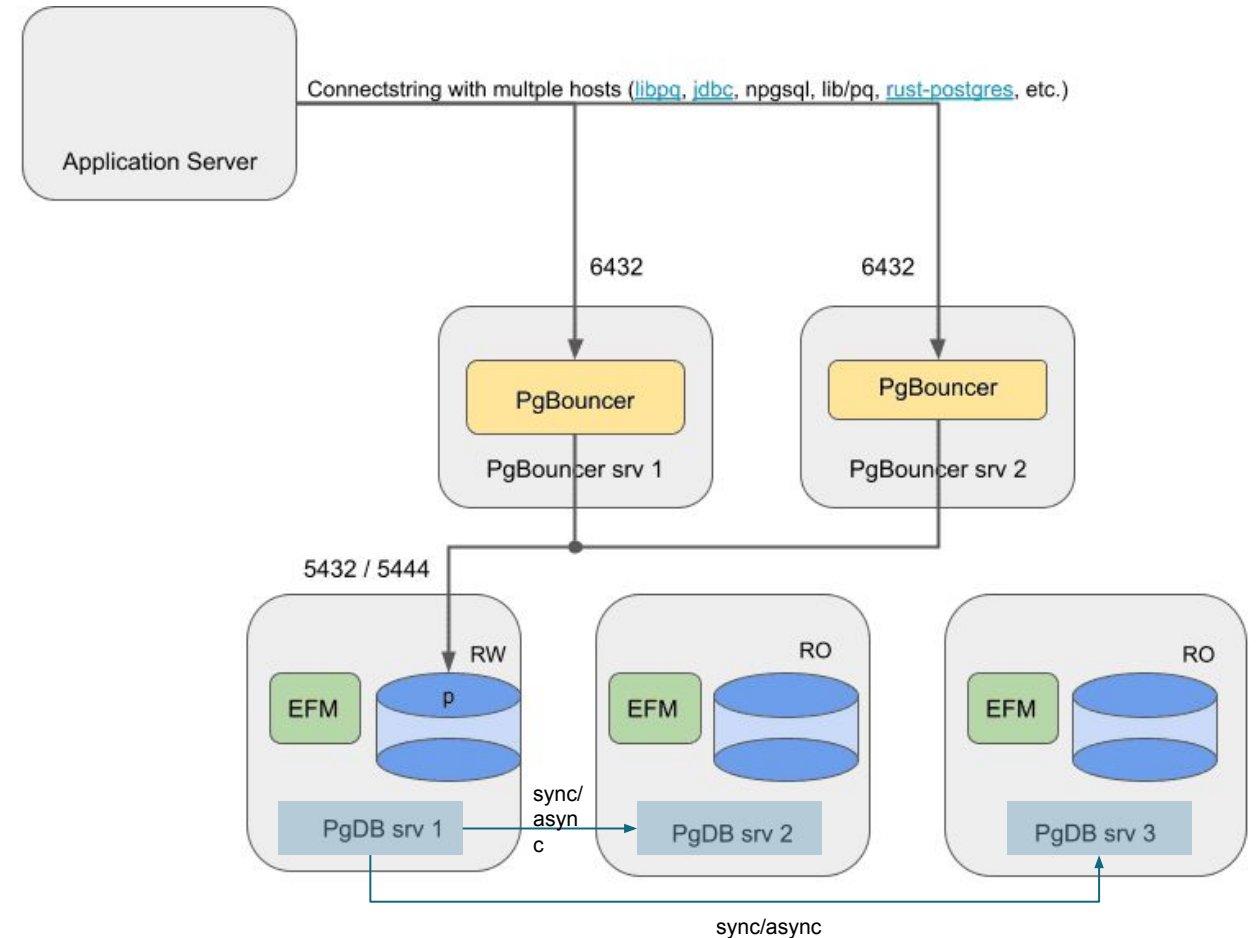
- EFM Failover Manager manages VIP addresses out of the box.
- VIP addresses allow applications to connect to a single IP address that is being routed to the primary database server.
- This architecture is the most basic solution to run when VIP addresses are available in your environment.



# Multi-node with connection pooler

Production, with no transaction loss acceptable. Target Availability: 99.99%

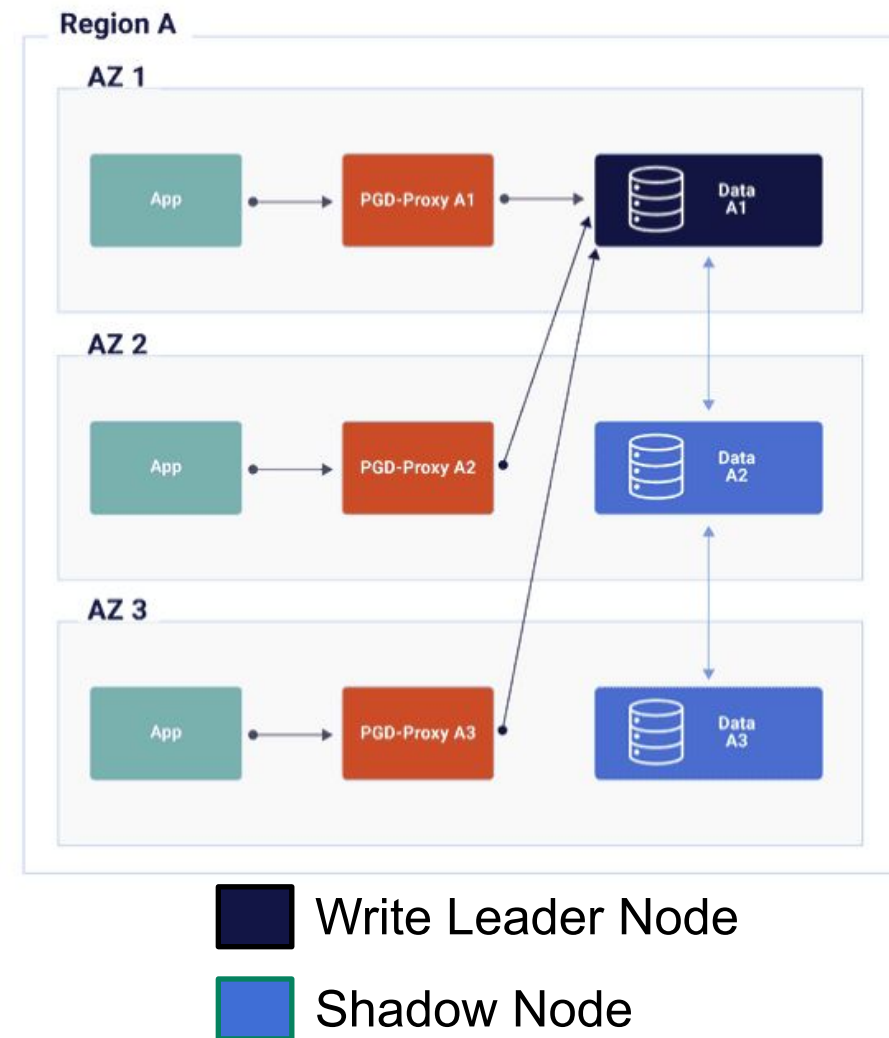
- PgBouncer adds capabilities such as connection pooling and the option to halt traffic.
- You can also use it as a proxy between the client and the Postgres database server.
- Route the traffic to the correct primary database server.
- Pgbouncer can help to optimize session management.



