CUFEP 2025

19th June 2025

Introduction


SEDONA
cetelem
BANQUE DE FRANCE
EUROSYSTÈME
LA BANQUE POSTALE
DIGORA
by deep
GRDF
stet
TRANSACTIS
MICHELIN
HSBC
axians
COVEA
groupe JVS
ÉCLAIREURS NUMÉRIQUES
MIRAKL
LA POSTE
PRÉFECTURE DE POLICE
Liberté
Égalité
Fraternité
LOXAM
AIRBUS
FICHORGA
17 96
cegedim
Smals
ICT for society
LOMBARD ODIER
LOMBARD ODIER DARIER HENTSCH
TotalEnergies
AFP
Institut national de la stat
et des études économiqu
Insee
Mesurer pour comprendre
AZQORE
CEGELEM

Meet the EDB Team

Hervé Timsit	Chief Revenue Officer
Cédric Amargier	ISV Account Executive
Cyril Bertrand	Senior Partner Account Executive
Cyrille Sauvain	Strategic Account Executive
Eric Pillon	Strategic Account Executive
Eric Labaune	Senior Customer Success Manager
Gilles Maghami	Senior Technical Account Manager
Sergio Romera	Senior Manager, Sales Engineering
Lucie Zeng	Associate Sales Engineer
Raphaël Chir	Senior Sales Engineer
Sébastien Sire	Senior Consultant, Professional Services
Sophie Palmer	Marketing Manager

Agenda

09:30 – 10:00	Accueil café et petit-déjeuner ☕
10:00 – 10:15	Introduction
	Cyrille Sauvain, Strategic Account Executive, EDB
	Eric Pillon, Strategic Account Executive, EDB
10:15 – 10:30	Le rôle de OSS dans l'innovation
	Hervé Timsit, CRO, EDB
10:30 – 11:15	Session 1 : EDB Postgres AI & Stratégie
	Sergio Romera, Senior Manager, Sales Engineering, EDB
	Lucie Zeng, Associate Sales Engineer, EDB
11:15 – 11:45	Témoignage client : Postgres chez Transactis
	Moez Essaidi, Solutions & IT Architect, Ph.D.
11:45 – 12:05	Pause café ☕
12:05 – 12:50	Session 2 : EDB Postgres AI & Stratégie
	Sergio Romera, Senior Manager, Sales Engineering, EDB
12:50–14:20	Discussion en table ronde & Lunch
14.20 - 14.35	Session de feedback & tirage au sort
14.40 - 16.30	Visite guidée du musée d'Orsay

EDB POSTGRES AI

CUFEP

Herve Timsit
CRO

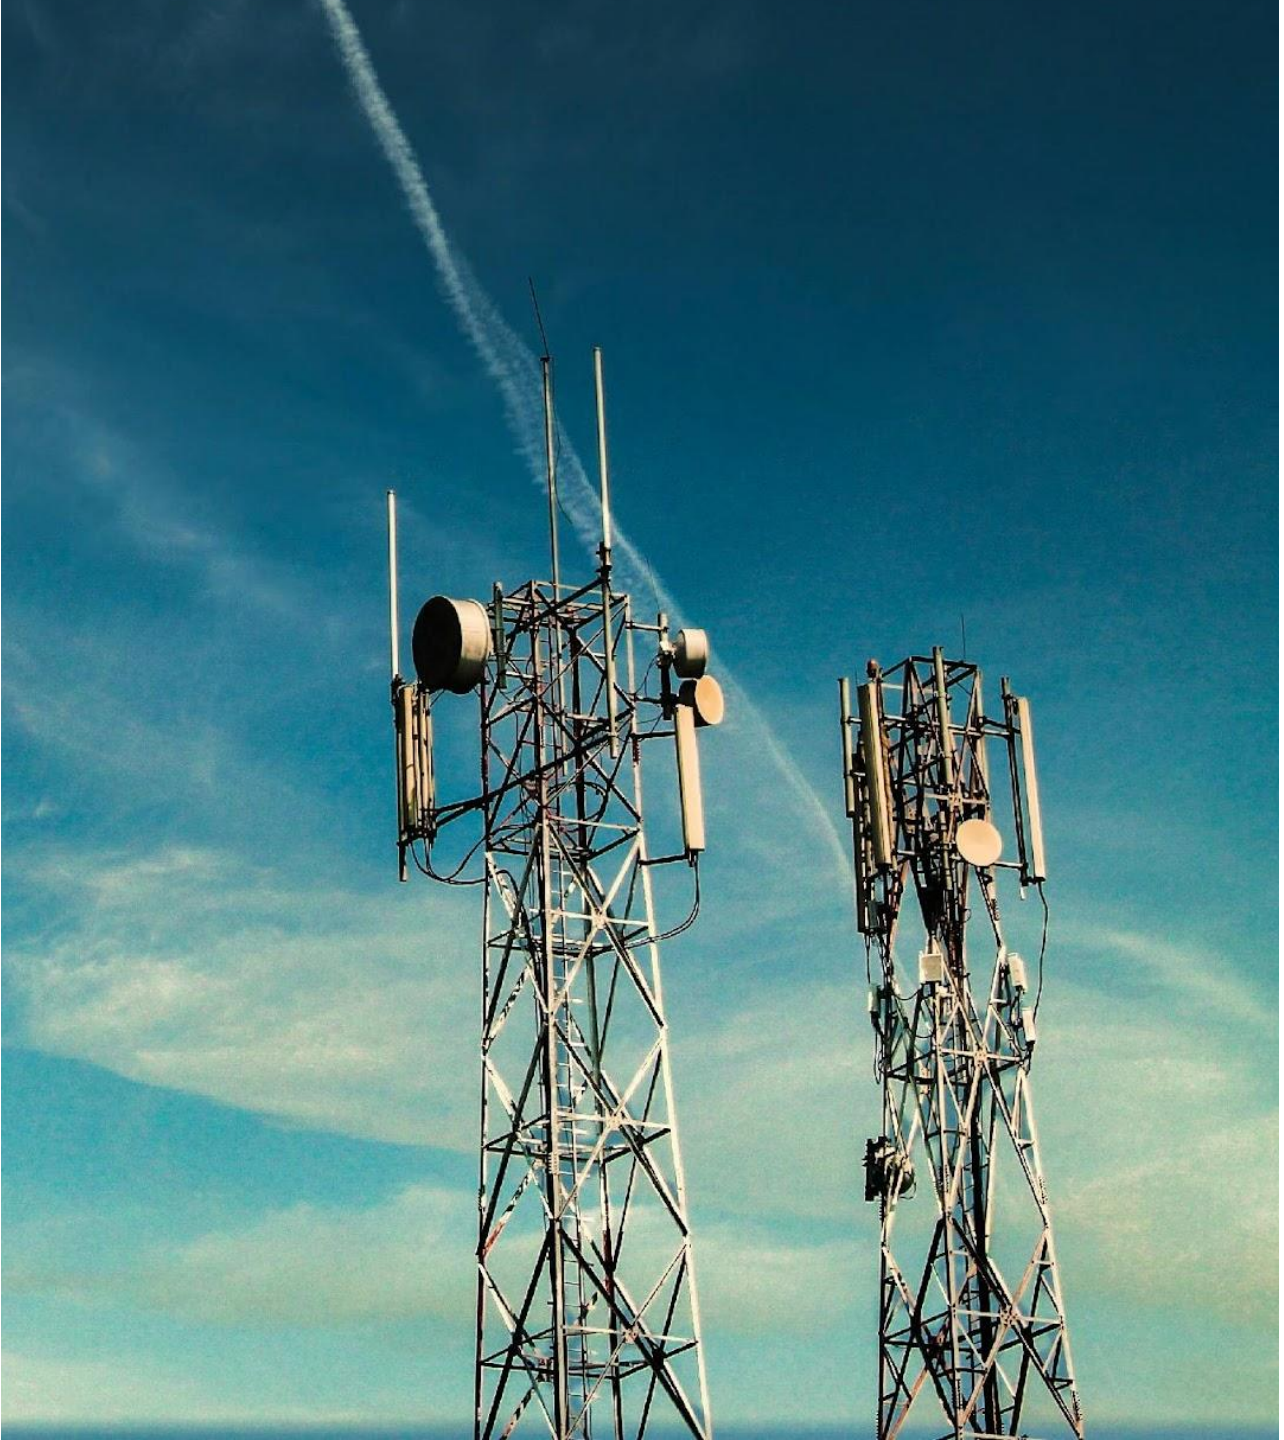

EDB POSTGRES AI

Une brève histoire des projets OSS

Movable Type (2013)

