Red Hat Summit

Connect

EDB and Red Hat OpenShift for Your Sovereign AI and Data Initiatives

Davide Tammaro | Principal Sales Engineer, EDB
Natale Vinto | Technical Director, Red Hat

EDB | Red Hat

Gen AI adoption rates across the enterprise are rapidly increasing

Gen AI is a core part of the products we use, at home and at work

Frontier AI services grow new capabilities at a relentless speed

Customize your preferred model using enterprise data to build an efficient, cost-effective solution.

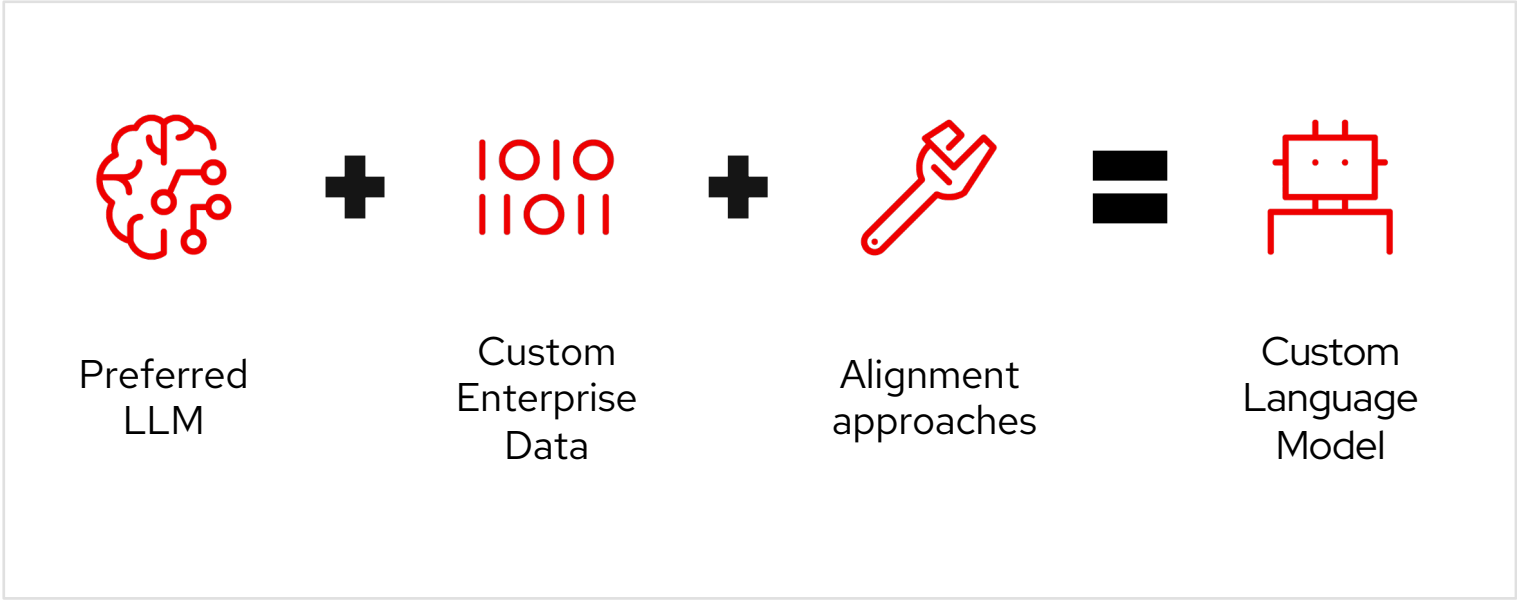

Red Hat

What's next: multiple approaches?

Build customized AI solutions that address domain specific business cases

Agentic AI

Agents, MCP servers, orchestration

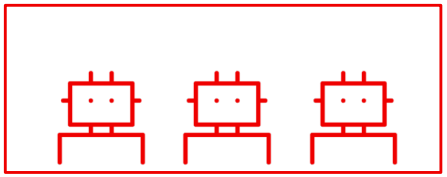

A new wave of AI application It's the framework that enables multiple agents to collaborate and adapt to solve complex problems on their own.

RAG

Retrieval Augmented Generation

Enhance Gen AI model generated text by retrieving relevant information from external sources, improving accuracy and depth of model's responses.

Fine tuning

InstructLab, OSFT, LoRA and QLoRA

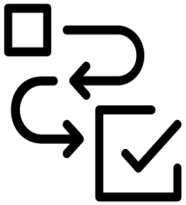

Customize a base model for specific tasks or private data, using a range of approaches—from full fine-tuning to parameter-efficient methods—to balance performance and efficiency.

Red Hat

Data Gravity

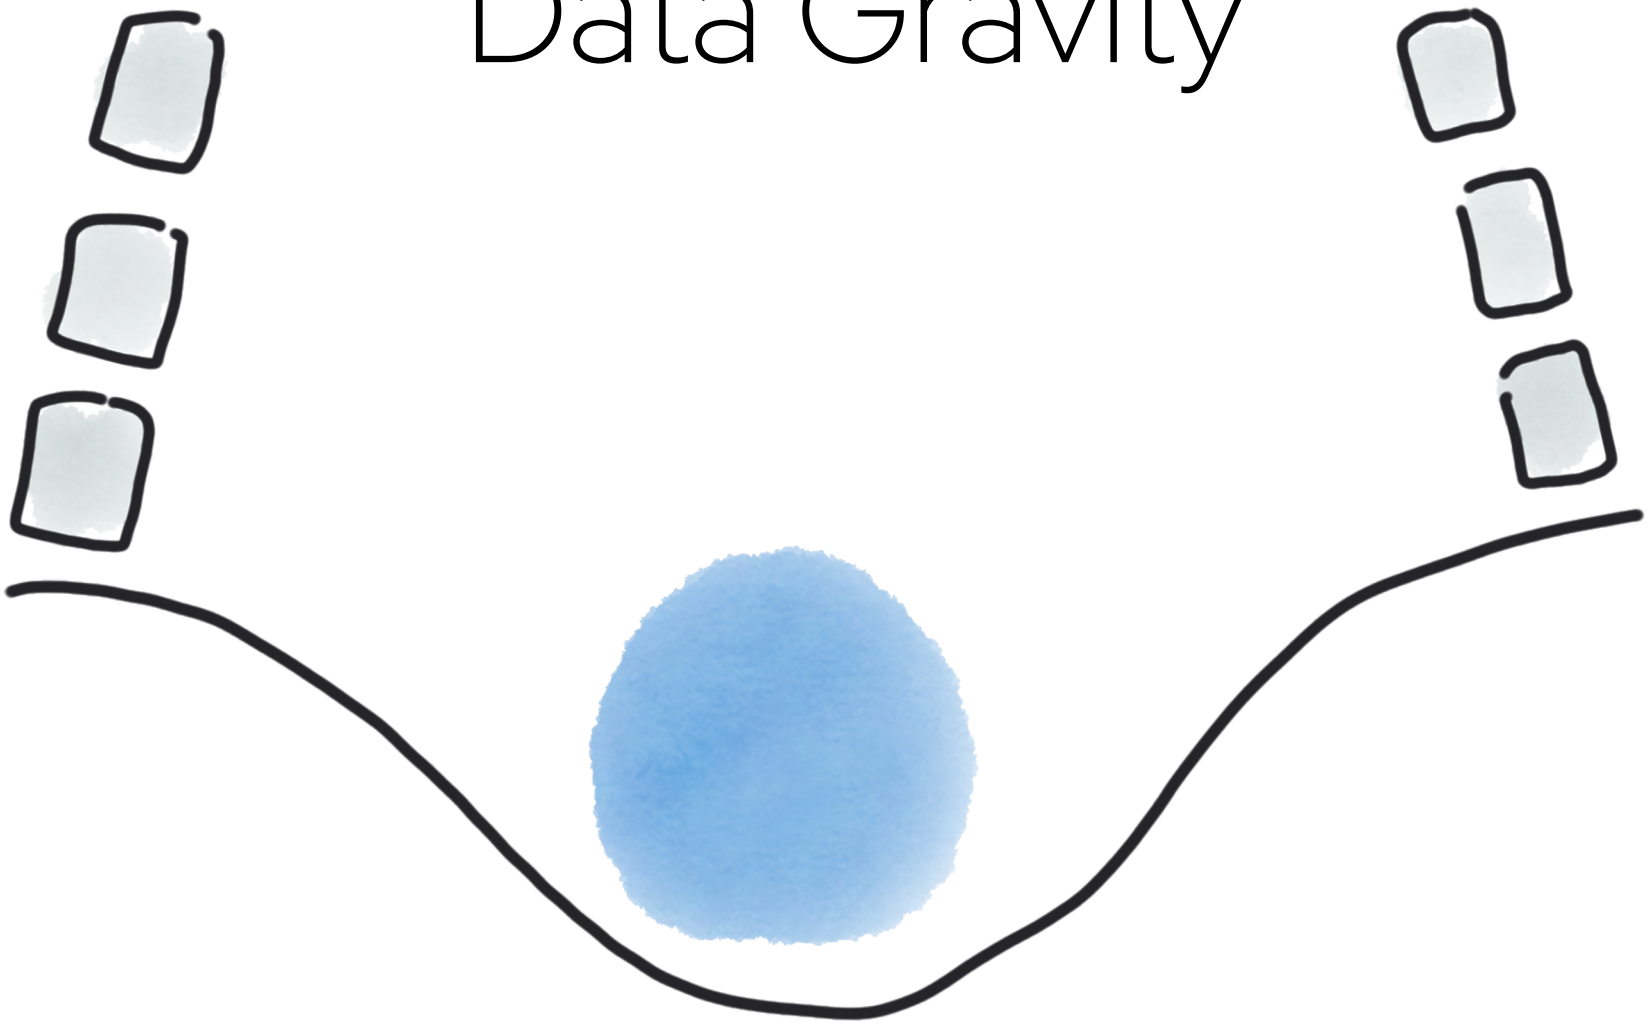

Scale and optimize your AI and application deployments

Existing investments must work in support of AI

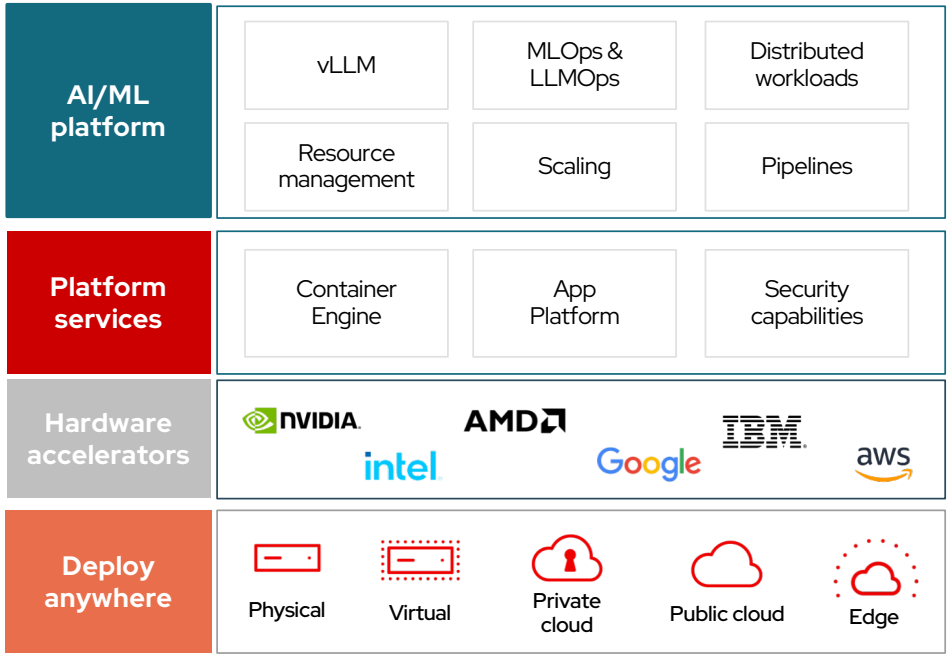

- **Integrate to real workflows** with access to data sources, workloads and applications.
- **Think of day 2 operations** for governance, management and automation.
- **Scale AI workloads dynamically** across hybrid cloud using Kubernetes, including horizontal and GPU scaling with automated resource management to meet fluctuating demands.