# Always-on single location

Production, with no transaction loss acceptable. **Target availability: 99.999%**

- Applications connect to any available PGD-Proxy node
  - PGD-Proxy routes application connection to the current **“write leader”** (vs. “shadow nodes”)
  - Eliminates conflict scenarios for applications not designed/updated for multi-master environments
- Redundant hardware for resilience during local failures
  - Multiple availability zones, 2 (or 3) PGD “data” nodes
  - “Witness” provides 3<sup>rd</sup> node for RAFT consensus
  - One PGD-Proxy node for each data node
- Other components (not shown):
  - Barman in alternate location used for DR
  - Postgres Enterprise Manager (PEM) for monitoring

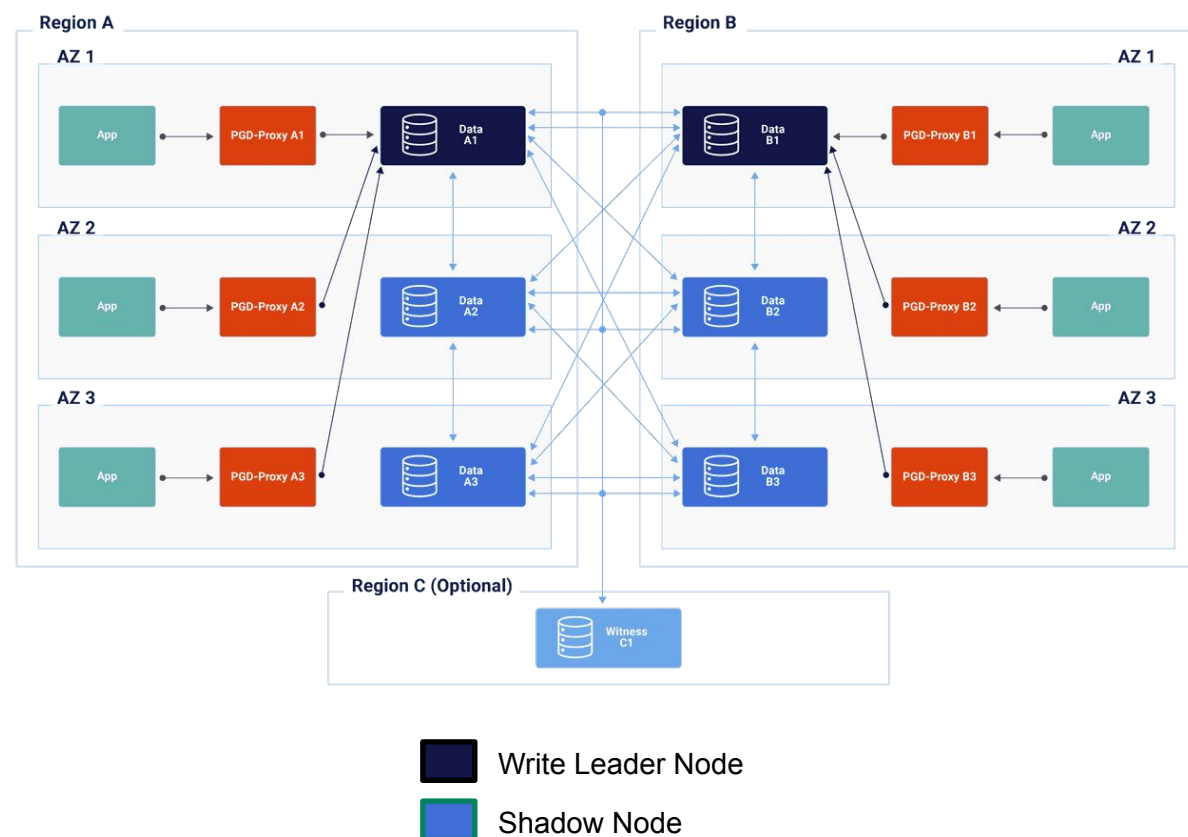




# Always-on multiple location

Extreme high availability and/or geo-distributed applications. **Target availability: 99.999% and RTO=0**

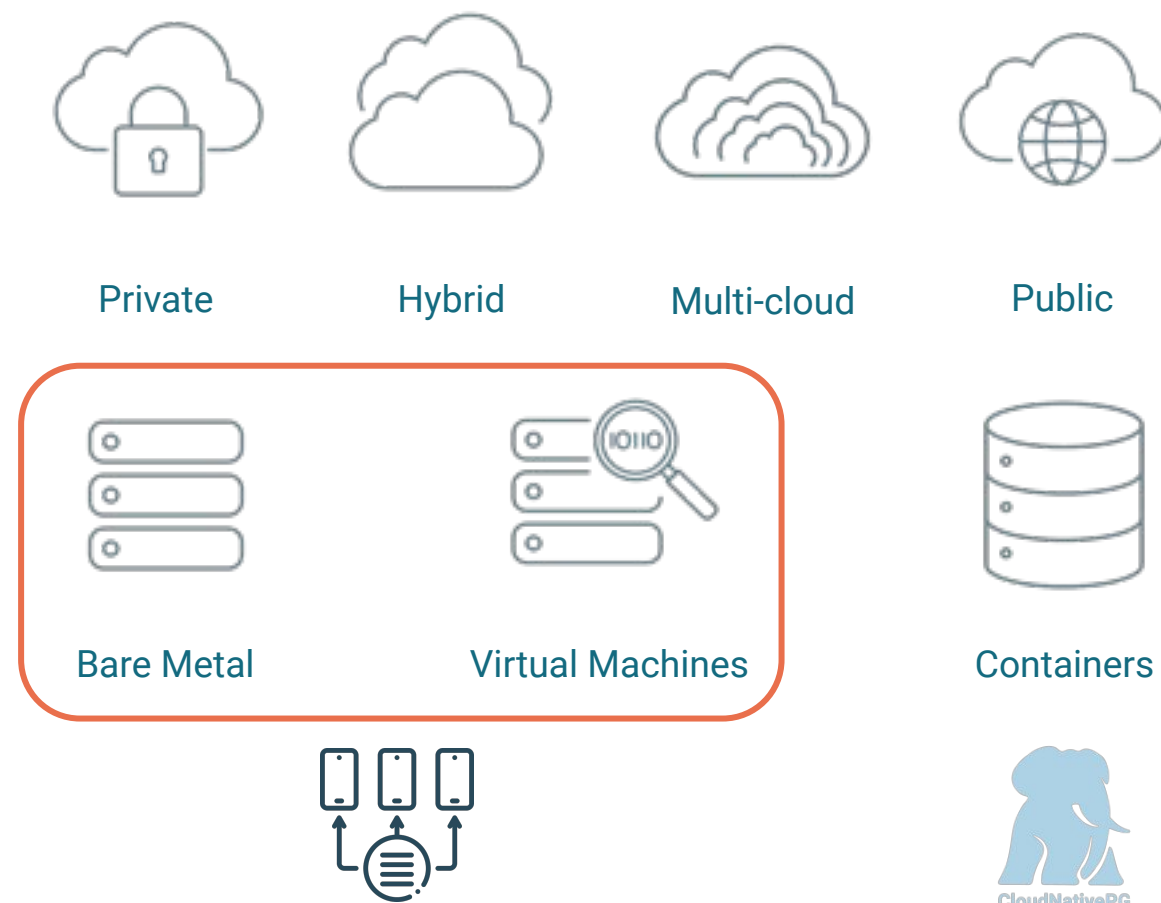
- Applications connect to any available PGD-Proxy node in their local (or remote) location
  - PGD-Proxy routes application connection to the current “**write leader**” in proxy’s location
  - Supports write leaders in each location or a single write leader across multiple locations
  - Supports Active/Active, Active/Passive, and Active/DR
- Redundant locations for resilience during local and regional failures
  - 4 or 6 PGD data nodes total (2 or 3 in each location)
  - “Witness” node in 3rd location for RAFT consensus
  - One PGD-Proxy node for each data node
- Other components (not shown):
  - Barman used for DR
  - Postgres Enterprise Manager (PEM) for monitoring