SugarCRM (2014)

Redis (2018)

MongoDB (2018)

Confluent (2018)

Cockroach Labs (2019)

Sentry (2019)

Elastic (2021)

HashiCorp (2023)

Snowplow (2024)

Greenplum (2024)

EDB POSTGRES AI

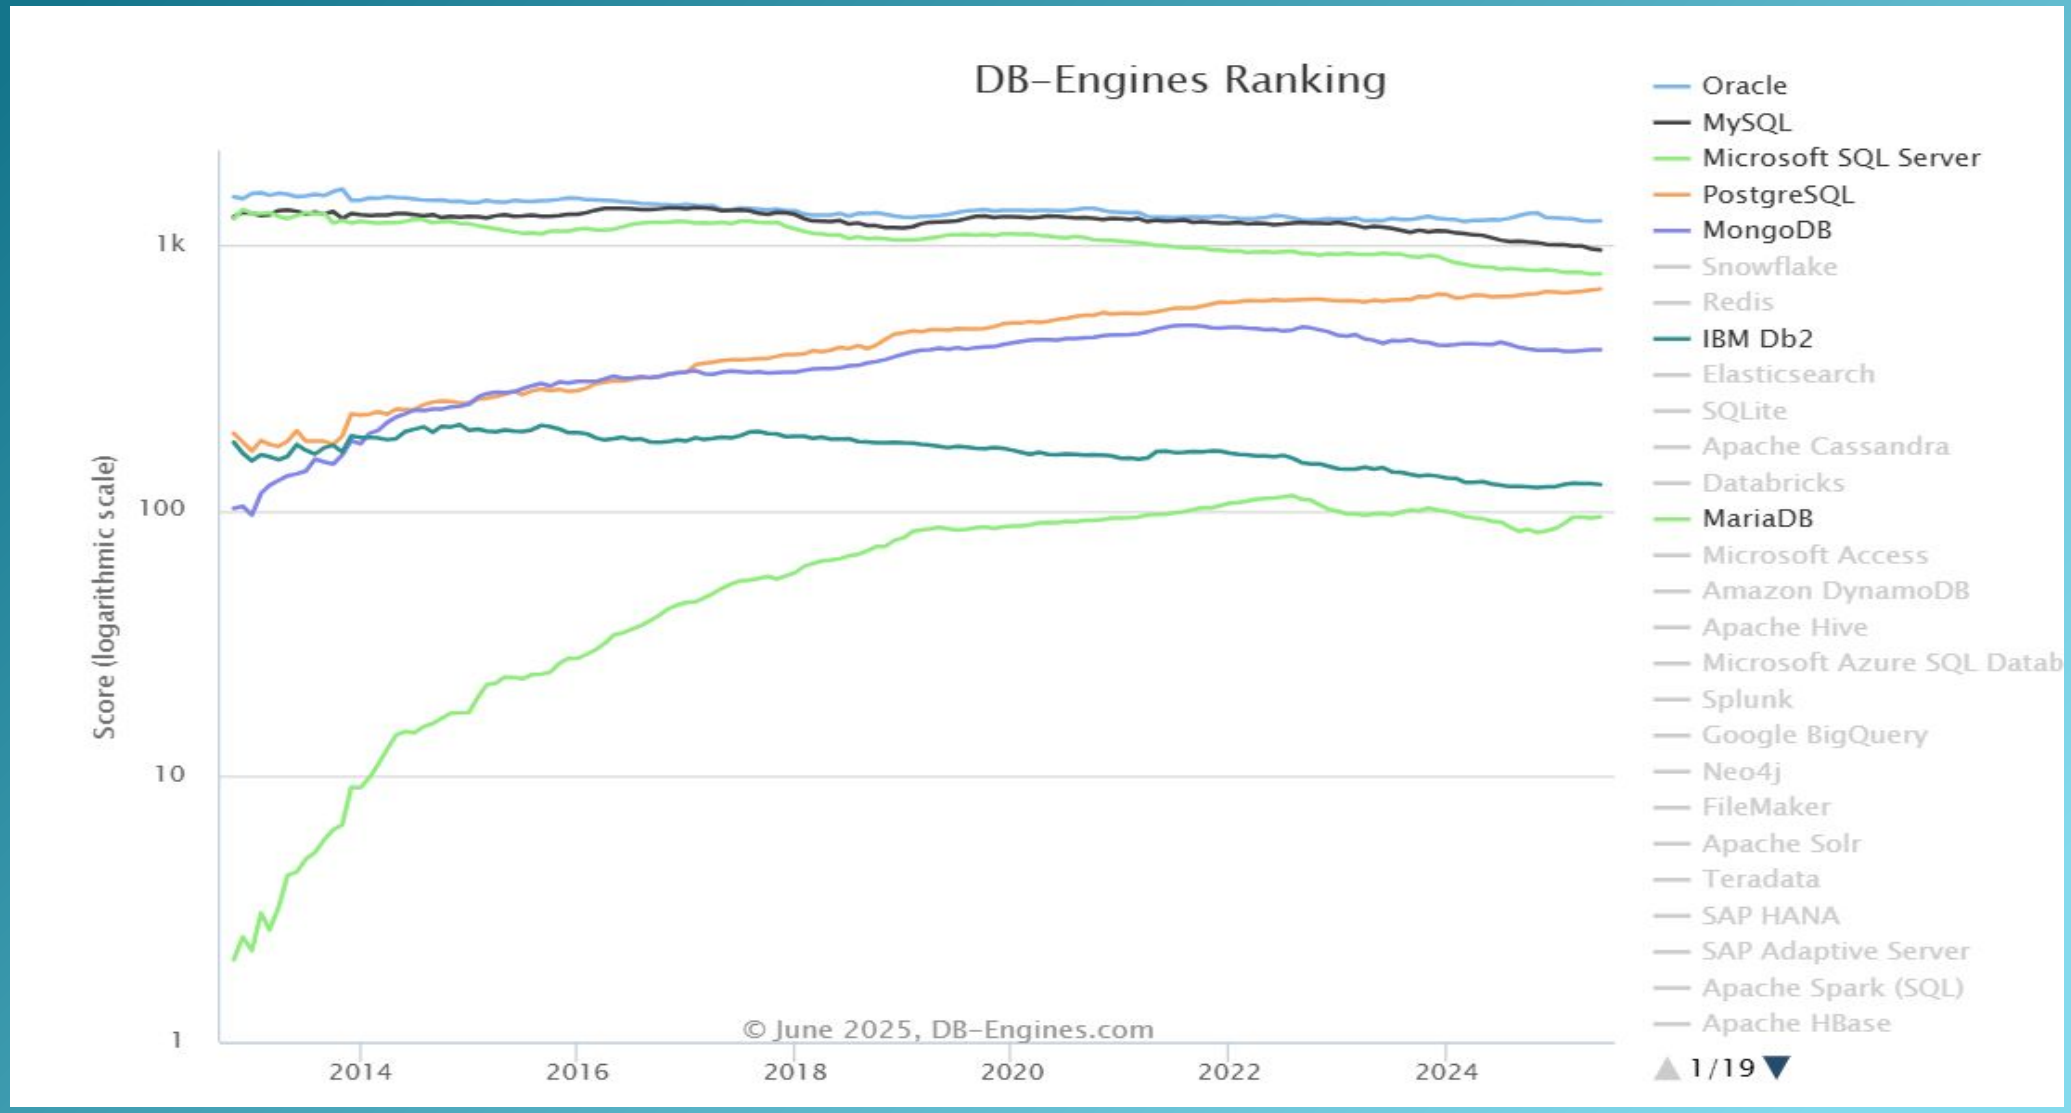

L’urgence d ‘innover sans s’enfermer

De changer

- Operating model
  (Cloud, Edge, On Prem)
- De Cloud
- De Hardware
- De stack (Bare Metal, VM, K8s)