Red Hat

AI safety, monitoring and observability

Track accuracy, biases, performance, and more

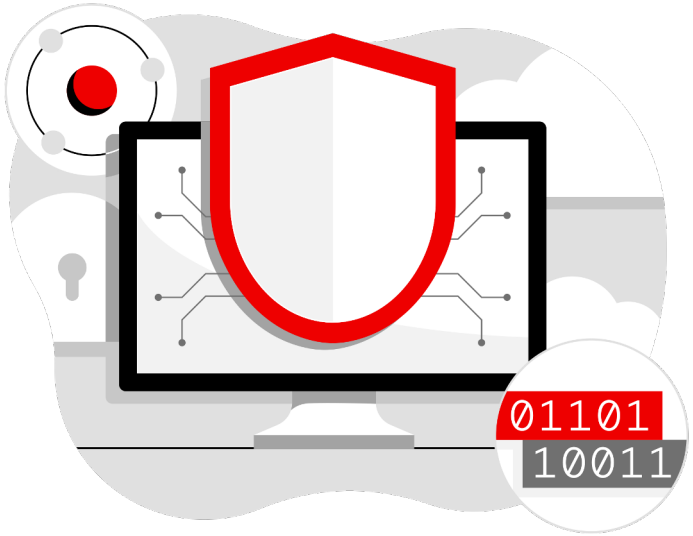

Detect bias and drift: Monitor for outcome disparities and training data differences.

Guardrails: A customizable framework to moderate interactions between users and generative AI.

Model monitoring: Track operations and performance metrics.

Accuracy evaluation: Measure model knowledge across topics.

Experiment tracking: Gain visibility and confidence when experimenting with models (visuals, metrics, UI/UX)

Red Hat OpenShift and Open Hybrid Cloud

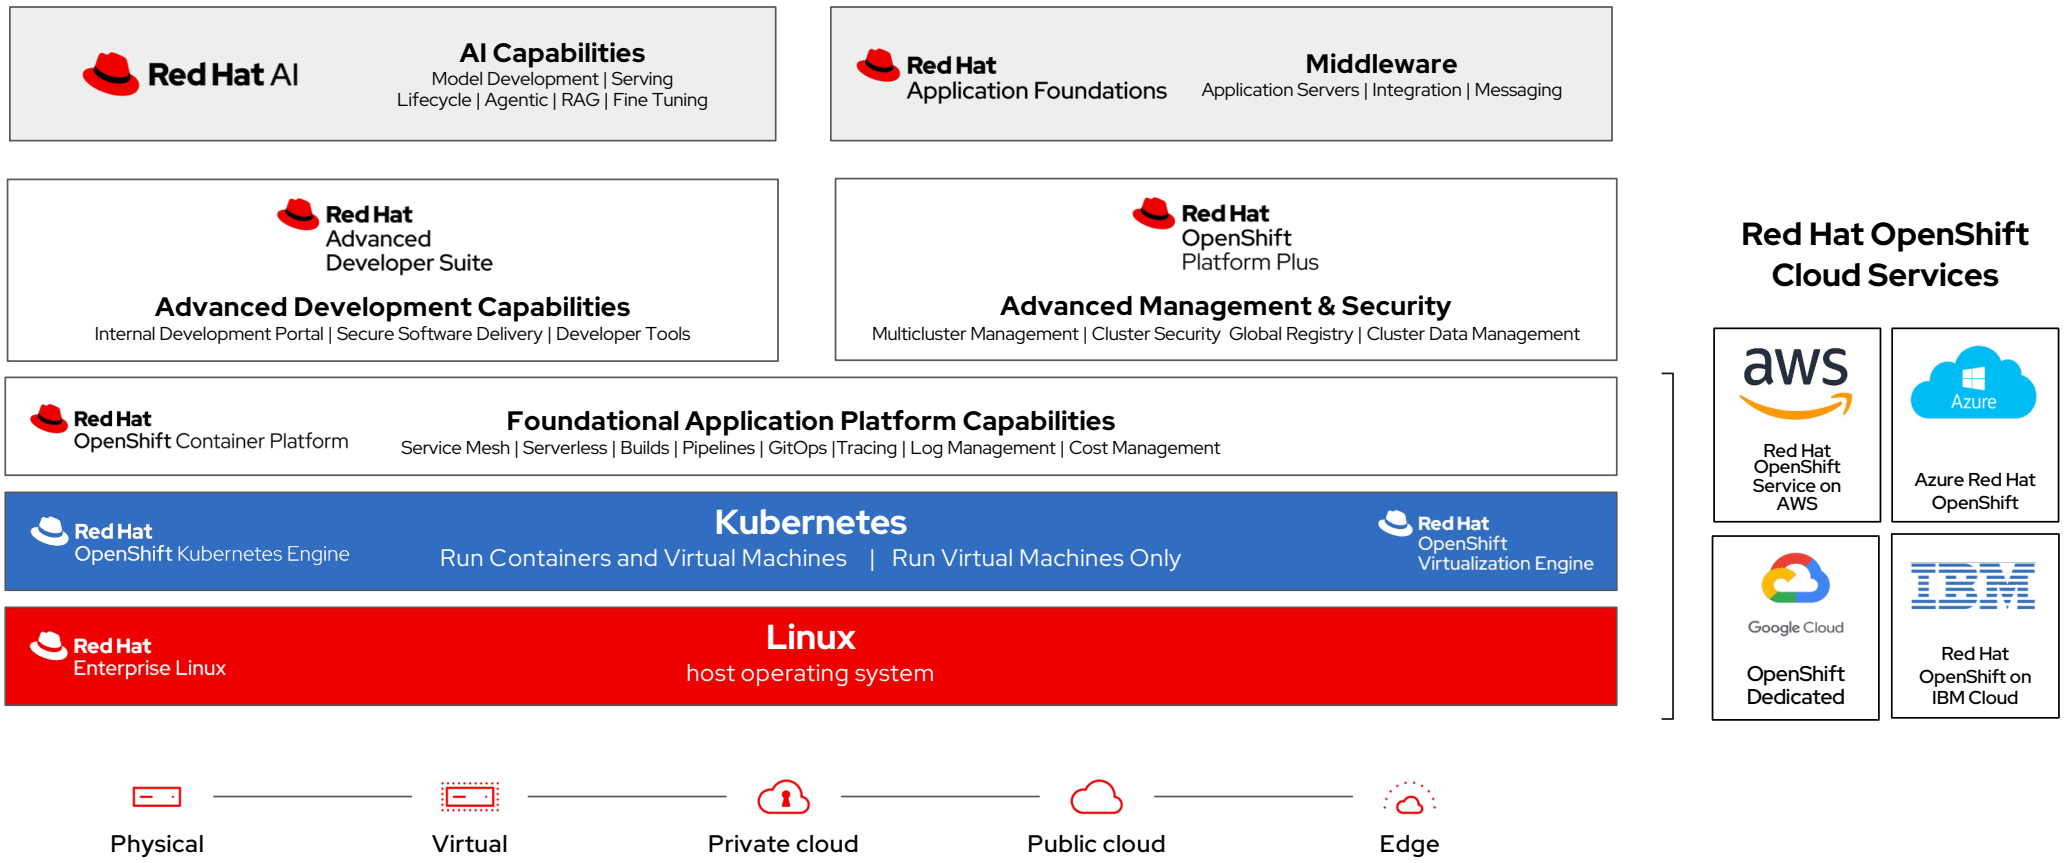

Red Hat OpenShift Cloud Services

Red Hat

Red Hat AI

Accelerate the development and delivery of AI solutions across hybrid-cloud environments