# Deployment and lifecycle automation

From self-managed to fully managed DBaaS in the cloud

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



# Trusted Postgres Architect (TPA)

- TPA is an orchestration tool that uses Ansible to deploy Postgres clusters according to EDB's recommendations.
- TPA embodies the **best practices** followed by EDB, informed by many years of hard-earned experience with deploying and supporting Postgres.
- These recommendations are as applicable to quick testbed setups as to **production environments**.



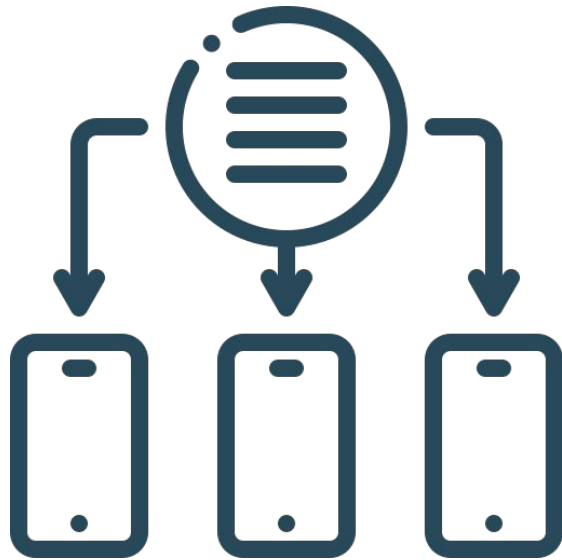
# What can TPA do?

- TPA is built around a declarative configuration mechanism that you can use to describe a Postgres cluster, from its topology right down to the smallest details of its configuration.
- TPA can:
  - **Provision servers** (e.g.: AWS EC2 or Docker); or you can deploy to existing servers)
  - **Configure** the operating system
  - **Install and configure Postgres and associated components** (PGD, barman, pgbouncer, repmgr, and various Postgres extensions)
  - **Run automated tests** on the cluster after deployments
  - **Deploy future changes to your configuration** (e.g., changing Postgres settings, installing and upgrading packages, adding new servers, and so on)



# How do I use it?

- Configure
- Provision
- Deploy



## Trusted Postgres Architect (TPA)

Open Source from EDB

```
>tpaexec configure mycluster \  
--architecture M1 \  
--postgresql 15 \  
--enable-patroni_
```



```
>tpaexec provision \  
mycluster_
```



```
>tpaexec deploy \  
mycluster_
```

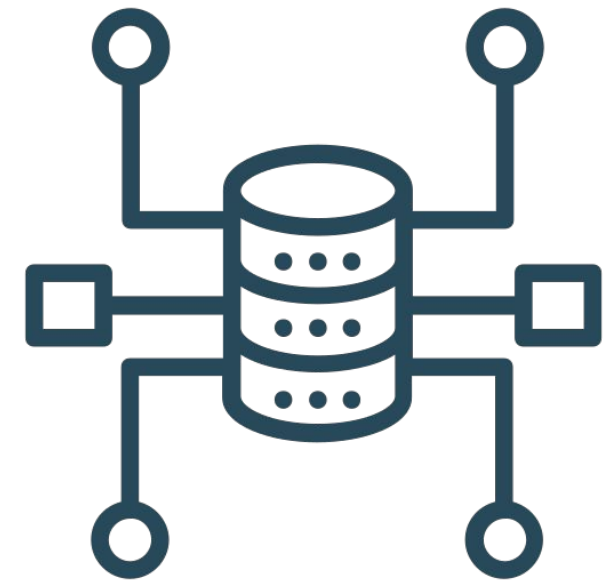


TPA deploys and configures  
robust Postgres architectures



# Configuration

- You must select an architecture and a platform for the cluster.
- An architecture is a recommended layout of servers and software to set up Postgres for a specific purpose.
- Architectures:
  - "M1" (Postgres with a primary and streaming replicas)
  - "PGD-Always-ON" (EDB Postgres Distributed 5 in an Always On configuration).





# Configure

```
tpaexec configure ~/clusters/my-cluster \  
--architecture M1 \  
--postgresql 15 \  
--failover-manager efm \  
--platform bare \  
--hostnames-from ~/clusters/hostnames.txt
```

# Provision

```
tpaexec provision ~/clusters/my-cluster
```

# Deploy

```
tpaexec deploy ~/clusters/my-cluster
```

# Test

```
tpaexec test ~/clusters/my-cluster -v
```

EXAMPLE



# Reference Architectures simplified

Match architecture with requirements

Architecture	Multi-AZ	Multi-Region	HA SLA	RPO=0	RTO=0	Rolling Major Release upgrade*
Single Node	No	No	N/A	No	no	no
Primary/Replica MultiNode	Yes	no	99.99%	yes	no	no
Always-on Single Location	Yes	no	99.999%	yes	yes	yes
Always-on Multi Location	yes	Yes	99.999%	Yes	Yes	yes

**Major release upgrades consideration:** It is a complex task that may require a service outage timeframe longer than the high availability SLA, but it cannot be considered a Major Issue because it is something to be planned in the lifecycle.