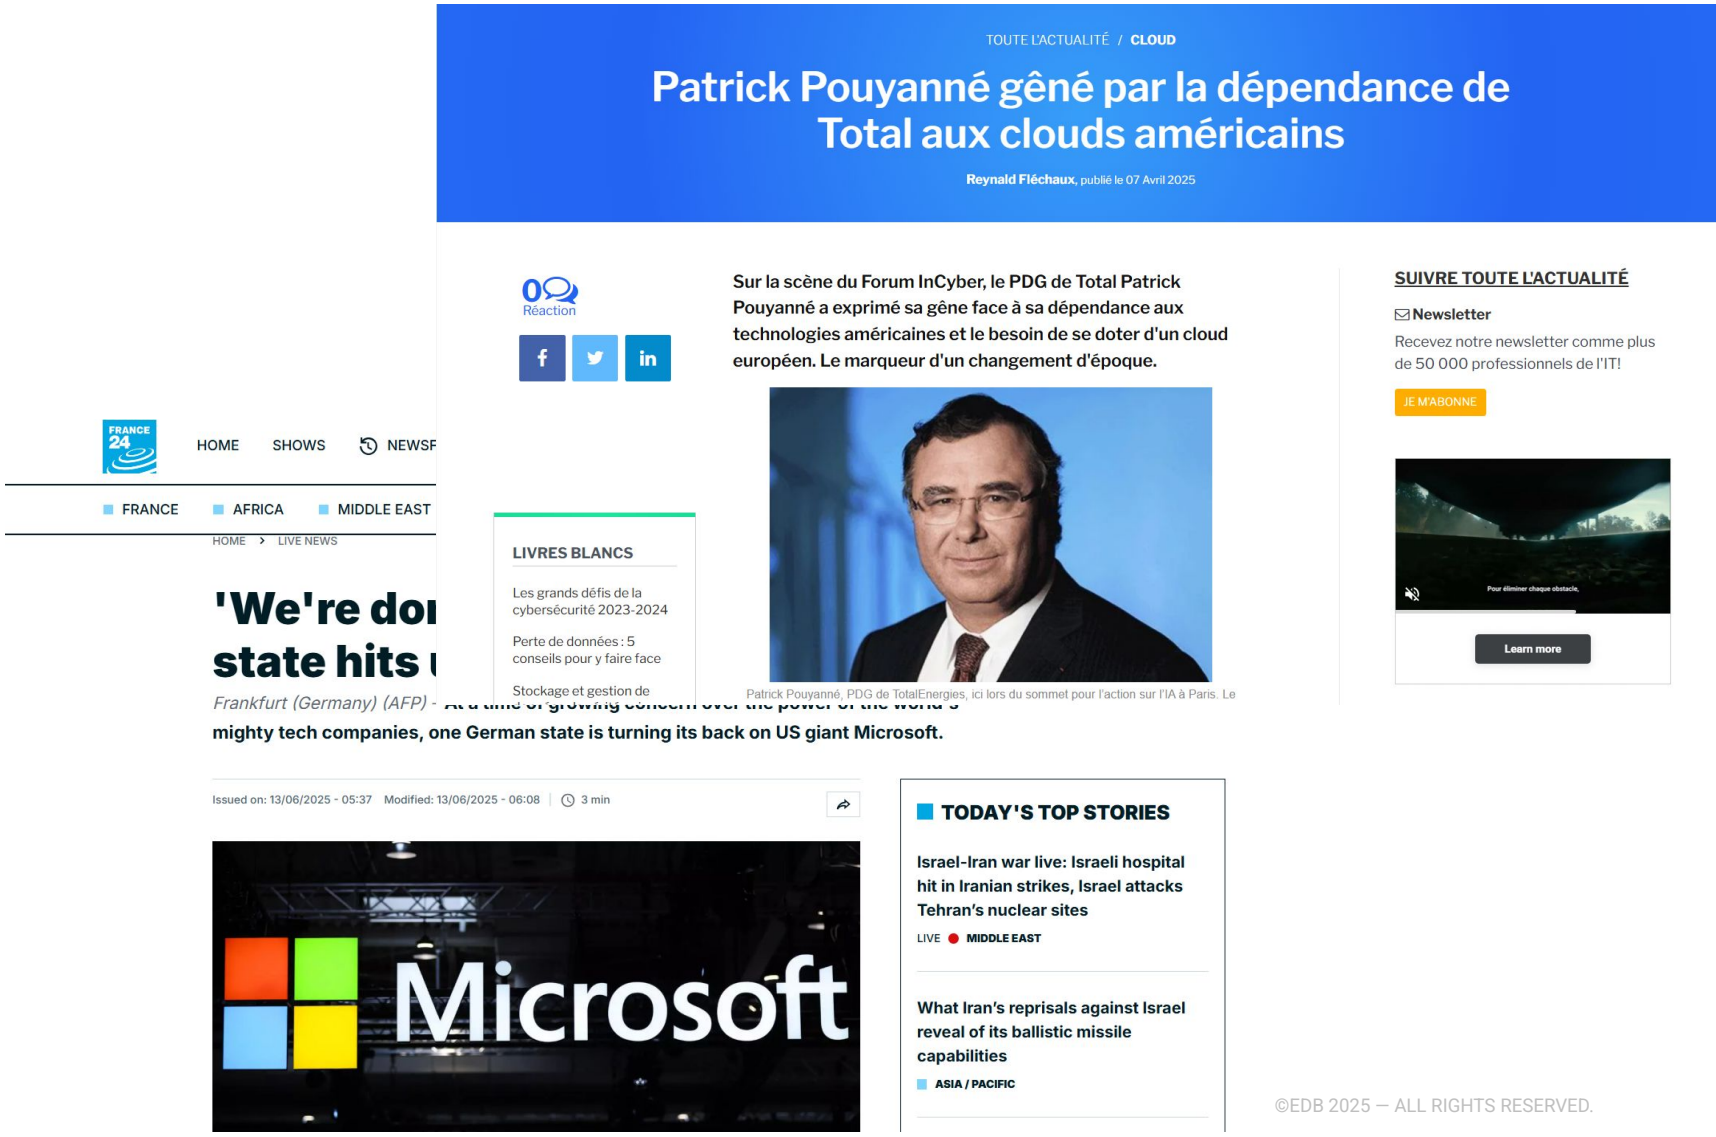

EDB POSTGRES AI

L’auto-intoxication

Hervé Timsit • You
CRO at EDB
2w •

In the realm of AI, the reliability of outputs hinges on the quality of data ingestion. For industries bound by regulations, the concept of sovereign data and AI emerges as the most logical path forward. Ensuring data sovereignty is paramount for robust AI outcomes in such sectors.