Increase efficiency with **fast, flexible and efficient inferencing**

Simplified and consistent experience for **connecting models to data**

Accelerate Agentic AI deployments

Flexibility and consistency when **scaling AI across the hybrid cloud**

Red Hat

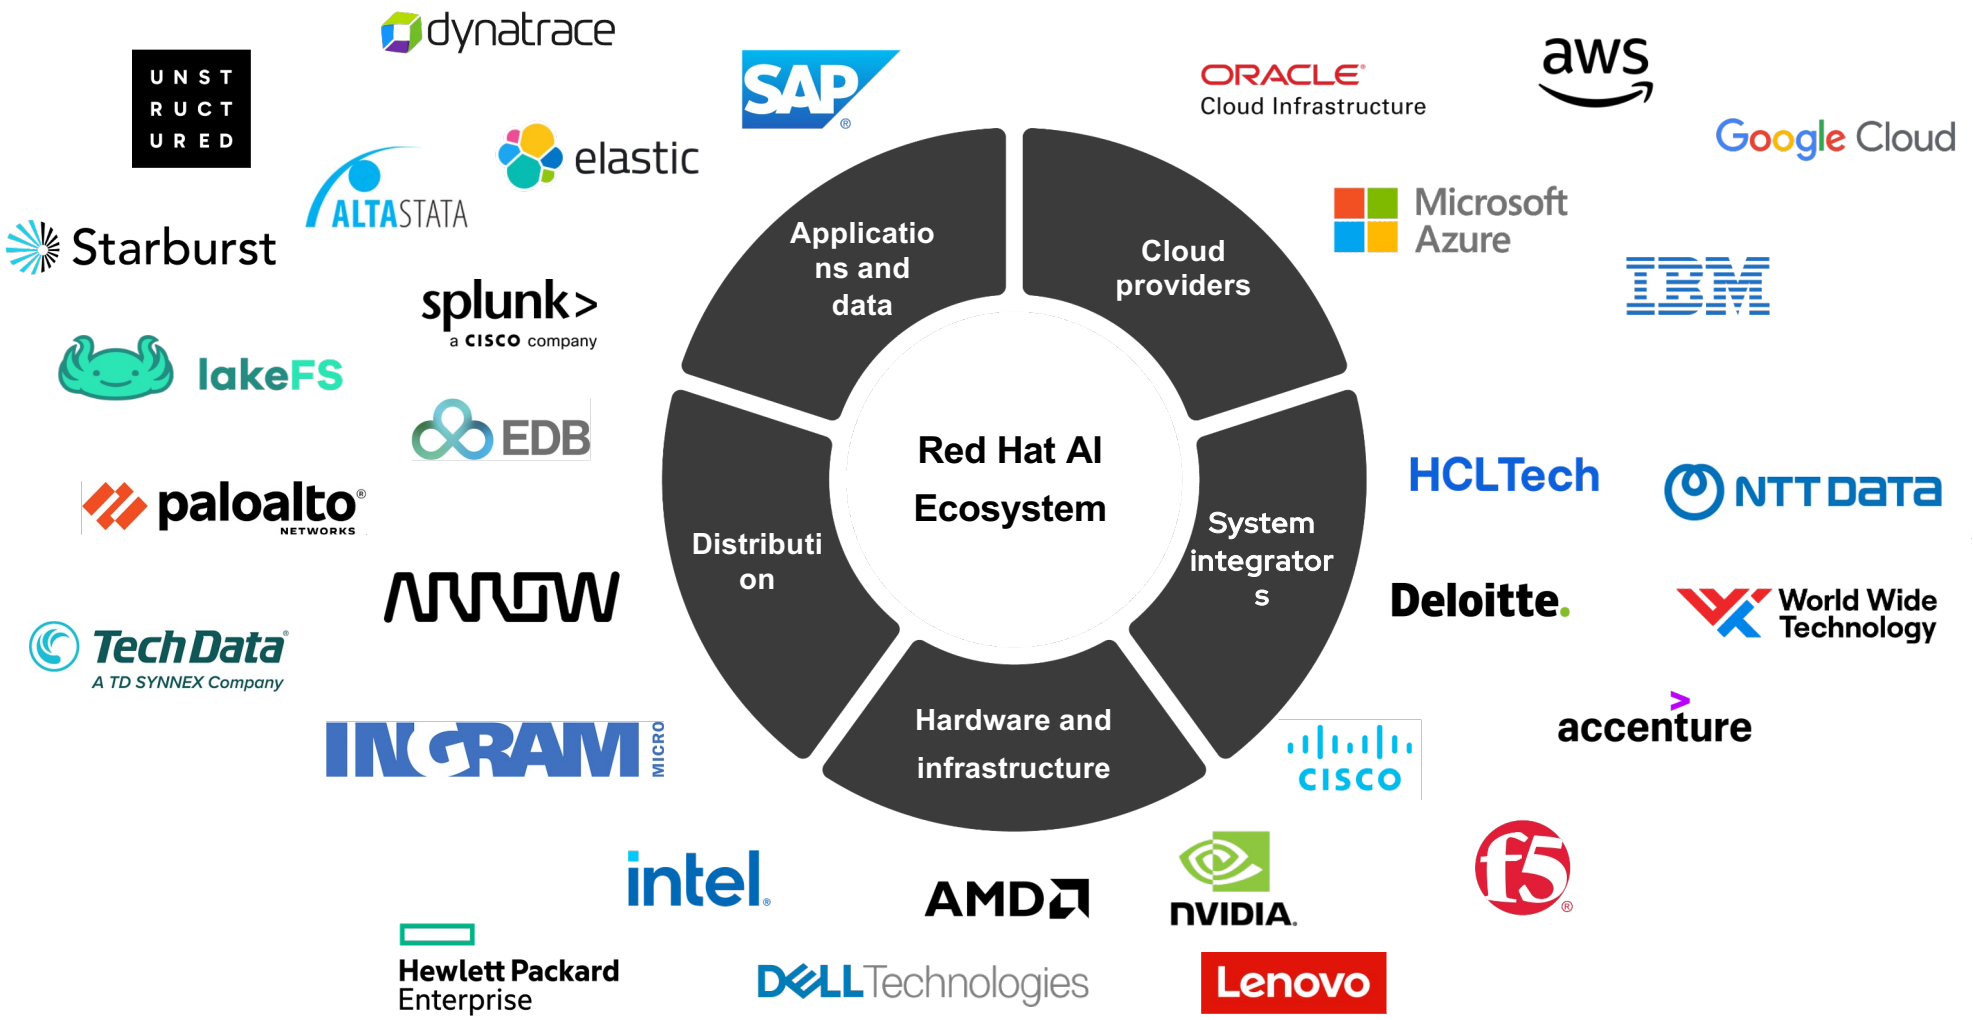
Applications and data
Cloud providers
System integrators
Hardware and infrastructure
Distribution
Red Hat AI Ecosystem
dynatrace
UNSTRUCTURED
SAP
ORACLE Cloud Infrastructure
aws
Google Cloud
elastic
ALTASTATA
Microsoft Azure
Starburst
IBM
splunk> a CISCO company
lakeFS
EDB
paloalto NETWORKS
HCLTech
NTT DATA
ARROW
Deloitte.
World Wide Technology
TechData A TD SYNNEX Company
accenture
INGRAM MICRO
CISCO
intel
AMD
NVIDIA
f5
Hewlett Packard Enterprise
DELL Technologies
Lenovo

Red Hat

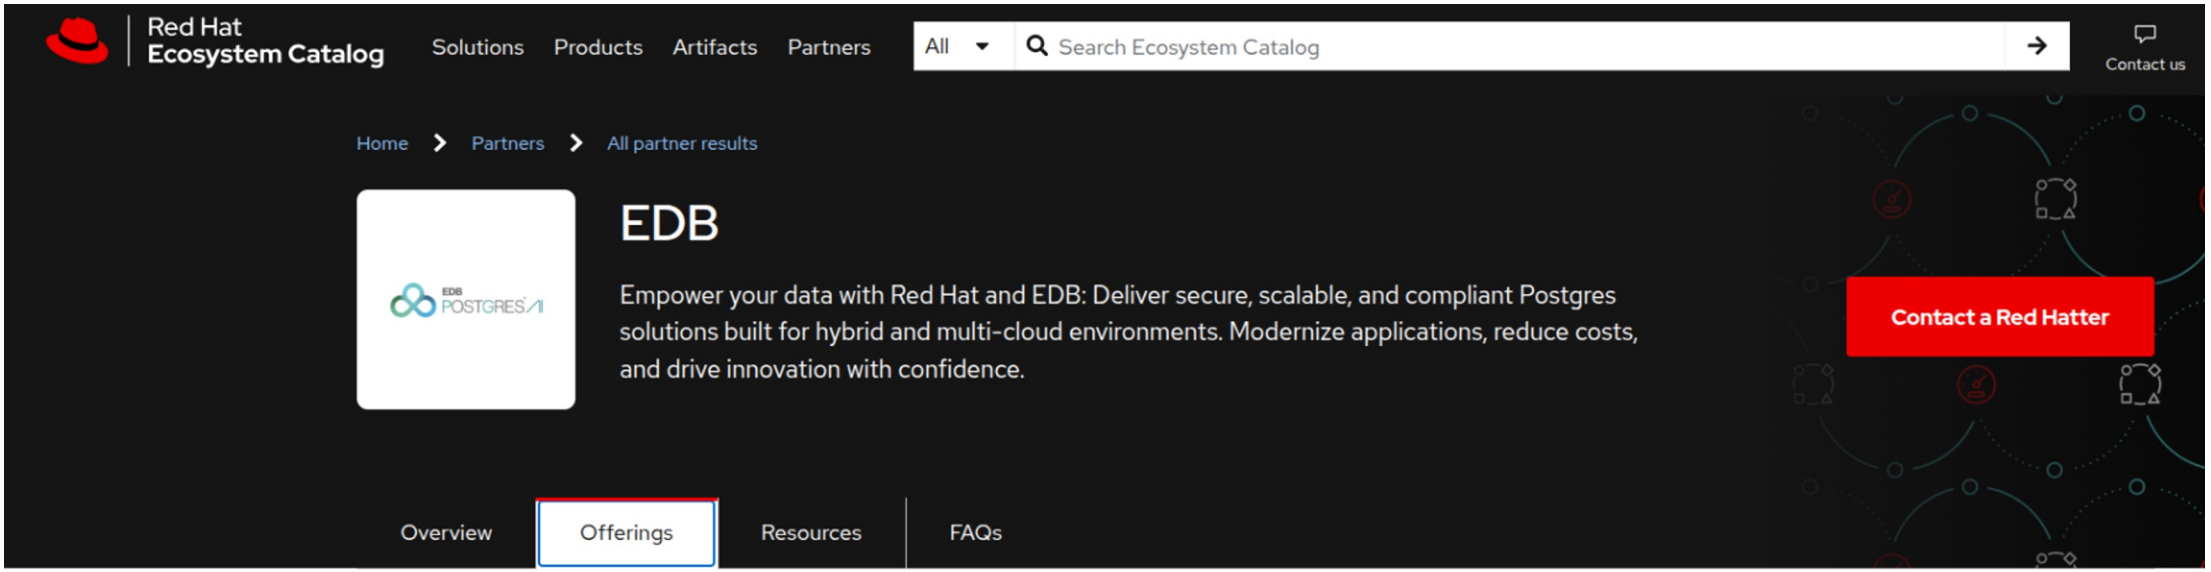

Search

1 - 6 of 6

EDB Postgres for Kubernetes

PostgreSQL Operator for mission critical databases in Openshift Container Platform

Containerized application Storage

EDB

Published 4 giorni fa

EDB Postgres Distributed

EDB Postgres Distributed for Kubernetes is an operator designed to manage EDB Postgres Distributed (PGD) v5 workloads on Kubernetes both in a single Kubernetes cluster or across multiple ones.

Containerized application Storage

EDB

Published 25 giorni fa

EDB Postgres AI

EDB

RAG (Retrieval augmented generation) provides extra info

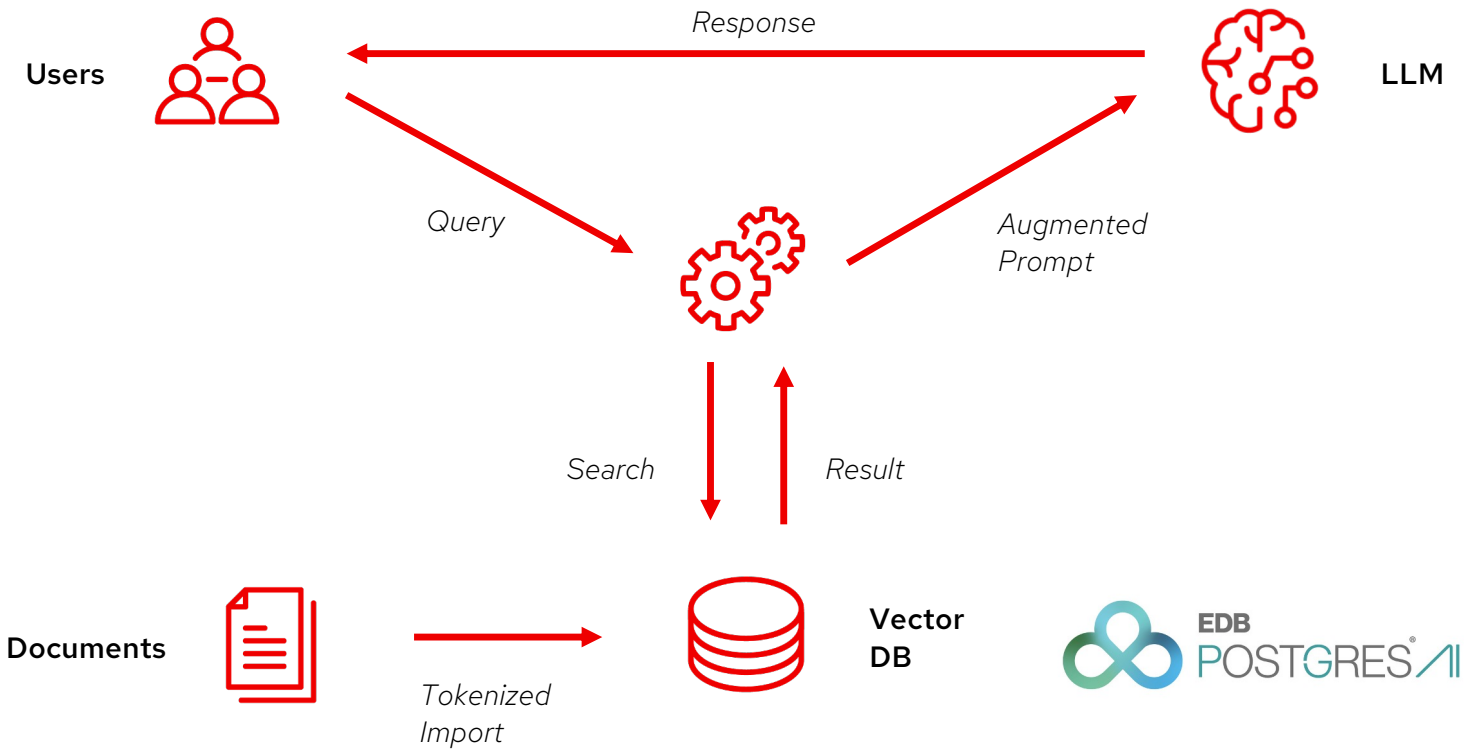

Red Hat

SOVEREIGN AI AND DATA PLATFORM FOR PRODUCTION-READY AI

EDB Postgres AI and Red Hat OpenShift AI work seamlessly together to provide a comprehensive platform