Thank you





# Accelerating the Value of Postgres with EDB Customer Services

Darko Milojkovic - Senior Manager, RDBA

Oliver Hofmann - Principal Technical Account Manager

Davide Tamaro - Senior Solution Engineer



## Why use EDB Customer Services?

EDB has largest PostgreSQL Global Technologists Team

**PS Team Dedicated** to Project Implementation Services

Solution Architects  
Database Consultants  
Migration Consultants  
Project Managers

**17+ years'** experience in successfully delivering Postgres migrations

Intellectual Property and tooling around **Migration Assessments and Automation** to accelerate your move to Postgres

**RemoteDBA, Training and TAM** Services for further scale and planning

**Partner Network**

# PROFESSIONAL SERVICES



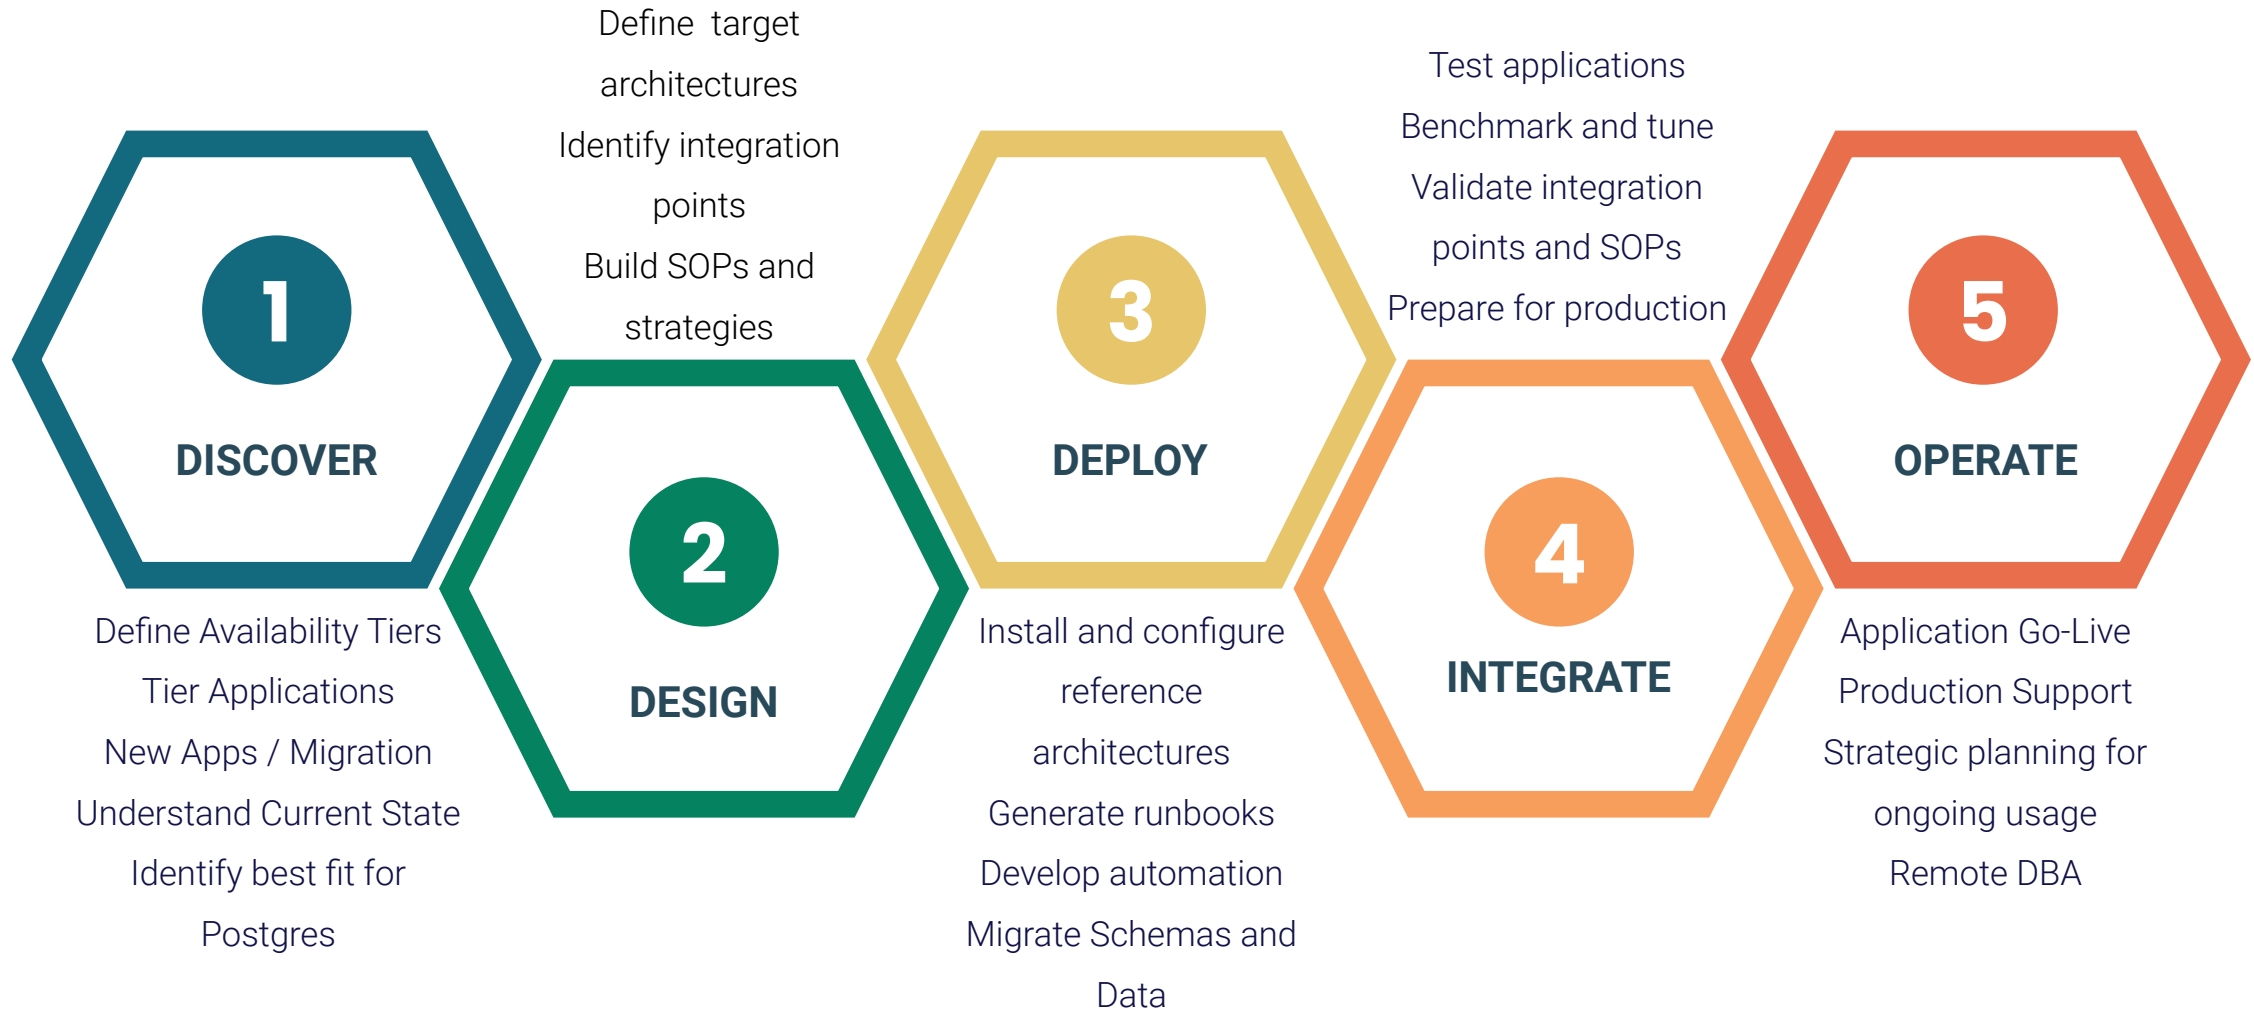


# THE EDB SERVICES APPROACH

- A database **never runs in isolation**
- Ensure a **robust Postgres platform** for migrated and new applications
- Scalability and reusability:
  - Architecture Documentation and SOPs
  - Deployment and Migration Automation
  - Data Migration Playbooks
  - Instructor Led Training
  - Ongoing Knowledge Transfer
  - Remote DBA for scale
  - Technical Account Management for strategic planning