Source: [Link to the article](https://lnkd.in/ezzi-9AF)

#sovereignAI #sovereigncloud #EDB #EDBPostrgesAI #Postgres

AI model collapse is not what we paid for
theregister.com

EDB POSTGRES AI

Il n y a pas d’innovation sans la liberté de créer

Tirez davantage de valeur de vos données grâce à une plateforme Omni-Data prête à l'emploi pour les charges de travail transactionnelles, analytiques et d'IA, avec une haute disponibilité, une évolutivité et une conformité prêtes à l'emploi.

EDB POSTGRES AI

CUFEP 2025

Part 1: EDB Postgres AI & Strategy
From Transactional to AI, Through Analytics at Scale

Sergio Romera, **Senior Manager, Sales Engineering South Team**
Lucie Zeng, **Associate Sales Engineer**

Agenda

Sovereign AI

WarehousePG

Analytics

EDB POSTGRES AI

Concepts

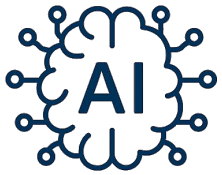

AI

Systems designed to simulate human-like intelligence, performing tasks that typically require human cognition. This **includes everything from simple decision rules to complex learning algorithms to understanding language and perception**. This is the umbrella term that includes all of the following.

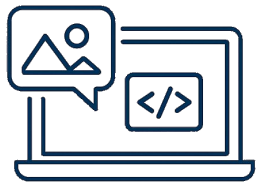

GenAI

A powerful subset of AI focused on **creating new content**, by learning patterns from existing data. Instead of just analyzing or classifying existing data, GenAI models can generate text, images, audio, video, or code that is original and often highly realistic.

Agentic AI

AI systems that **can autonomously set goals, plan actions, execute those plans, and adapt based on the results and environment**. These systems aren't just reactionary; they're proactively working towards objectives, potentially breaking down complex tasks into smaller steps without constant human intervention. **Think of them as AI workers that can drive their own tasks.**

Sovereign AI

Though the idea of "sovereignty" has typically referred specifically to nations, the concept has evolved to apply to any entity with an owned domain. Sovereign AI is when **an organization has complete control and ownership over its AI infrastructure, data, models, and capabilities** — serving needs like data residency, security, ethical alignment, and strategic independence.

EDB POSTGRES AI

EDB Postgres® AI

Database

- ENTERPRISE POSTGRES SERVER
- ORACLE COMPATIBLE SERVER
- COMMUNITY POSTGRESQL SERVER
- MULTI-MODEL EXTENSIONS
- MANAGEMENT & OBSERVABILITY
- KUBERNETES OPERATORS
- SUPPLY CHAIN SECURITY
- MIGRATION TOOLS
- HIGH AVAILABILITY

Analytics Accelerator

- ANALYTICS ENGINE
- LAKEHOUSE CONNECTOR
- SUPPORT FOR GREENPLUM WORKLOADS

AI Factory

- VECTOR ENGINE
- AI PIPELINE
- GENAI BUILDER
- AGENT STUDIO
- MODEL SERVING

Hybrid Management

- HYBRID OBSERVABILITY
- HYBRID DBAAS
- DISTRIBUTED HA (99.999%)
- MIGRATIONS

Deploy Anywhere

- SOVEREIGN DATA AND AI FACTORY
- HYBRID SOFTWARE
- MANAGED PLATFORM
- INTEGRATION PARTNERS: AWS, GCP, AZURE, RED HAT OPENSHIFT, IBM, SUPERMICRO, NVIDIA, GRIPTAPE

Sovereign Assurance

- MANAGED SERVICES
- PILOT TO PRODUCTION SLAs
- OUTCOME EXECUTION

EDB POSTGRES AI

Concepts

Data warehouse

A data warehouse is a centralized repository that stores integrated data from one or more disparate sources.

It is designed for reporting and data analysis and typically holds historical data structured for fast query performance, often following a star or snowflake schema.

Data warehouses are optimized for read-heavy analytical queries rather than transactional processing.

Data Lake

A data lake is a centralized repository that allows you to store all your structured and unstructured data at any scale.

Unlike a data warehouse, it stores raw data in its native format, without a predefined schema.

This flexibility enables various analytical approaches, including big data processing, machine learning, and real-time analytics.

Data Lakehouse

A data lakehouse is a new data architecture that combines the benefits of data lakes (flexibility, low cost, ability to handle raw data) with the benefits of data warehouses (structured transactions, schema enforcement, robust query performance).

It typically uses open file formats and provides ACID (Atomicity, Consistency, Isolation, and Durability) transactions, data governance, and strong performance for both analytics and machine learning workloads, often built on top of a data lake.

Columnar Data

A columnar data format stores data table-wise by columns rather than by rows. This means all values for a single column are stored contiguously.

This format is highly efficient for analytical queries that often access only a subset of columns, as it minimizes disk I/O and improves compression rates, leading to faster query execution in analytical databases.

Some of the popular Columnar Databases are Amazon Redshift, MariaDB, Snowflake Data Cloud, Microsoft Azure Cosmos DB, and many more.

Concepts

DataFusion is an **extensible query engine** written in Rust that uses Apache Arrow as its in-memory format.

Delta Lake is an **open-source table format** and transaction layer that enhances modern data lakes with ACID guarantees, schema enforcement, time travel, and scalable performance.

Delta Lake adds database-like reliability and consistency to object storage systems such as S3, GCS, and Azure Data Lake Storage.

Open format built on **Apache Parquet** with a transaction log

Apache Iceberg is a high-performance, **open table format** designed for data lakes built on object storage. It provides scalable metadata management, schema evolution, and ACID transactions.

Apache Parquet is an open source, **column-oriented data file format** designed for efficient data storage and retrieval. It provides **high performance compression** and encoding schemes to handle complex data in bulk and is supported in many programming language and analytics tools.