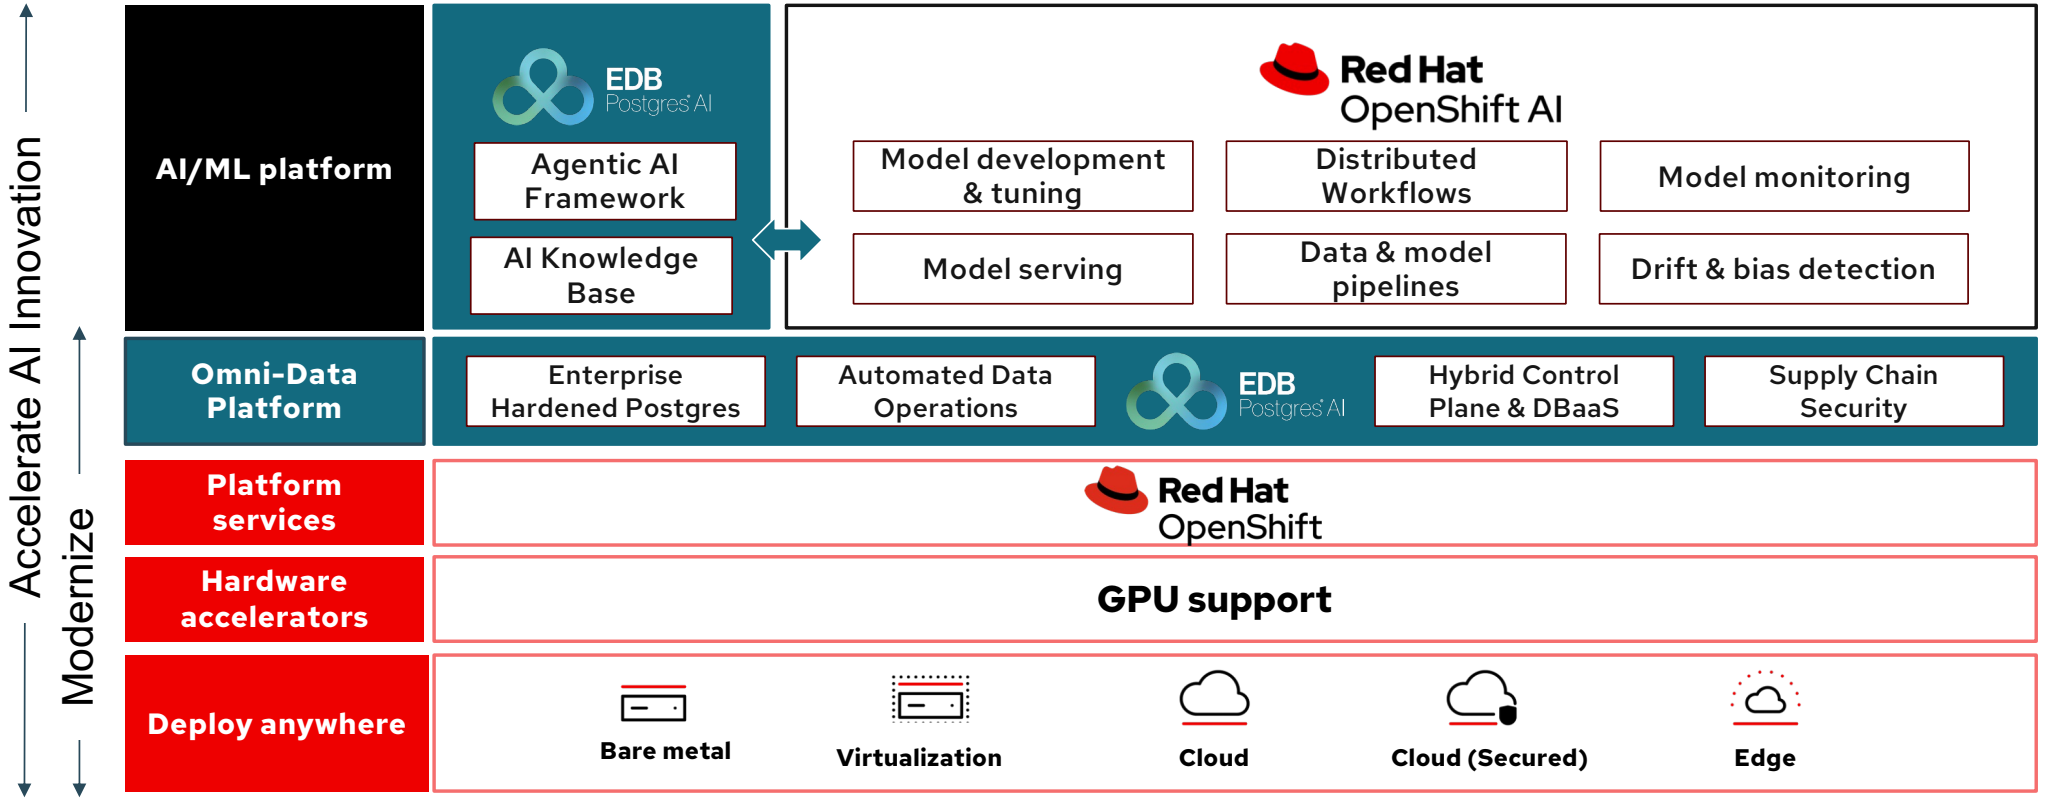

EDB

Red Hat

EnterpriseDB Postgres AI

Database

Database Servers

- Enterprise Postgres
- Oracle Compatible
- Community PostgreSQL

- Multi-Model Extensions
- Management & Observability (Basic)
- Kubernetes Operators
- Supply Chain Security
- Migration Toolkit
- High Availability

Analytics Accelerator

- Analytics
  Columnar Query Engine
  MPP Warehouse
- Lakehouse Storage
  Delta Tables, Iceberg
- Support for Greenplum Workloads

AI Factory

- Vector Engine
- AI Pipeline
- GenAI Builder
- Agent Orchestrator
- Model Serving

Hybrid Management

- Hybrid Observability
- Hybrid DBaaS
- Distributed HA (99.999%)
- Migration Center

Deploy Anywhere

ENGINEERED SYSTEM | HYBRID SOFTWARE | MANAGED CLOUD

Deployment Partners: AWS, GCP, Azure, Red Hat, IBM, Supermicro

Integration Services (CX)

Solution Architecture | Pilot to Production SLAs | Managed Services

The Postgres® Vitality Index

A MEASURE OF TOTAL INVESTMENT IN ONGOING VITALITY AND FUTURE OF POSTGRES, BASED ON THE CONTRIBUTIONS OF TEN LEADING COMPANIES.

EDB VITALITY SHARE

METHODOLOGY

Scores were calculated using publicly available data sources from across the PostgreSQL community. The sample set includes the top nine most commercially recognized companies, based on their recorded contributions. Each company's percentage of contribution was determined as a proportion of the total contributions within this sample set, with all percentages expressed out of 100%.

*AVG. ACROSS FOUR CATEGORIES COMPRISED OF 16 FACTORS

EDB POSTGRES AI

VITALITY INVESTMENT

EDB INVESTS AS MUCH IN POSTGRES VITALITY AS THE NEXT NINE LEADING COMPANIES COMBINED*.

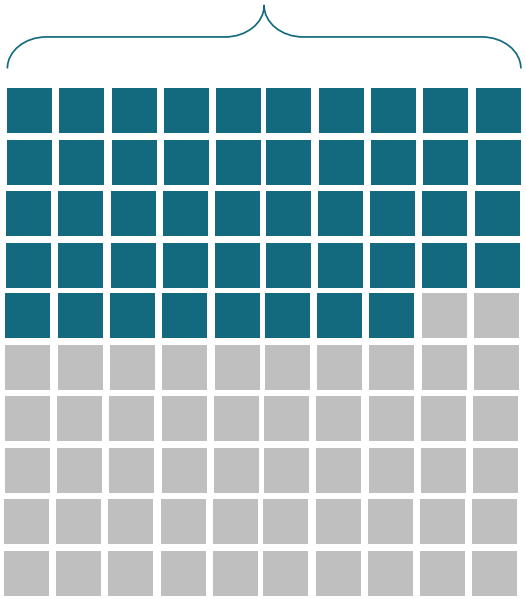

CONTRIBUTIONS OF NEXT 9 LEADING COMPANIES

ELEMENTS OF THE POSTGRES VITALITY INDEX

Postgres Core Contributions

Evaluates impact based on authored and committed code, PostgreSQL 17 patch authorship, and recognized feature contributions.

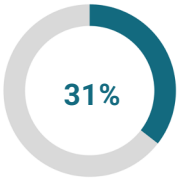

Supported Ecosystem Tools & Extensions

Assesses contributions through delivered packages and installers, maintained build farm machines, and leadership in Release Management Teams over the past five years.

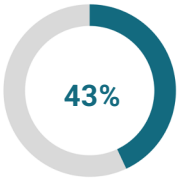

Community Nurture

Includes new patch authors, emerging PostgreSQL contributors, recognized community members, and NPO PostgreSQL Board representation.

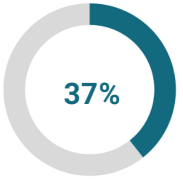

Enterprise Augmentation

Evaluates hybrid support options, mission critical availability (99.999% uptime), seamless online upgrades, transparent data encryption (TDE), and extensibility for analytical and AI workloads.

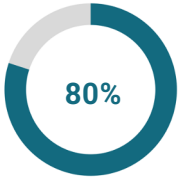

Disclaimer

For simplicity, "Kubernetes" is used as a general term to refer to concepts that also apply to Red Hat OpenShift.

Why shall I consider running Postgres on in Kubernetes?

Innovation, data sovereignty, cloud neutrality, AI, data portability, vendor lock-in mitigation, and TCO reductions are fueling the growing adoption of Postgres databases on Kubernetes.

Kubernetes is a great orchestrator and we train it to orchestrate Postgres…

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 masking	Scale out / scale down
PostGIS	Logging	Volume Snapshots	EDB Hybrid Manager	Virtual Private DB	Minor rolling updates
	Pooling			SQL Inject protection	Major updates

Evolution of PostgreSQL in containers

From Docker system containers to Kubernetes native databases with CloudNativePG

2013/3: Docker is released. Postgres runs mainly for testing in system containers

2015/7: Kubernetes 1.0 is released. Stateless applications only.