# EDB SERVICES LIFECYCLE FRAMEWORK

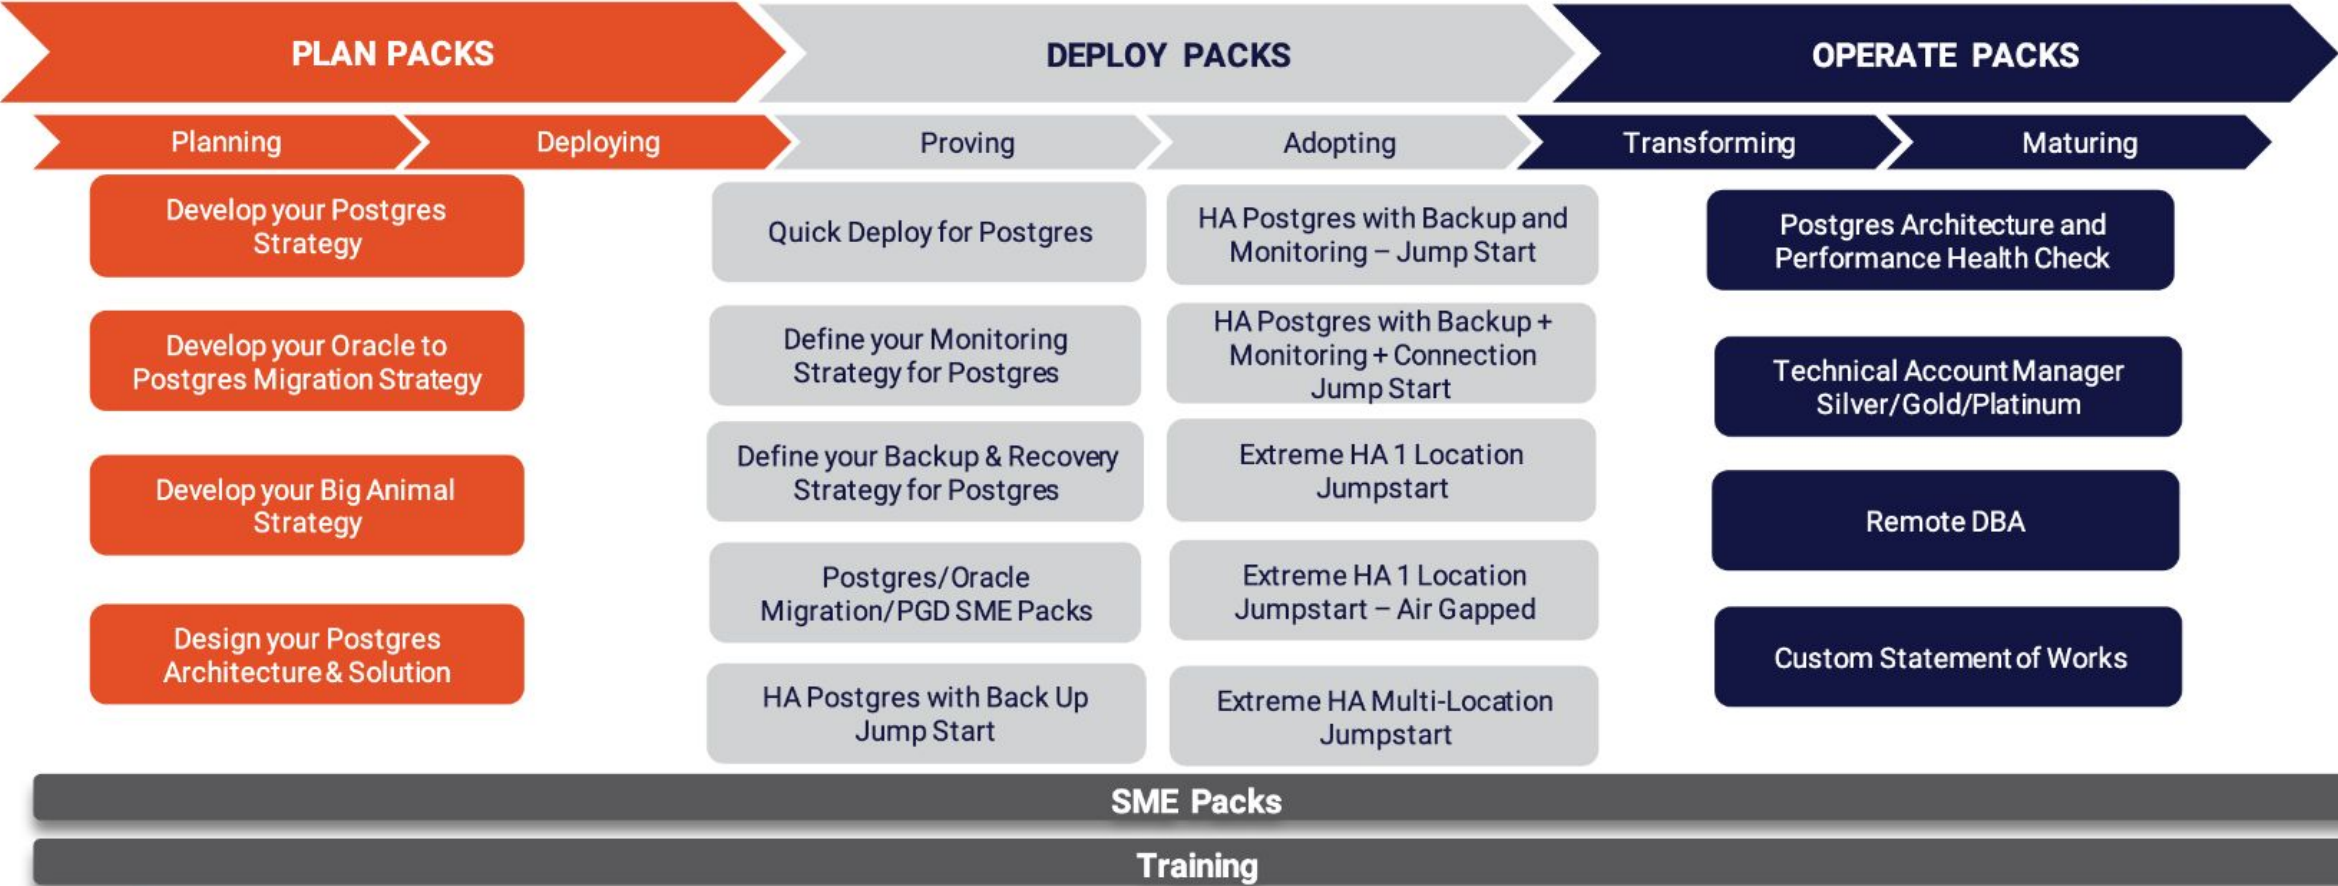


# TRAINING



# CX/PS Catalog - Lifecycle Offerings

CX/PS Services are Designed to accelerate our Customers Postgres Journey



# Why Choose EDB Training?

Training from the World's Leading Postgres Experts

“ EDB training helps drive our business efficiency and boost our team's skills, performance, and results. We're excited that our team is now maximizing their potential. ”

CTO, Leading Financial Institution







## Instructor-led

- Private Group Training
- Instructor-led
- Essentials and Advanced courses
  - EDB Postgres Advanced Server
  - EDB Postgres Distributed
  - PostgreSQL
- Hands-on Labs