Concepts

Data warehouse

Data Lake

Data Lakehouse

using

Columnar Data

EDB POSTGRES AI

Concepts

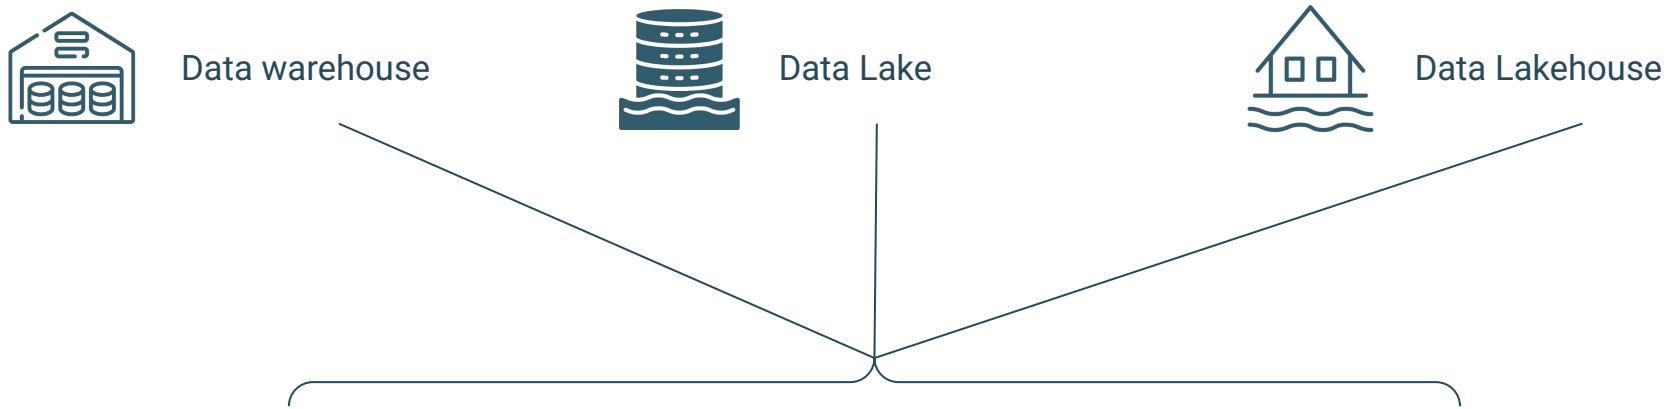

Columnar Data

A columnar data format stores data table-wise by columns rather than by rows. This means all values for a single column are stored contiguously.

This format is highly efficient for analytical queries that often access only a subset of columns, as it minimizes disk I/O and improves compression rates, leading to faster query execution in analytical databases.

Some of the popular Columnar Databases are Amazon Redshift, MariaDB, Snowflake Data Cloud, Microsoft Azure Cosmos DB, and many more.

Data & AI Challenges

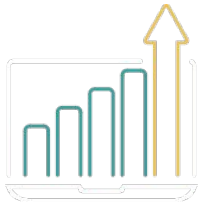

Fast Data Growth

Driven by digital transformation, IoT, business expansion, and increased AI adoption.

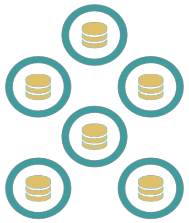

Data Silos & Legacy Monolithic systems

Slow innovation, poor scalability, high maintenance cost, and limited AI integration.

Data Security & compliance

Databases are prime targets—SQL injection, weak encryption, and shadow DBs are major threats.

Complex & Hybrid infrastructure

High operational overhead, costly integrations, and limited observability.

Getting AI into production is hard

Over 80% of IT leaders say it takes 6–24 months to go from AI pilot to production.

EDB POSTGRES AI

Open Source is the way forward for Innovation

open, decentralized and deeply collaborative

Increases speed of innovation

Open development accelerates discovery and progress vs. closed proprietary systems.

Open-source communities like **PostgreSQL** and **Hugging Face** evolve rapidly through global collaboration.

Democratizes access

Open sharing removes barriers to emerging AI technologies.

Frameworks like **TensorFlow**, **PyTorch**, and **LangChain** make advanced AI tools available to all—startups to enterprises.

Trust (Openness and Transparency)

Users see how data is used—supporting privacy and sovereignty.

Open models (e.g., **LLaMA**, **Mistral**) and transparent data pipelines allow full visibility into AI decisions and data usage.

Reduces costs and Avoid vendor lock-in

Flexible, customizable, and cost-effective—no dependency on a single vendor.

Platforms like **Kubernetes**, **Linux**, and **PostgreSQL** enable flexible, cloud-agnostic deployments

EDB POSTGRES AI

Open Source Software market analysis

Insights into the current trends, challenges, priorities shaping OSS adoption among organizations worldwide.

96% Either increased or maintained their use of OSS* in 2024

34% reported a significant increase

39% Is using Kubernetes, grown remarkably since 2021 (18%) more than double)

51% Is using PostgreSQL database

29% Is challenged with Installation, upgrades, and configuration issues

34% Is challenged with keeping up with updates and patches

*OSS: Open Source Software

Source*: 2025 State of Open Source Report by OpenLogic
OSS Usage, Market Trends, and Analysis