2016/11: Operator pattern by CoreOS

2017/3: Crunchy Data releases the first Postgres operator based on Patroni

2017/12: Statefulsets are introduced in Kubernetes 1.9 (*1 year after beta in 1.5*)

2018/8: Zalando releases their operator

2019/4: Local persistent volumes are introduced in Kubernetes 1.14

2019/8: The Cloud Native initiative at EDB (2ndQuadrant at that time) begins

2021/2: EDB launches Cloud Native Postgres

2022/5: EDB open sources CloudNativePG

2024/10: CloudNativePG reaches 4500 stars on GitHub (#1 Postgres operator)

2025/1: **CloudNativePG** becomes a **CNCF** project entering the **Sandbox**

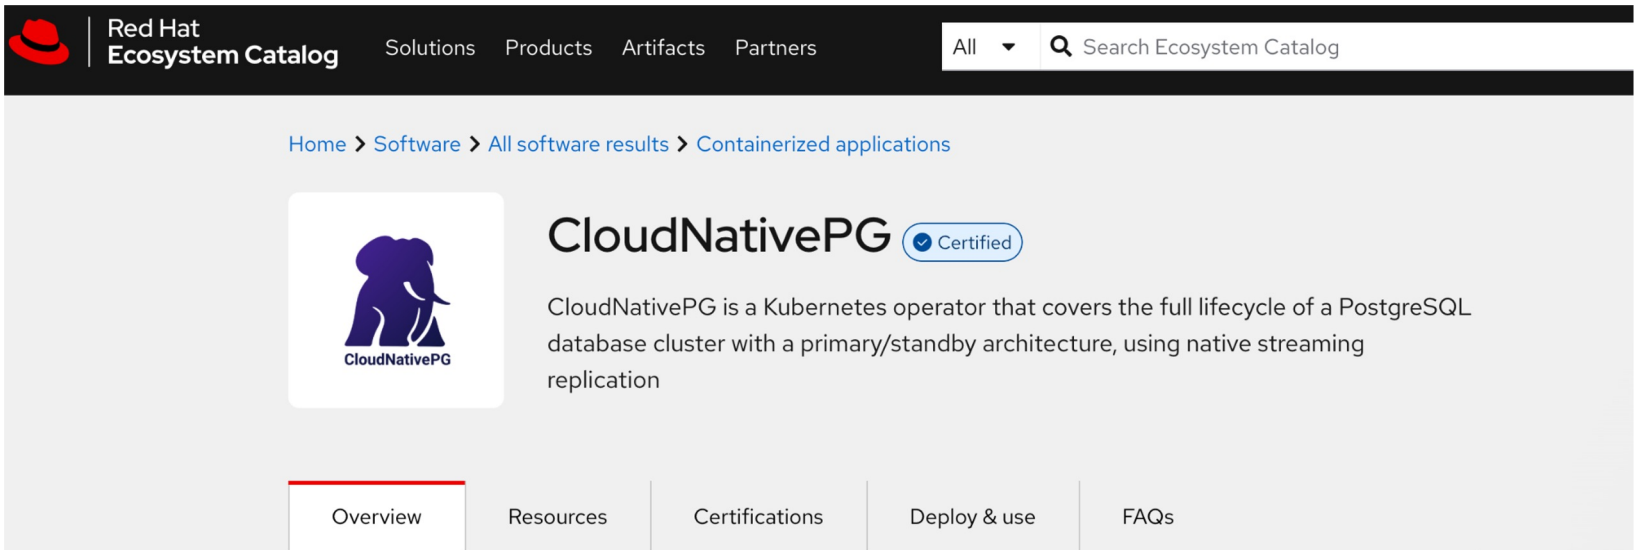

CloudNativePG is a Kubernetes operator that covers the full lifecycle of a PostgreSQL database cluster with a primary/standby architecture, using native streaming replication

→ Self-Healing and automated failover

In case of detected failure on the primary, the operator will change the status of the cluster by setting the most aligned replica as the new target primary. As a consequence, the instance manager in each alive pod will initiate the required procedures to align itself with the requested status of the cluster, by either becoming the new primary or by following it. In case the former primary comes back up, the same mechanism will avoid a split-brain by preventing applications from reaching it, running `pg_rewind` on the server and restarting it as a standby. Self-healing is enhanced by the automated recreation of a standby: in case the pod hosting a standby is removed, the operator initiates the procedure to recreate a standby server.

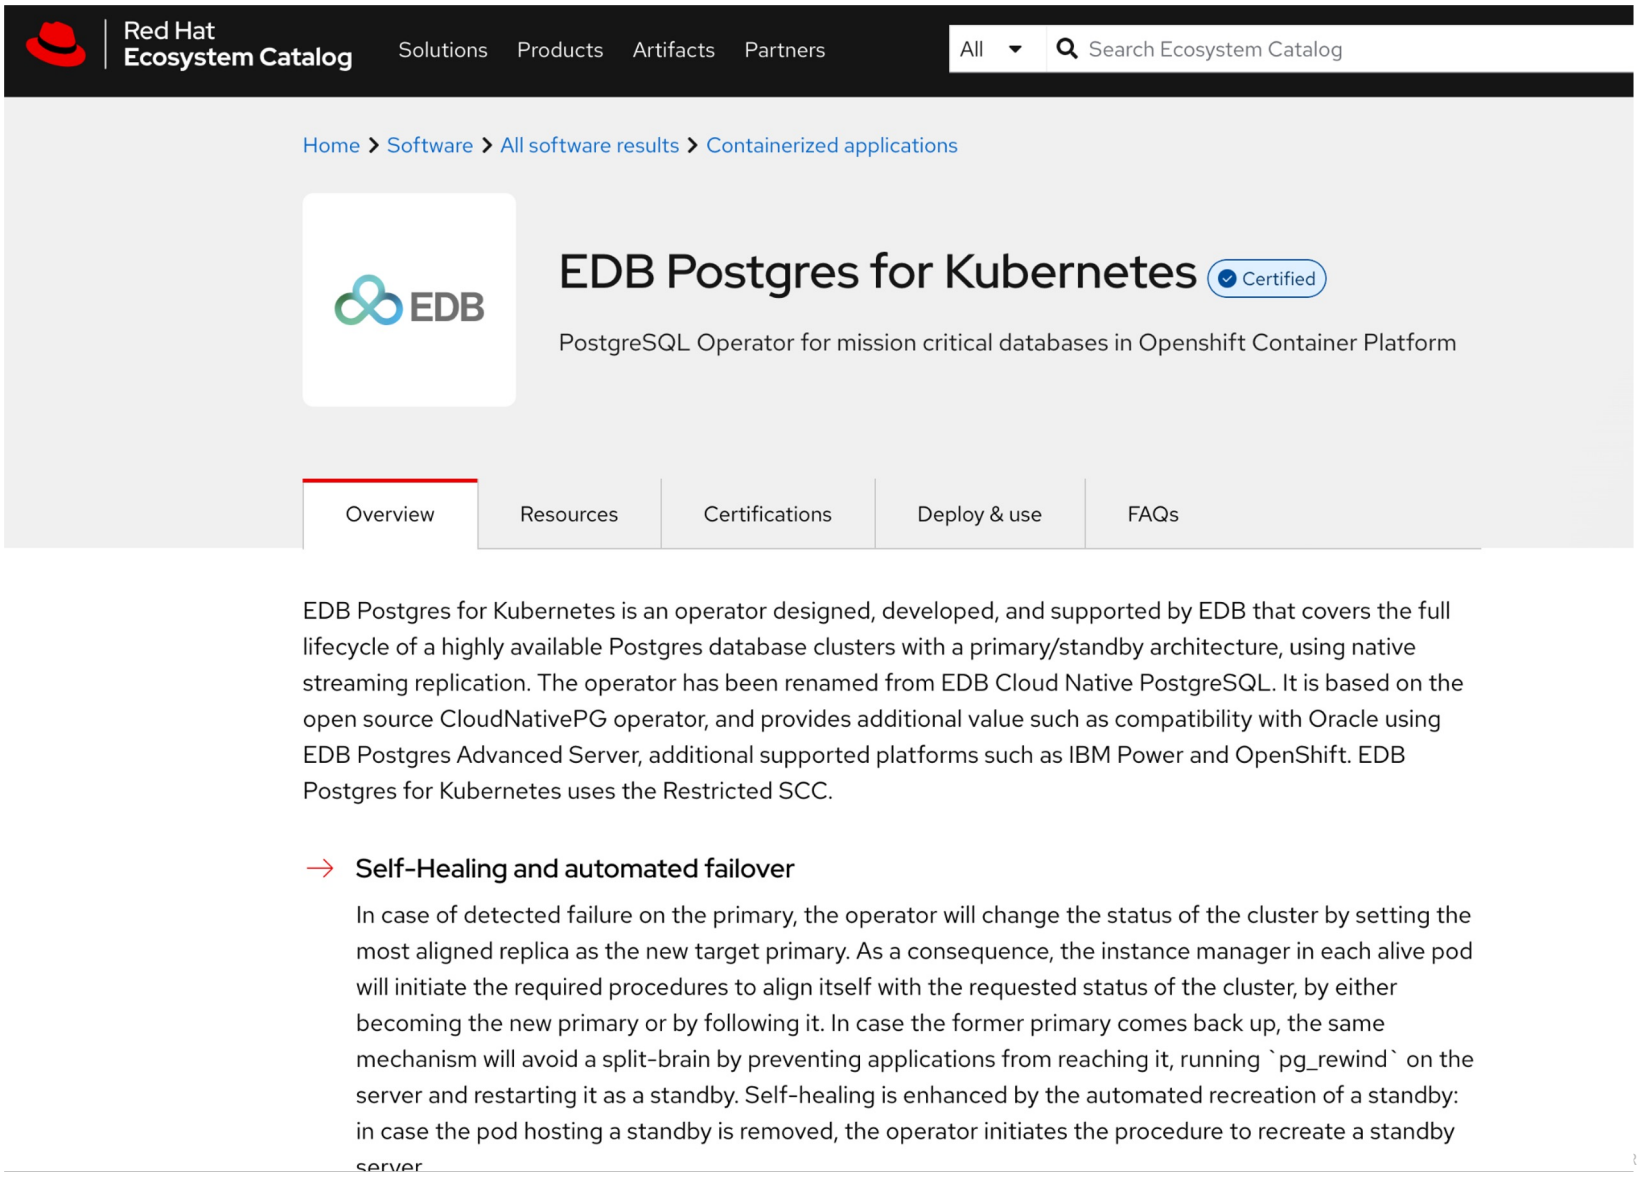
Red Hat
Ecosystem Catalog
Solutions
Products
Artifacts
Partners
All
Search Ecosystem Catalog
Home > Software > All software results > Containerized applications
EDB
EDB Postgres for Kubernetes
Certified
PostgreSQL Operator for mission critical databases in Openshift Container Platform
Overview
Resources
Certifications
Deploy & use
FAQs
EDB Postgres for Kubernetes is an operator designed, developed, and supported by EDB that covers the full lifecycle of a highly available Postgres database clusters with a primary/standby architecture, using native streaming replication. The operator has been renamed from EDB Cloud Native PostgreSQL. It is based on the open source CloudNativePG operator, and provides additional value such as compatibility with Oracle using EDB Postgres Advanced Server, additional supported platforms such as IBM Power and OpenShift. EDB Postgres for Kubernetes uses the Restricted SCC.
Self-Healing and automated failover
In case of detected failure on the primary, the operator will change the status of the cluster by setting the most aligned replica as the new target primary. As a consequence, the instance manager in each alive pod will initiate the required procedures to align itself with the requested status of the cluster, by either becoming the new primary or by following it. In case the former primary comes back up, the same mechanism will avoid a split-brain by preventing applications from reaching it, running `pg_rewind` on the server and restarting it as a standby. Self-healing is enhanced by the automated recreation of a standby: in case the pod hosting a standby is removed, the operator initiates the procedure to recreate a standby server.
RIGHTS RESERVED.