- Learn more at [enterprisedb.com/training](https://enterprisedb.com/training)

Have questions or ready to get started?

**Contact Us**





# REMOTE DBA



# Why Choose EDB Remote DBA Service?



Around-the-clock  
assurance

On-premises and  
in the cloud



Premium  
management

More than  
monitoring alone



Timely, affordable,  
and reliable

Proactive,  
on-demand help



Same Postgres  
everywhere

With ultra  
responsive SLO's



Augment your team

PostgreSQL  
Certified  
Engineers  
available 24/7



# EDB Remote DBA Outcomes



Major Cost Savings &  
High ROI's



Disaster Recovery &  
Business Continuity Plan



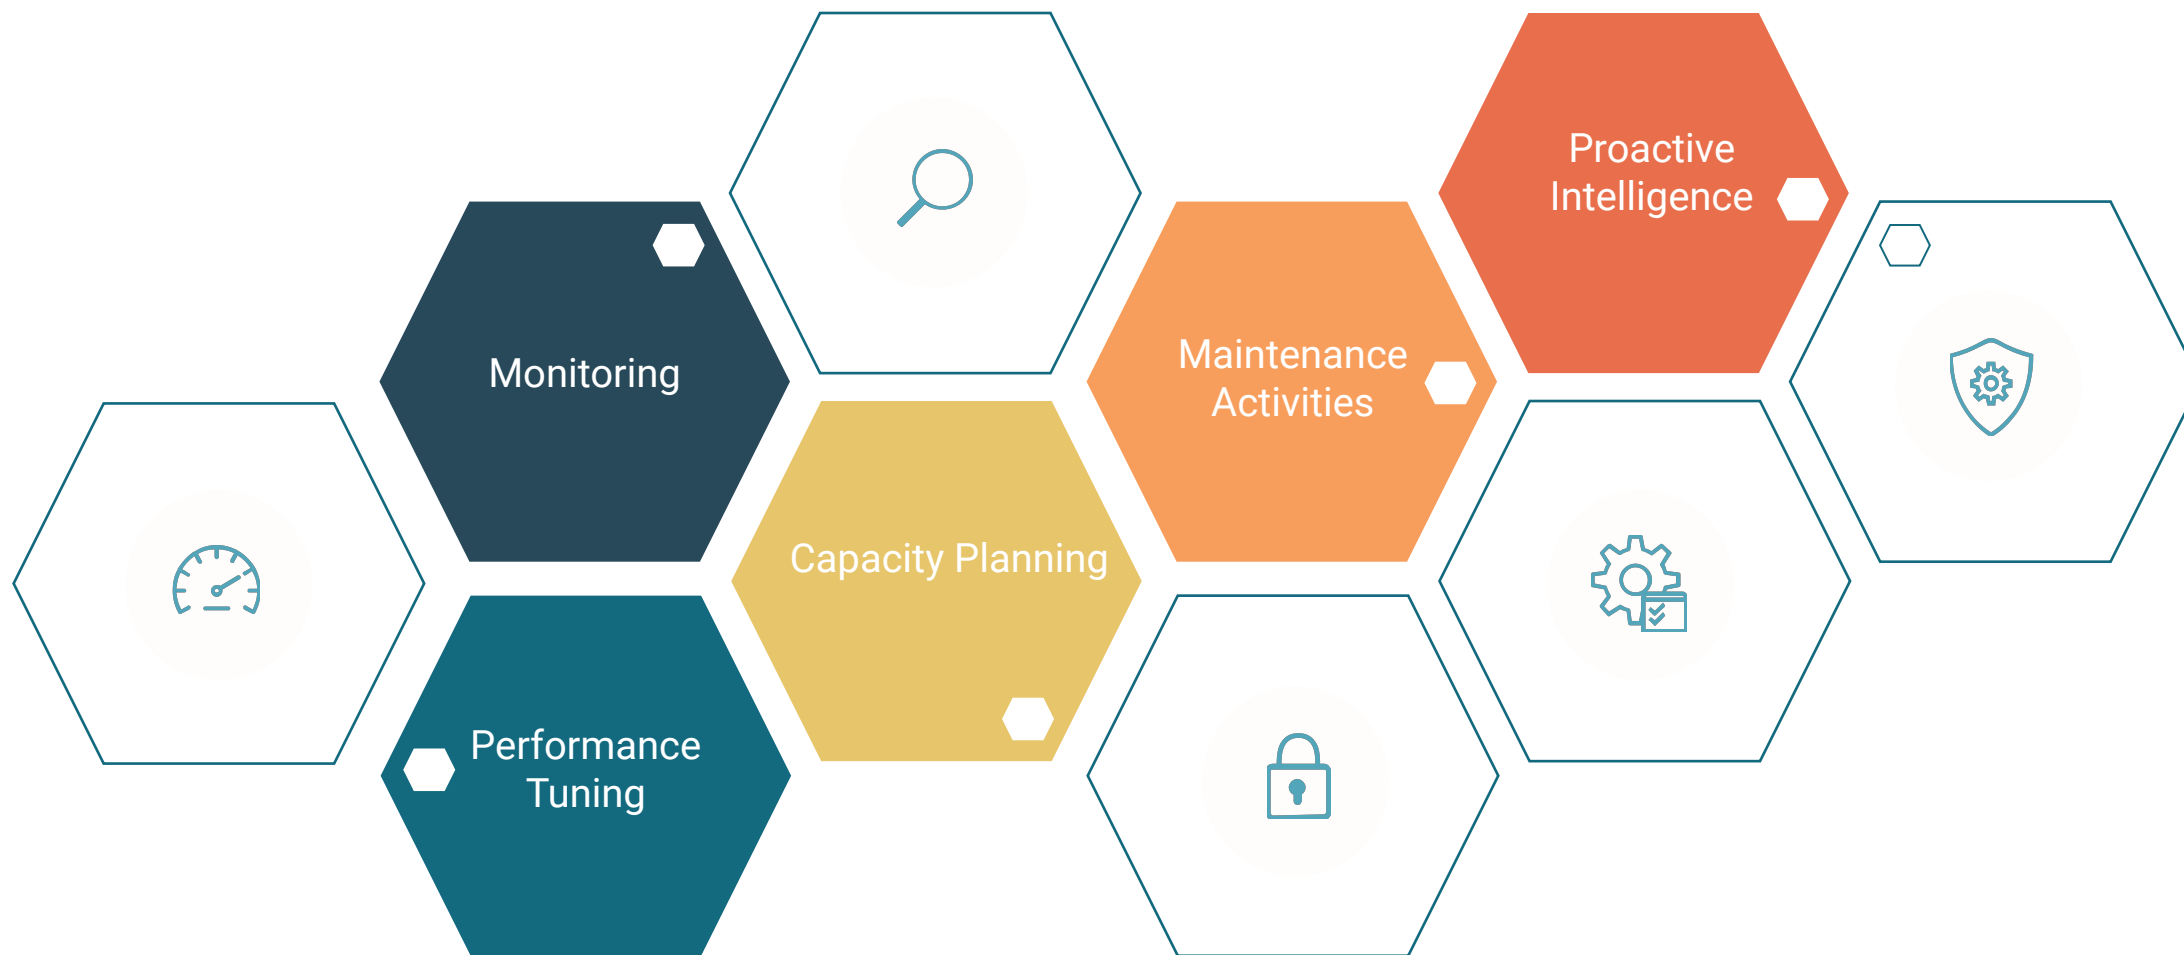
Proven Ability to Execute &  
Certified Expertise