EDB POSTGRES AI

Postgres usage in Kubernetes

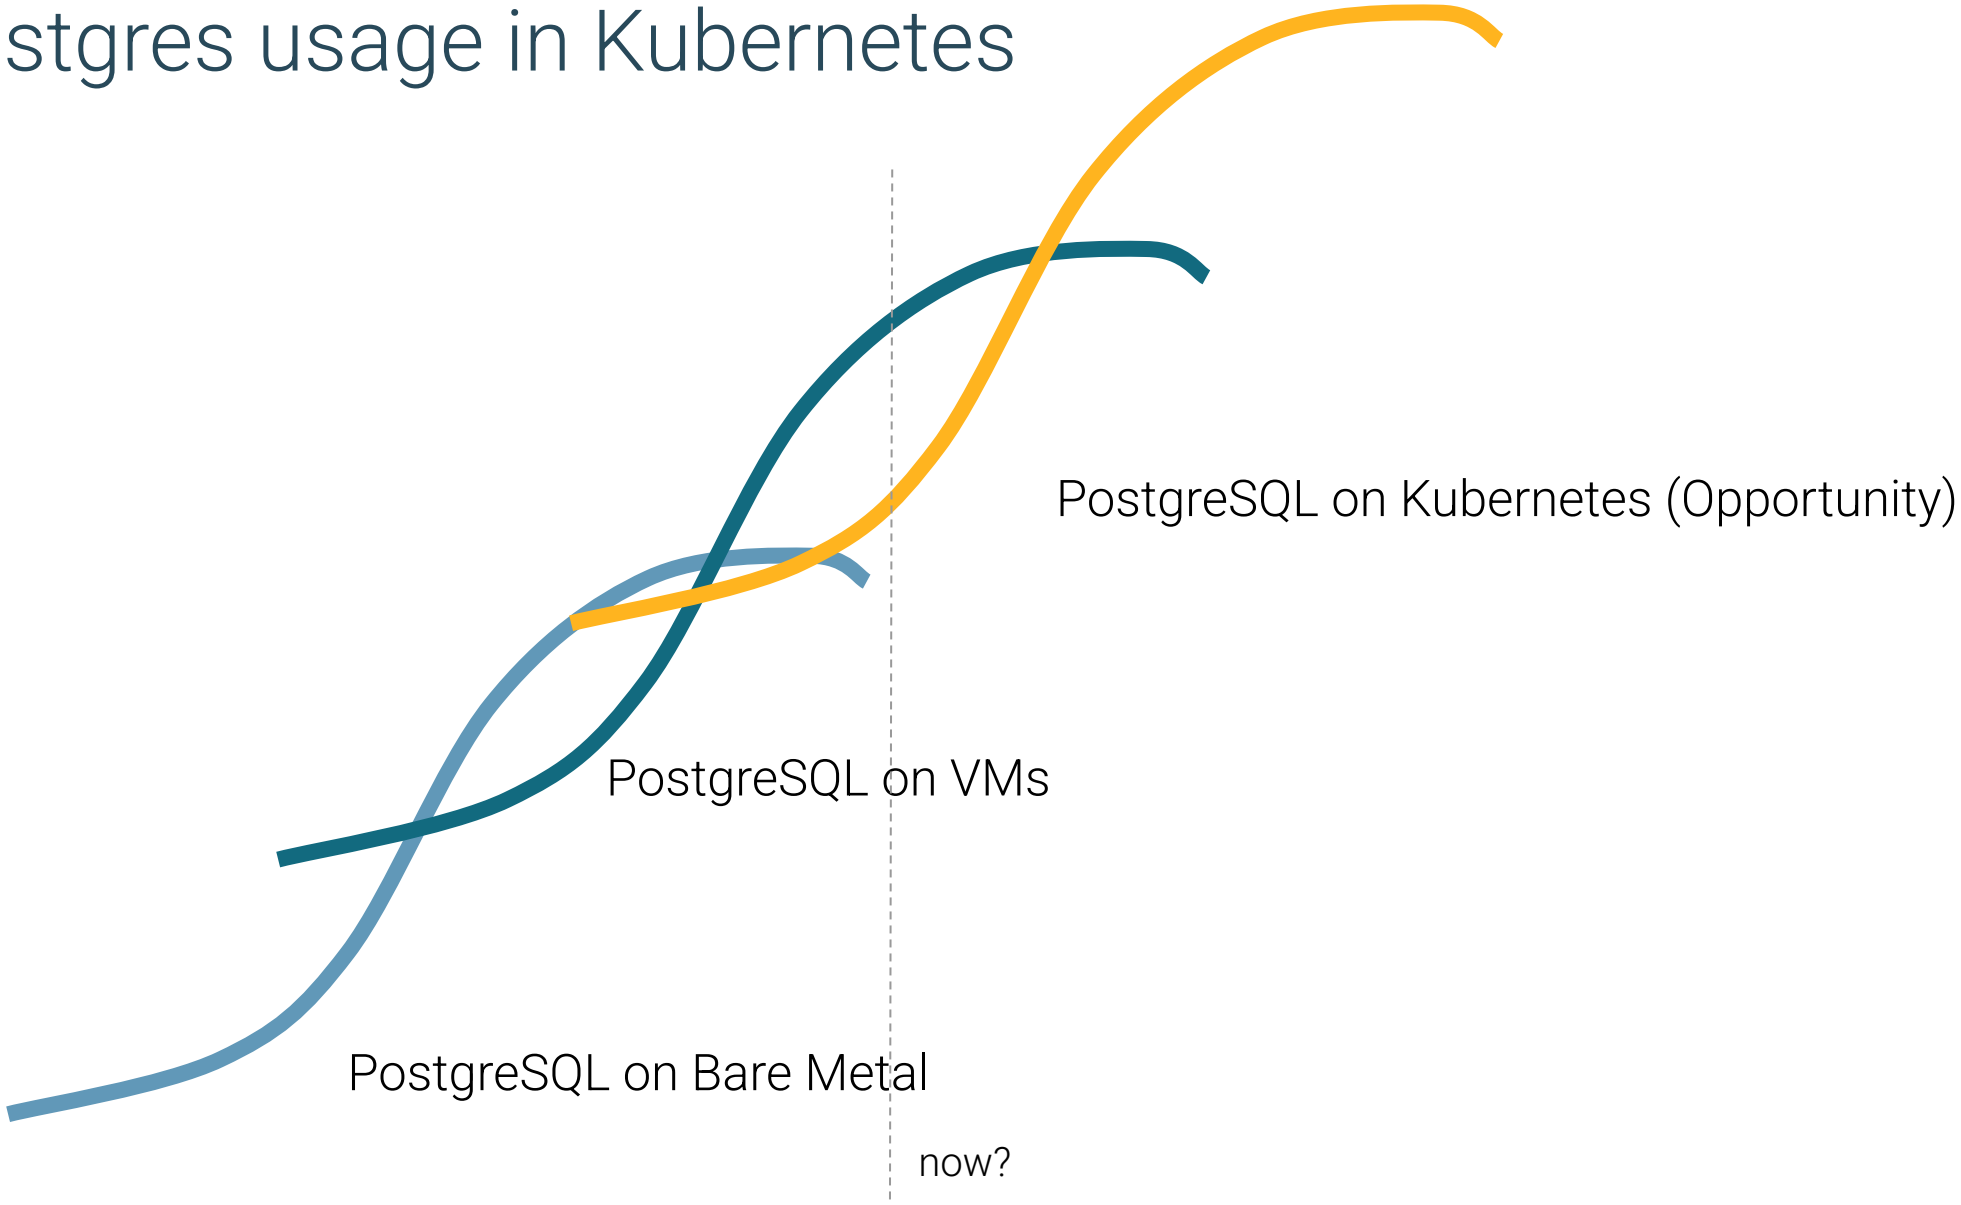

EDB POSTGRES AI

EDB Postgres AI Platform

Superpowers

1 - SOVEREIGN CONTROL
Cloud service automation in private customer-controlled environment (Hybrid DBaaS)

2 - ENTERPRISE SECURITY
Protect your data: Transparent data encryption, SQL protection, audit

3 - SECURE SUPPLY CHAIN
Mitigate open-source vulnerabilities. Software Bill of Materials

4 - LEGACY MODERNIZATION
Oracle compatibility mode and migration assistance.
80% Licensing Drop

5 - GLOBAL HIGH AVAILABILITY
Up to 99.999% uptime. Active/active distributed clusters serving global users

6 - CLOUD-NATIVE DEPLOYMENT
The #1 most popular Kubernetes operator. PG on K8s with confidence

7 - HYBRID OBSERVABILITY
Monitoring and management across clouds and on prem.Up to 30% productivity boost

8 - Analytics & AI ACCELERATION
Power GenAI and Agentic AI with sovereign Postgres and Analytics Accelerator. Up to 30x faster queries.

9 - Analytics & ML At Scale
Consolidate petabyte-scale analytics workloads based on Postgres with “Greenplum”

EDB POSTGRES AI

WarehousePG

EDB POSTGRES AI

WarehousePG

WarehousePG (WHPG) is an advanced, fully featured, open source data warehouse, based on Greenplum® Database and PostgreSQL.

WarehousePG provides powerful and rapid analytics on petabyte scale data volumes.

Uniquely geared toward big data analytics, WarehousePG is powered by the world's most advanced cost-based query optimizer delivering high analytical query performance on large data volumes.

EDB POSTGRES AI

What is WarehousePG?

Key Terminology

WHAT

Lakehouse

- Centralized data repository for business intelligence and reporting.
- Structured data/Semi-Structured/Non-Structured
- Optimized for high-performance analytical queries at petabyte scale.

HOW

Massively Parallel Processing (MPP)

- Distributes data and queries across multiple nodes for fast execution.
- Each node processes tasks independently, then aggregates results.
- Enables scalability for petabyte-scale analytics.

WHY IT'S RELEVANT FOR US

Postgres-Based

- Based on Postgres.
- WarehousePG is essentially a bunch of Postgres instances working together.
- Advantage over traditional data warehouse with Postgres extension compatibility for semi-structured and unstructured data support.

A Brief History of Greenplum

Greenplum Open Sourced **2015**

Broadcom Makes Greenplum Closed Source **May 2024**

2003

Bizgres

Greenplum, Inc. sponsors OSS Bizgres Project

2006-2009

Greenplum

Rebrands Bizgres MPP development to Greenplum Database

2010

GREENPLUM A DIVISION OF EMC

Acquired by EMC

Focus on Hardware Appliances

2013

Pivotal Greenplum

Spun out as an asset of Pivotal Software

Focus on Platform as a Service, Consulting

2015 - 2019

GREENPLUM DATABASE

Greenplum Database goes Open Source

2020 - 2023

VMware Tanzu™ Greenplum®

Acquired by VMware

Focus on Virtualization, Application Delivery

**2023 - **

vmware by Broadcom

Acquired by Broadcom

The end of open source Greenplum

EDB forks **WarehousePG** from the last open source version of Greenplum and launches **EDB Postgres AI - Support for Greenplum Workloads**

April 2025

WarehousePG

EDB POSTGRES AI

EDB POSTGRES AI

What is WarehousePG?