The PostgreSQL `Cluster` resource

CloudNativePG

```
apiVersion: postgresql.cnpg.io/v1
kind: Cluster
metadata:
  name: clapton
spec:
  instances: 3
  postgresql:
    synchronous:
      method: any
      number: 1
  storage:
    size: 40Gi
  walStorage:
    size: 10Gi
```

EDB Postgres for Kubernetes

```
apiVersion: postgresql.k8s.enterprisedb.io/v1
kind: Cluster
metadata:
  name: clapton
spec:
  instances: 3
  postgresql:
    synchronous:
      method: any
      number: 1
  storage:
    size: 40Gi
  walStorage:
    size: 10Gi
```

USE CASE: SOVEREIGN AI

PRODUCTION-READY AI WITH FULL CONTROL OVER DATA

Challenges

Regional AI governance requirements (e.g. EU AI Act)
Cloud sovereignty concerns
Global GPU access

EDB Postgres AI Solution»

Models and data are entirely contained in the system
Custom models for local languages, norms, and ethics
No public cloud dependencies
Pre-configured with Supermicro and NVIDIA GPUs

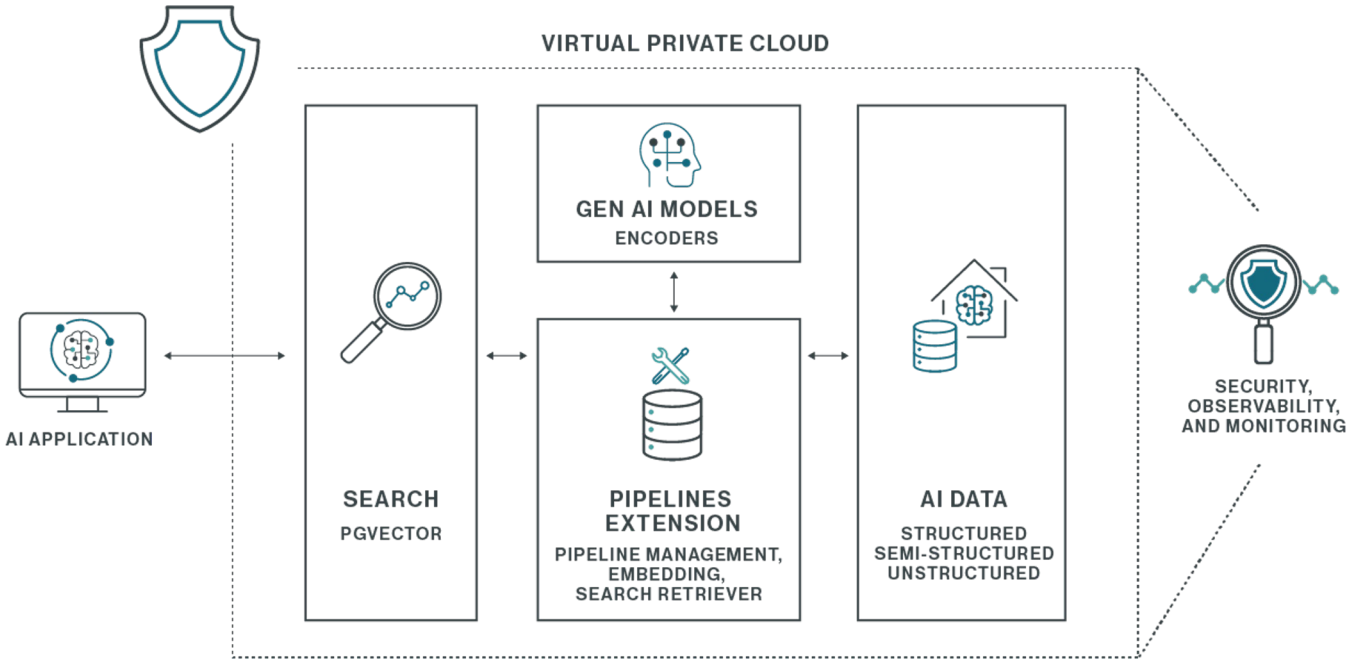

EDB POSTGRES AI

Intelligent Optimization

REDUCE COSTS AND CARBON FOOTPRINT, WHILE BOOSTING PERFORMANCE

Challenges

- Expertise required for tuning
- High costs to meet demand (the “scale by credit card” dilemma)
- ESG tracking and goals

EDB Postgres AI Solution»

- **Optimize in minutes**, without expert intervention
- **8x faster app performance** with intelligent recommendations
- **Reduce costs** to scale
- **Measure and meet ESG goals**,

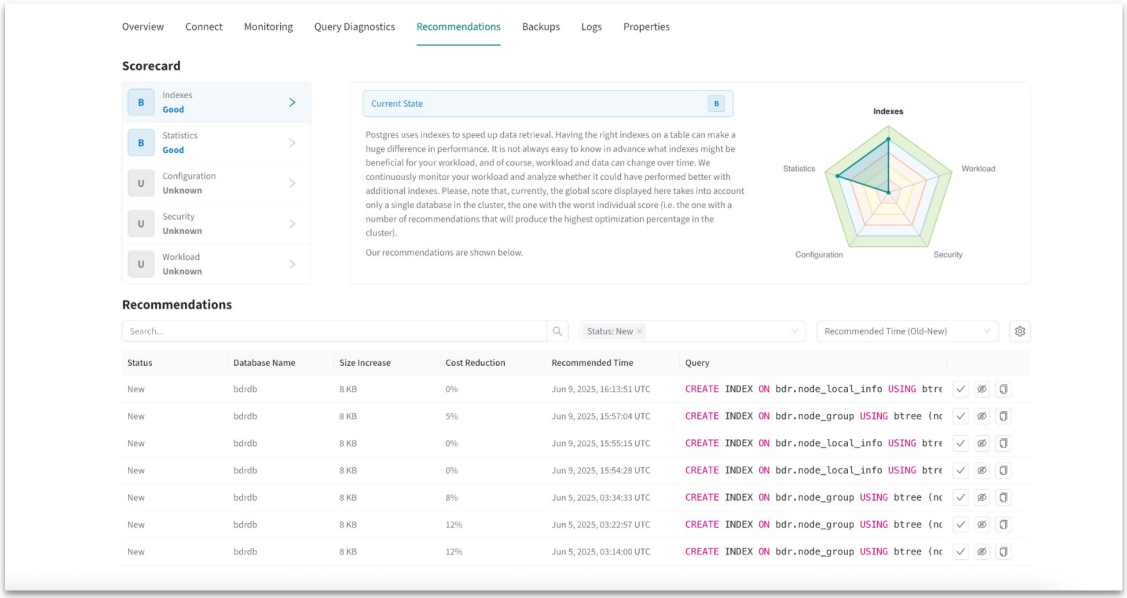

LEARN MORE >>

See how to view and apply recommendations from EDB Postgres AI to reduce costs and carbon footprint while boosting performance.

Service High Availability

ALWAYS-ON DATA PLATFORM FOR TIER 1 AND GLOBAL APPLICATIONS

Challenges

- Query latency
- Customer drop-off
- Complex HA configuration
- Business continuity

EDB Postgres AI Solution»

- **Simple HA:** Provision 4 or 5 9's HA clusters in a few clicks
- **Online upgrades**, no downtime
- **Query diagnostics:** Resolve bottlenecks 5x faster
- **Intelligent recommendations:** Up to 8x faster app

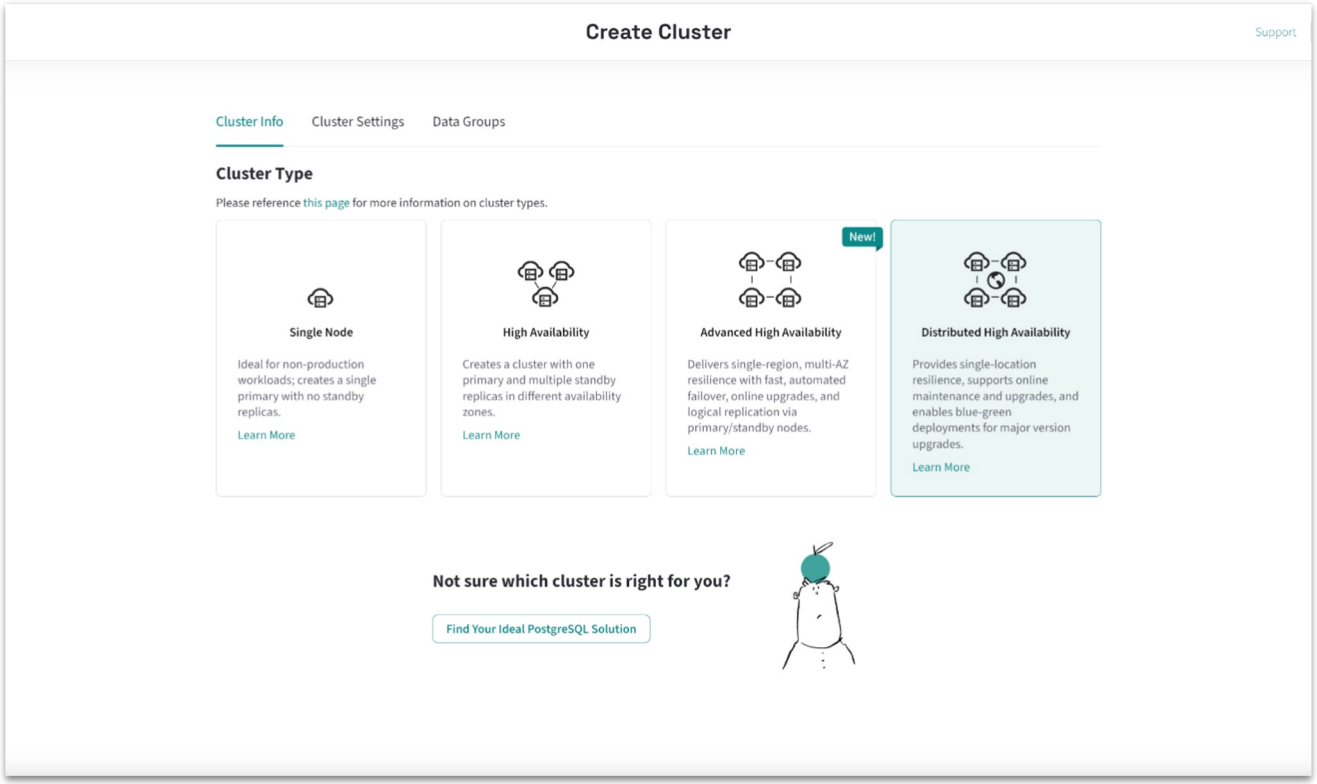

WATCH DEMO >>

See how to deploy a highly available EDB Postgres AI cluster from a template in just a few clicks.