Rapidly Scalable Service &  
SLO Adherence



# EDB Remote DBA Service details



# SOC 2 Certification and benefits



Improves  
Security  
Posture



Build  
Customers  
Trust



Reduces Risk  
of Data  
Breaches



Serves a  
Marketing  
Differentiator



Demonstrates  
Commitment  
to Security



# EDB Remote DBAs are Customer Focused

95.4%

Customer Satisfaction  
Score

100%

Initial response SLO achievement for  
Sev 1 and Sev 2

100%

Remedy SLO achievement for Sev 1  
and Sev 2

96%

Retention rate



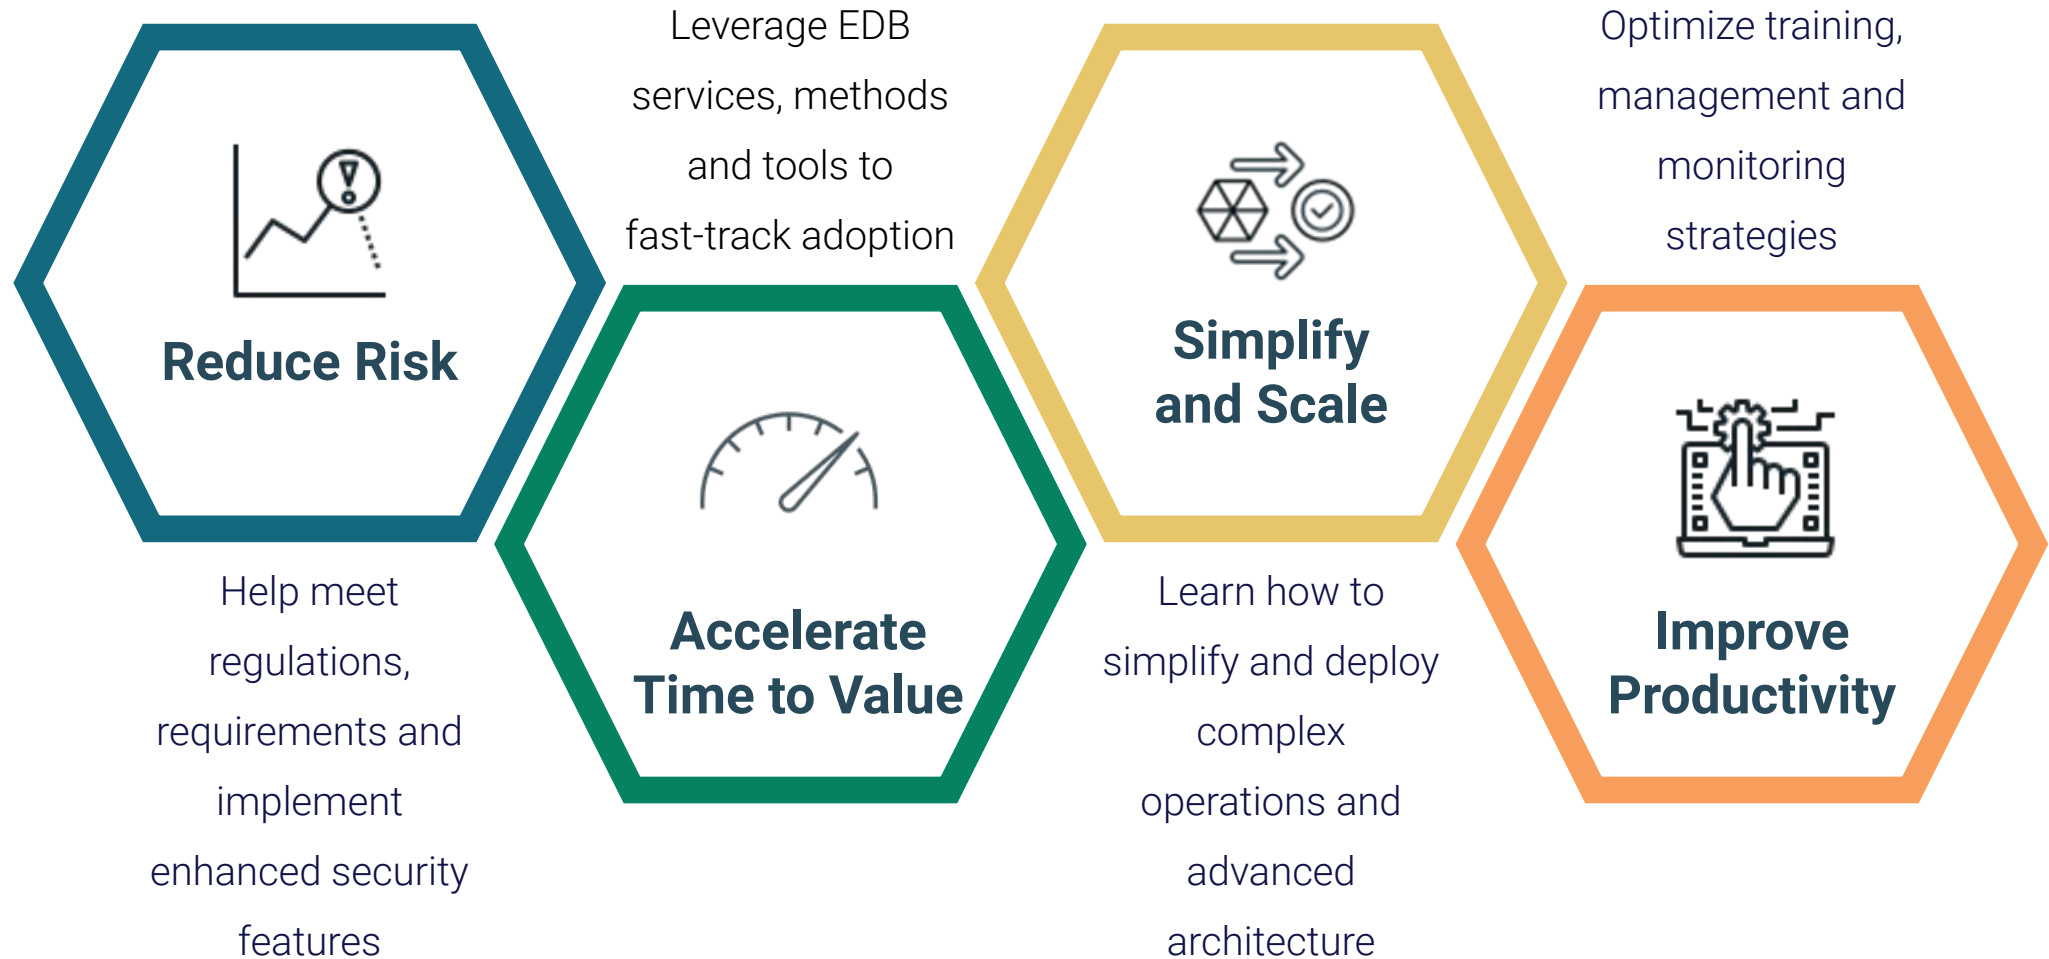
# TAM - Technical Account Management

TAM is a paid EDB technical resource who partners with the client's technical team and proactively guides them towards a successful execution of their technical initiatives.



# Technical Account Manager Value


Proactive assistance helps prepare for technology changes and ensures customer gets the most from Postgres, their EDB product, solution, and support



# Technical Account Manager Delivery and Outcomes



**Accelerate innovation with database initiative prioritization strategies**



**Case/Support Management to reduce downtime and increase client satisfaction**



**Upscale client's workforce and retain talent through coaching**



**Increase adoption with planning and selecting proprietary and Open Source Tools**



**Approach to Large Projects and Major Version Upgrade Strategy**



**Ensuring Postgres success by facilitating client's Postgres roadmap**



# Common questions

**I already have support,  
why do I need TAM?**

TAM delivers proactive strategic guidance and planning to optimize your Postgres estate and return the highest ROI. EDB technical product support is reactive break/fix support.