WarehousePG in Action

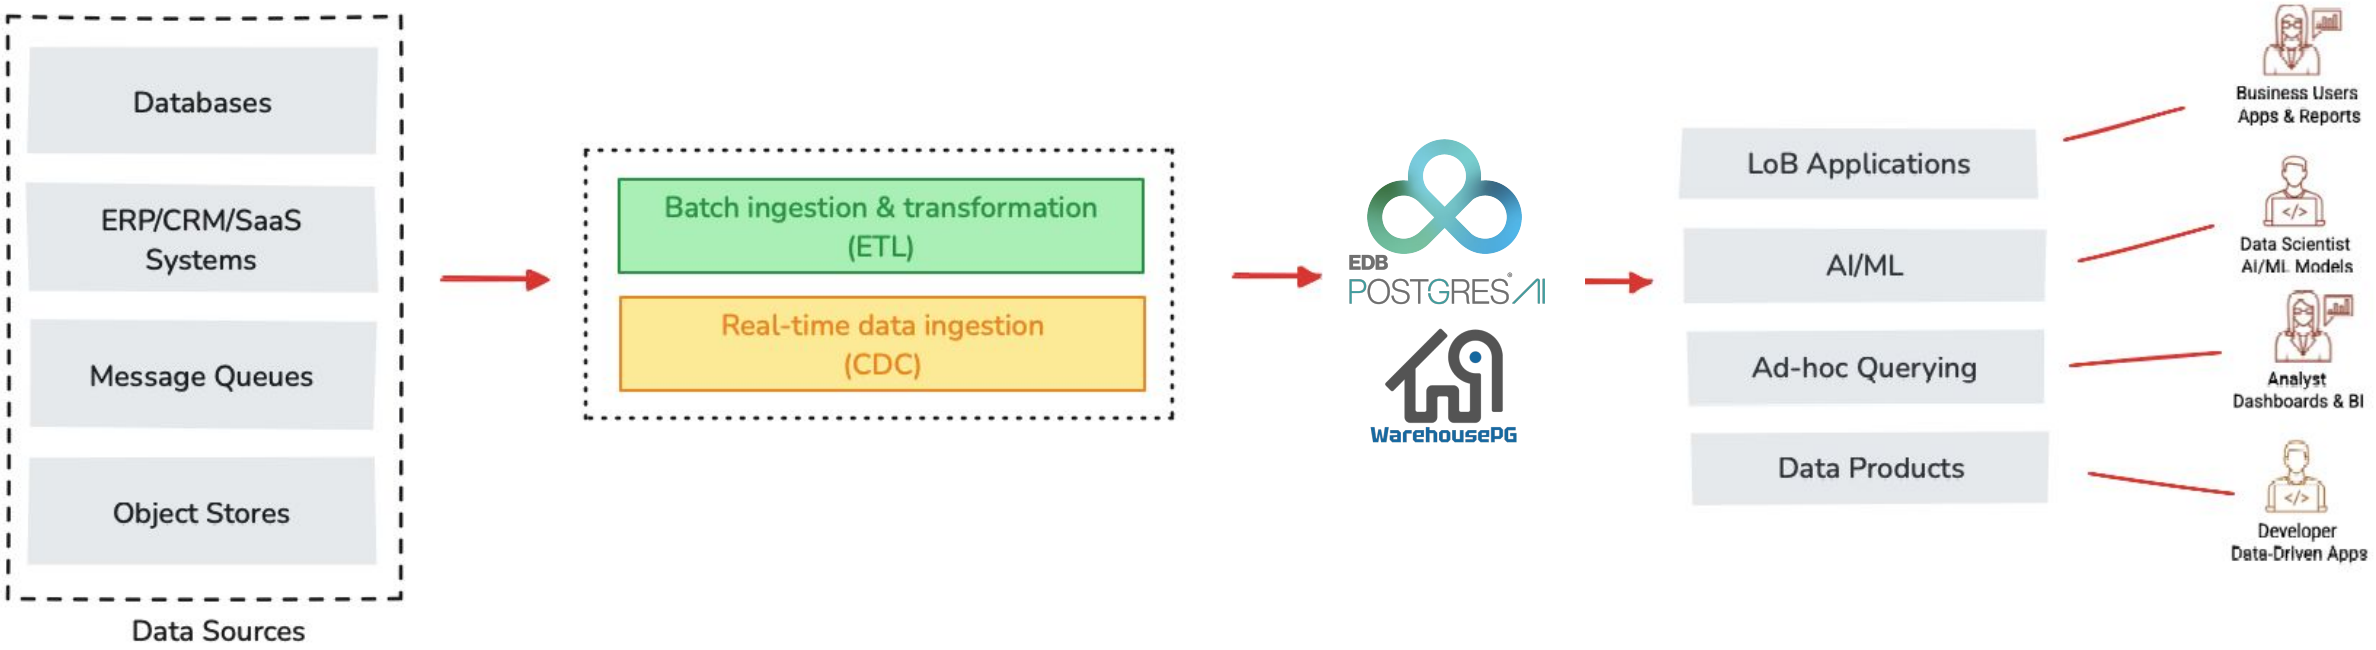

Drive Advanced Analytics with EDB WarehousePG

Petabyte Scale

Massively Parallel Processing (MPP) Cluster topology on multiple Segments

In-DB Machine Learning

Accelerate innovation, run at scale without needing to transfer data to other tools

Parallel Query Distribution

Master (coordinator) to optimize and parallel plan query distribution on multiple segments

GPUs Support

Leveraging GPU powered ML frameworks like PyTorch, TensorFlow via e.g. CUDA

Heterogeneous Data Access

Enables parallel, high throughput data access and federated queries through PXF (platform extension framework)

ELT Transformation Support

Accelerate data transformation pipelines through 3rd party push down integration to WarehousePG