Unified Query Engine

TRANSFORM POSTGRES INTO A TRUSTED, SOVEREIGN ENGINE FOR SEARCH — IN POSTGRES, ACROSS THE LAKEHOUSE, AND ANYWHERE WITH AGENTS

Challenges

- Too many query engines
- Governance gaps
- Legacy silos and modern data stack fragmentation

EDB Postgres AI Solution»

- **One platform** for querying across all data
- **Familiar Postgres interface**, no new skill sets required
- **Real-time data processing:** improve efficiency, flexibility, scalability, and security

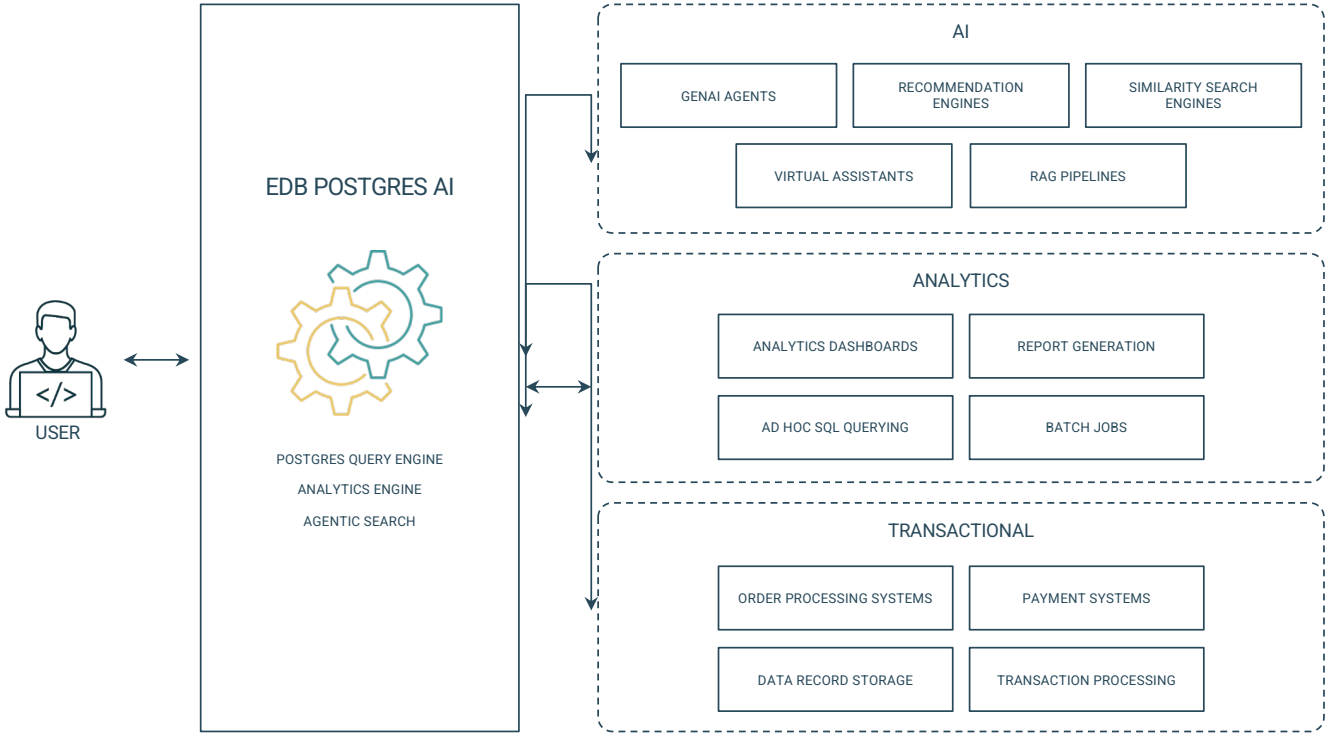

Analytics and AI Ecosystem Integration

TURN FAMILIAR, TRUSTED POSTGRES INTO A FRONTEND FOR LAKEHOUSE AND AI ECOSYSTEMS

Challenges

Data silos across multiple ecosystems slow time to insights, spiral costs

Operational, security, and governance issues

EDB Postgres AI Solution»

Simplified management across the data and AI landscapes

Accelerated time to insights with centralized querying of data in Postgres and lakehouses

Speed and cost-efficiency with automated data tiering for up to

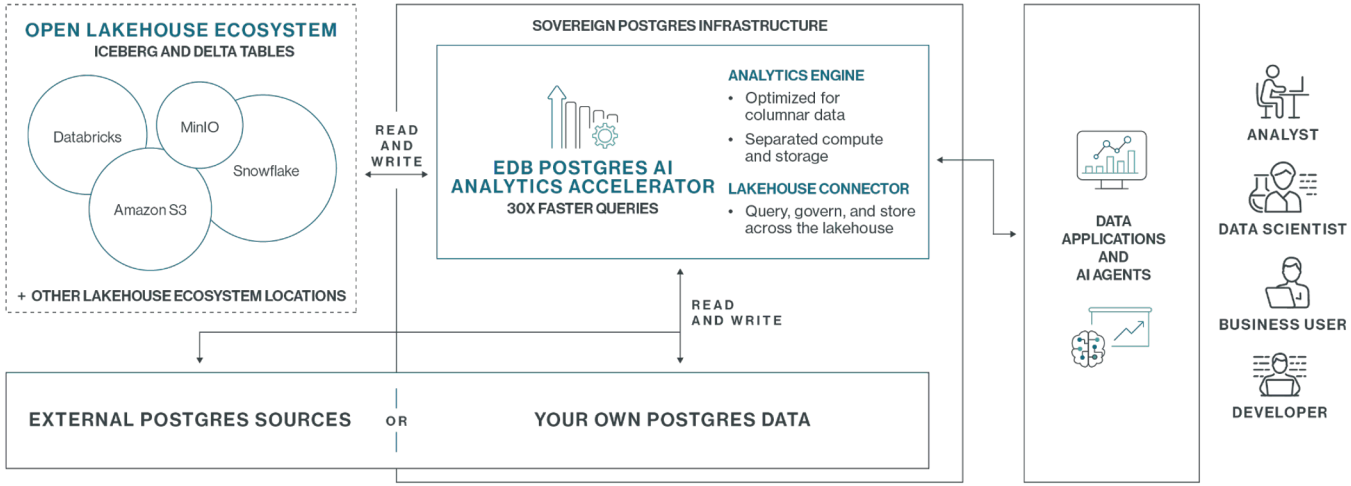

USE CASE: LEGACY APP MODERNIZATION

BREAK FREE FROM LEGACY CONSTRAINTS TO SUPPORT NEXT GEN INNOVATION

Challenges

- Legacy constraints
- Multi-model data and open standards
- Compliance issues

EDB Postgres AI Solution »

Outcomes

- Next gen unified platform
- Faster time to market
- Compliant open source

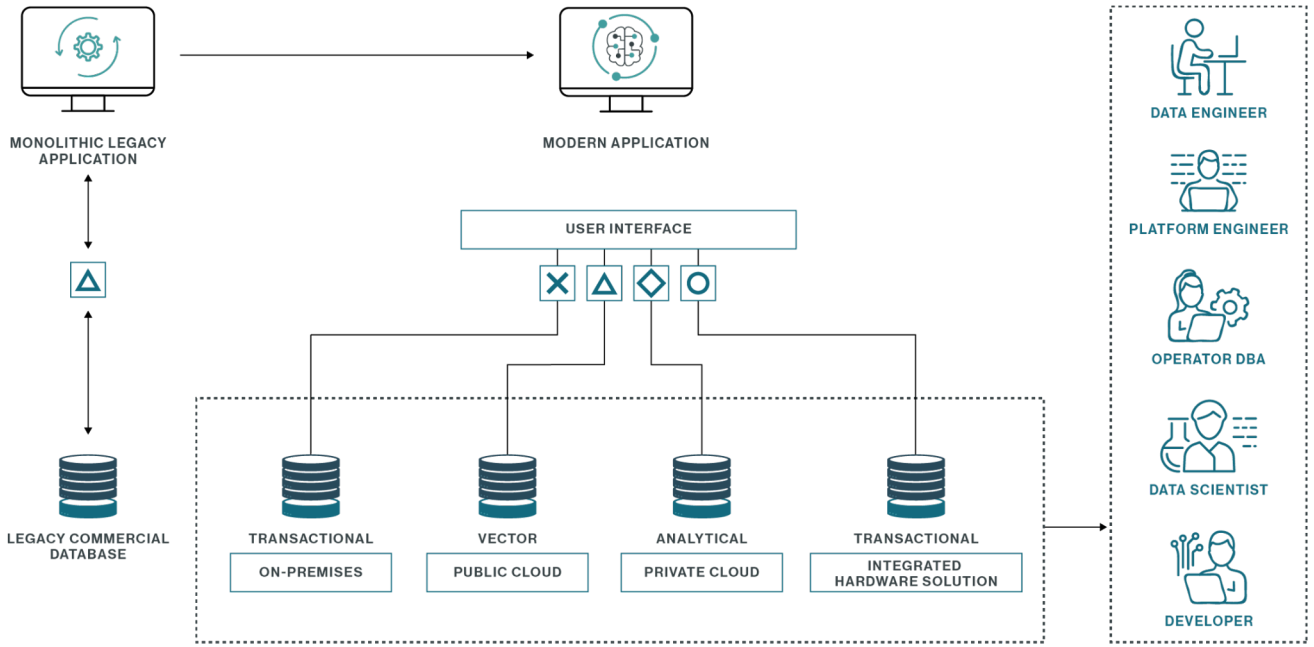

EDB POSTGRES AI

USE CASE: OMNI-DATA PLATFORM

PRODUCTION-READY AI WITH FULL CONTROL OVER DATA

Challenges

- Data silos
- Complex, sprawling infrastructure
- High-cost management

EDB Postgres AI Solution »

Outcomes

- Multi-model support for next gen applications
- 30x faster insights
- Empowered decision making

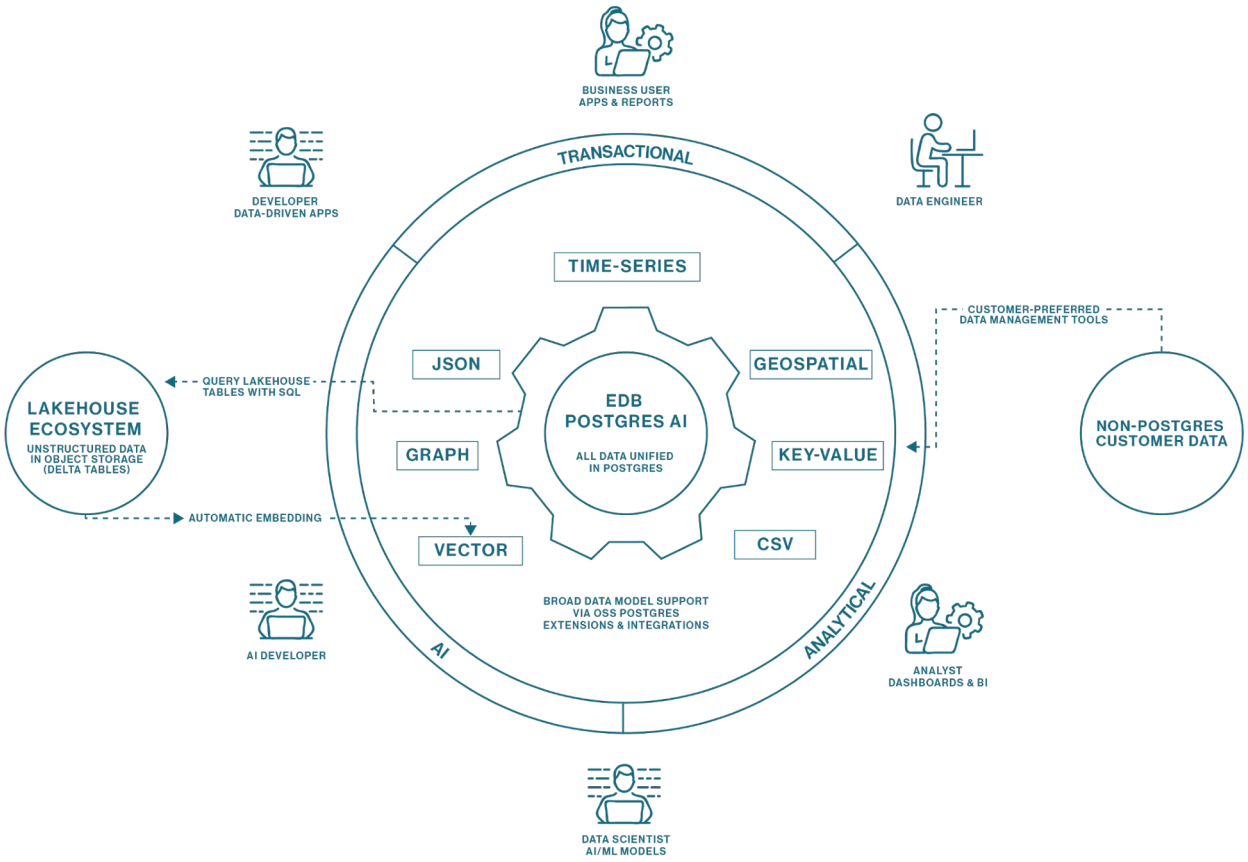

EDB POSTGRES AI

USE CASE: HYBRID DATABASE-AS-A-SERVICE (DBaaS)