**Do unused hours  
carryover?**

Your TAM is expected to be routinely engaged with you which reduces the prospect of unused hours that do not carryover.

**I have a Customer  
Success Manager, why do  
I need a TAM?**

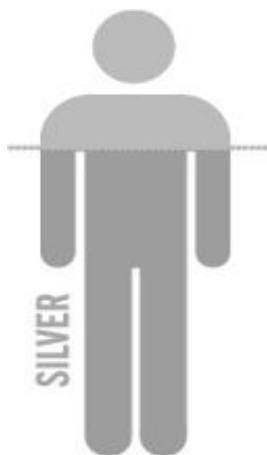
The CSM is responsible for operational engagement between the customer and EDB focused on business outcomes and value. A TAM is responsible for engagement to help drive technical outcomes and strategize with customers to maximize technical value from EDB solutions.

**What is an EDB Executive  
Sponsor?**

The TAM Platinum Executive Sponsor is a senior leader that helps ensure effective communication, strategic alignment and an avenue for escalation.



# Technical Account Management Options



Hours of work / month	Up to 16	Up to 40	Up to 60
Review cadence	Monthly	Weekly	Weekly
Planning	Quarterly, online	Quarterly, on-site	Twice quarterly, on-site
EDB Executive Sponsor	-	-	√







**EDB**

Postgres® for the AI Generation

# Grazie e Arrivederci Presto

## Contacts

- Marco Cesare, Enterprise Business Leader  
Email: [marco.cesare@enterprisedb.com](mailto:marco.cesare@enterprisedb.com)  
Phone: +39 335 6519644
- Roberto Serravento, Public sector Business Leader  
Email: [roberto.serravento@enterprisedb.com](mailto:roberto.serravento@enterprisedb.com)  
Phone: +39 348 0712205
- Mario Vetrano, Cloud Specialist EMEA  
Email: [mario.vetrano@enterprisedb.com](mailto:mario.vetrano@enterprisedb.com)  
Phone: +39 392 0652539

## Case Studies

- [Banca Popolare di Sondrio Streamlines Innovation with Postgres and EDB](#)
- [Zucchetti Works with EDB to Implement Cost-Effective, Scalable Database Archiving Solution for Electronic Invoices](#)

[Website](#)

[Follow us](#)