CLOUD AGILITY IN A SECURE, SELF-HOSTED MODEL

Challenges

- Deployment flexibility
- Developer agility
- Complex management

EDB Postgres AI Solution »

Outcomes

- Hybrid deployment
- Faster app development
- Operational efficiency and TCO reduction

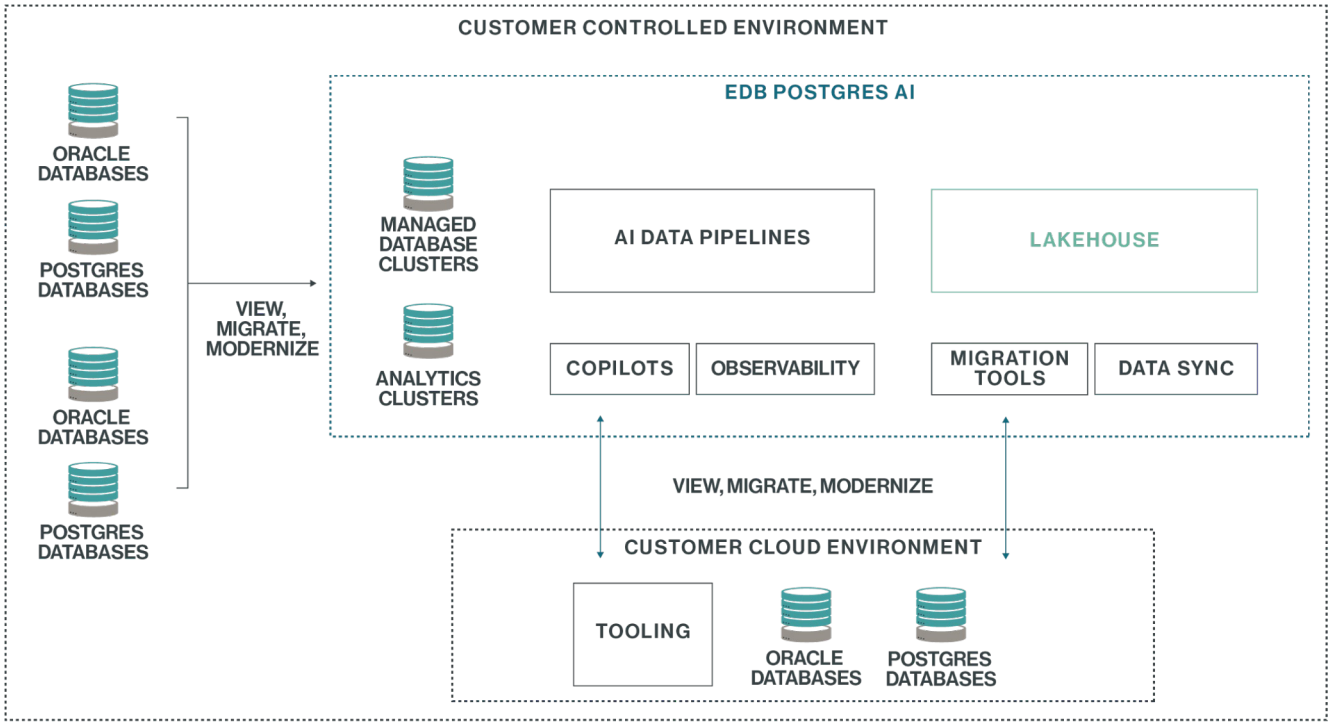

EDB POSTGRES AI

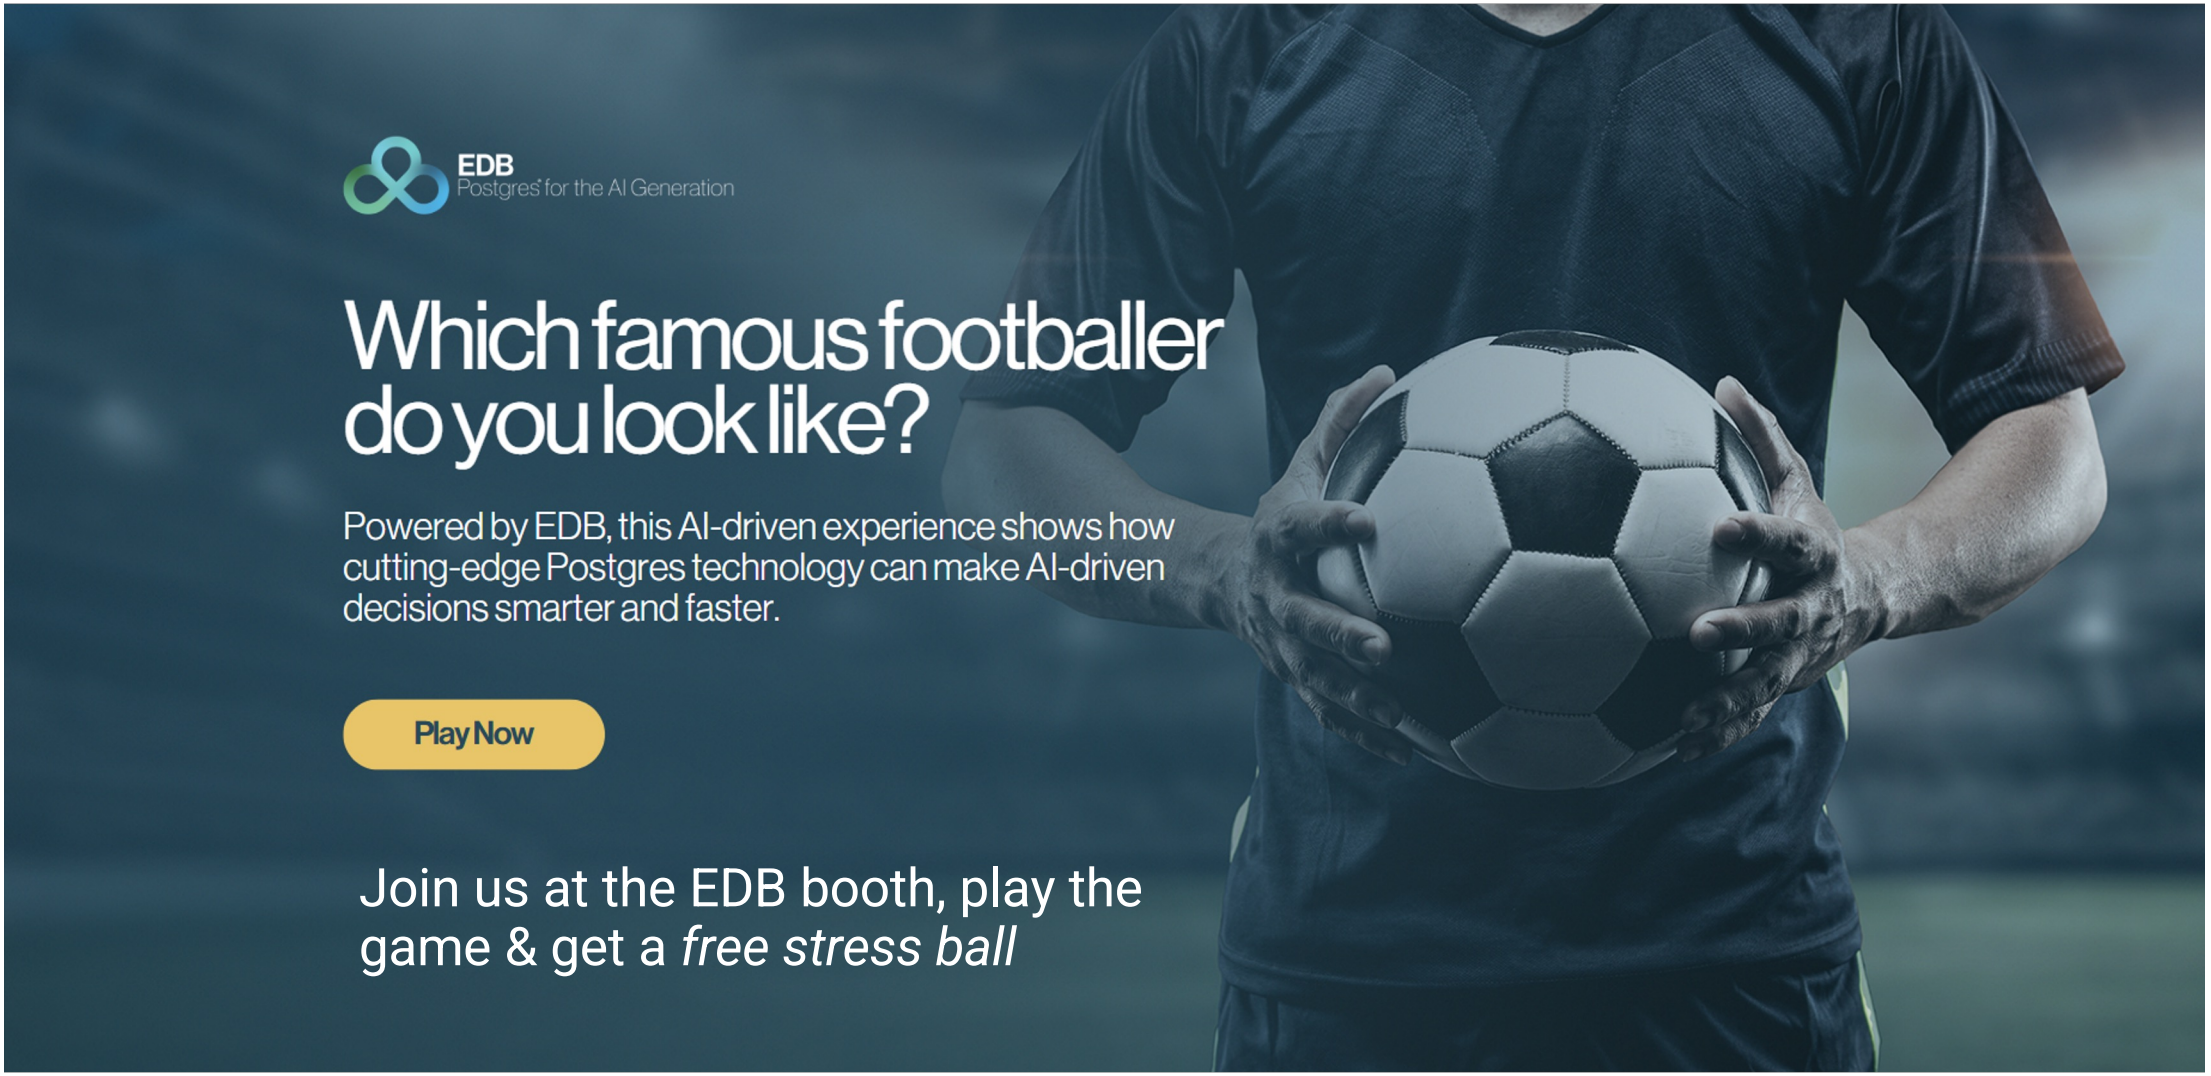

EDB POSTGRES AI

Red Hat Summit

Connect

Grazie

 linkedin.com/company/red-hat

 facebook.com/redhatinc

 youtube.com/user/RedHatVideos

 twitter.com/RedHat