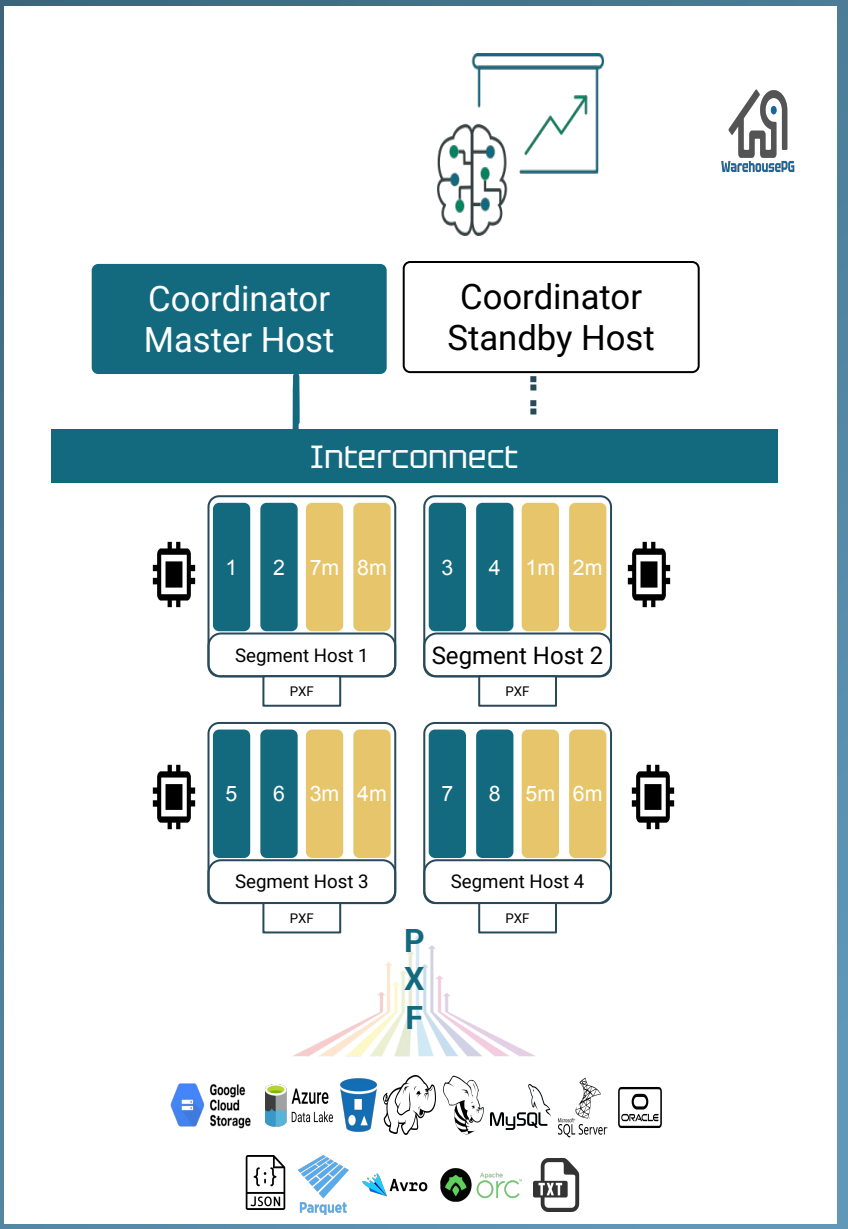

WarehousePG Highlights

Petabytes scale with high-speed /-concurrency performance, via MPP Cluster topology

Postgres Fork based on PG 12.12

Master (coordinator) for query optimization and parallel plan and distribution

Multiple Segments

- Segment is a Postgres Process which holds data
- Own CPUs, Hard disk and RAM
- Scalable

HA achieved through Standby for Master and Mirror for Segments

High-speed Interconnection for continuous data processing pipelining

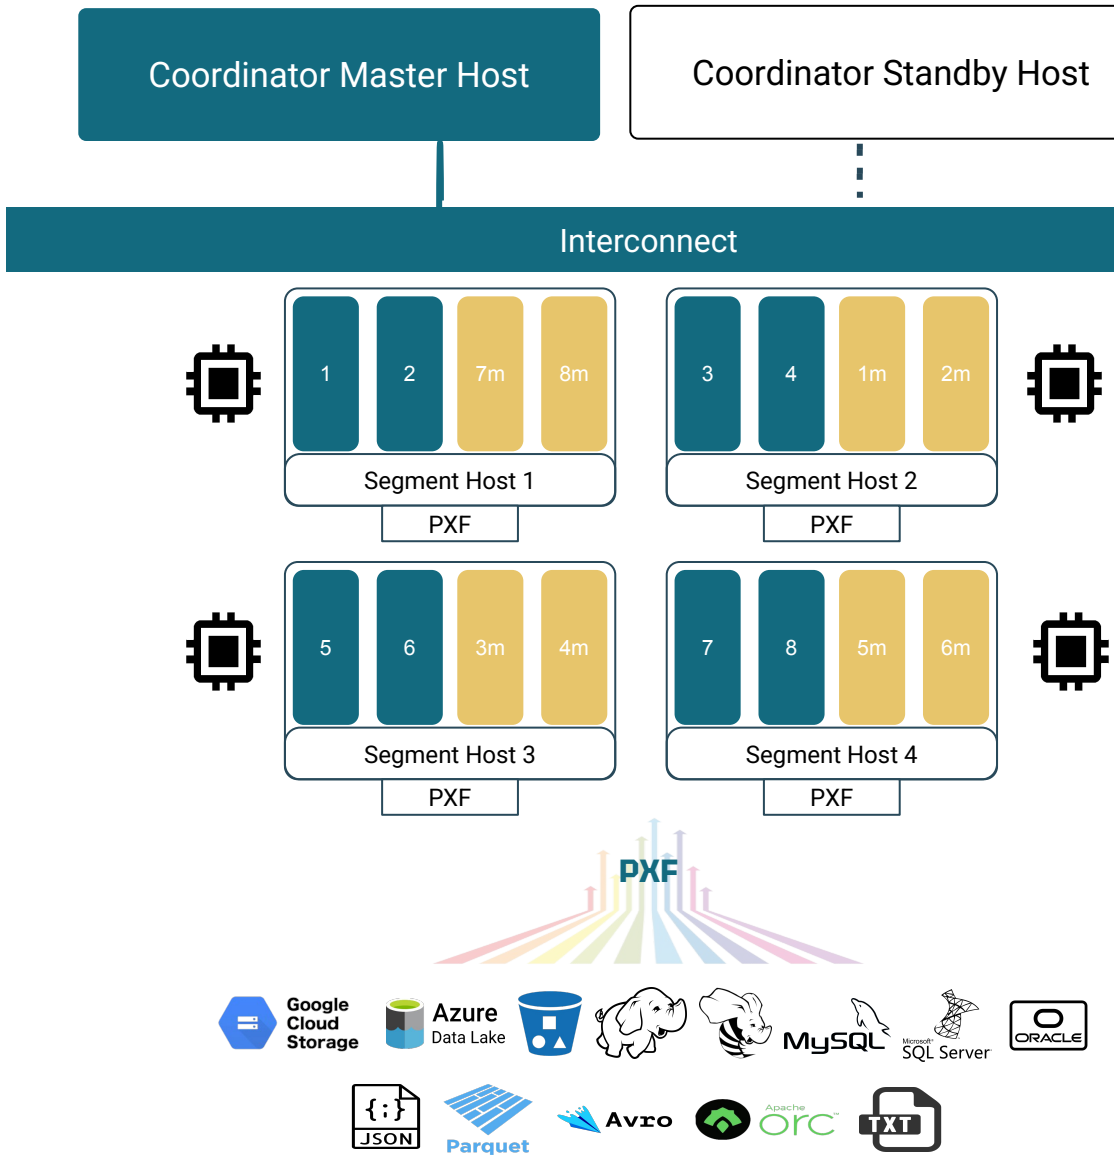

Primary Segments
Mirror Segments

Analytic Accelerator

EDB POSTGRES AI

EDB Postgres AI Analytics Accelerator

Postgres Lakehouse is a new type of Postgres cluster (it's really just one node) that you can provision with on the EDB Postgres® AI Platform.

It includes a vectorized query engine based on Apache DataFusion for fast queries over columnar Lakehouse tables in object storage (using the Delta Lake protocol).

EDB POSTGRES® AI

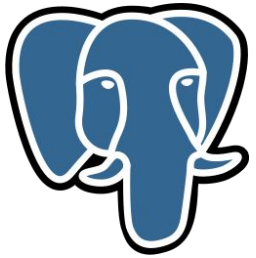

pgvector

APACHE DATAFUSION™

What is EDB Postgres AI Analytics Accelerator?

Key Terminology

WHAT

Analytics Accelerator

- Use the Analytics Accelerator to explore the analytical capabilities built on EDB Postgres®.
- This accelerator helps you understand core concepts, explore key technologies such as EDB Postgres® Lakehouse, and learn how to implement analytics with EDB Hybrid Manager (HM).

HOW

Mix of Open Source Technologies

- Apache DataFusion
- Apache Iceberg
- Delta tables

WHY IT'S RELEVANT FOR US

Postgres-Based

- Based on Postgres.
- Analytics Accelerator is essentially Postgres mixed to other technologies.
- Performance and AI integration.

EDB Postgres AI Analytics Accelerator Sample Architecture

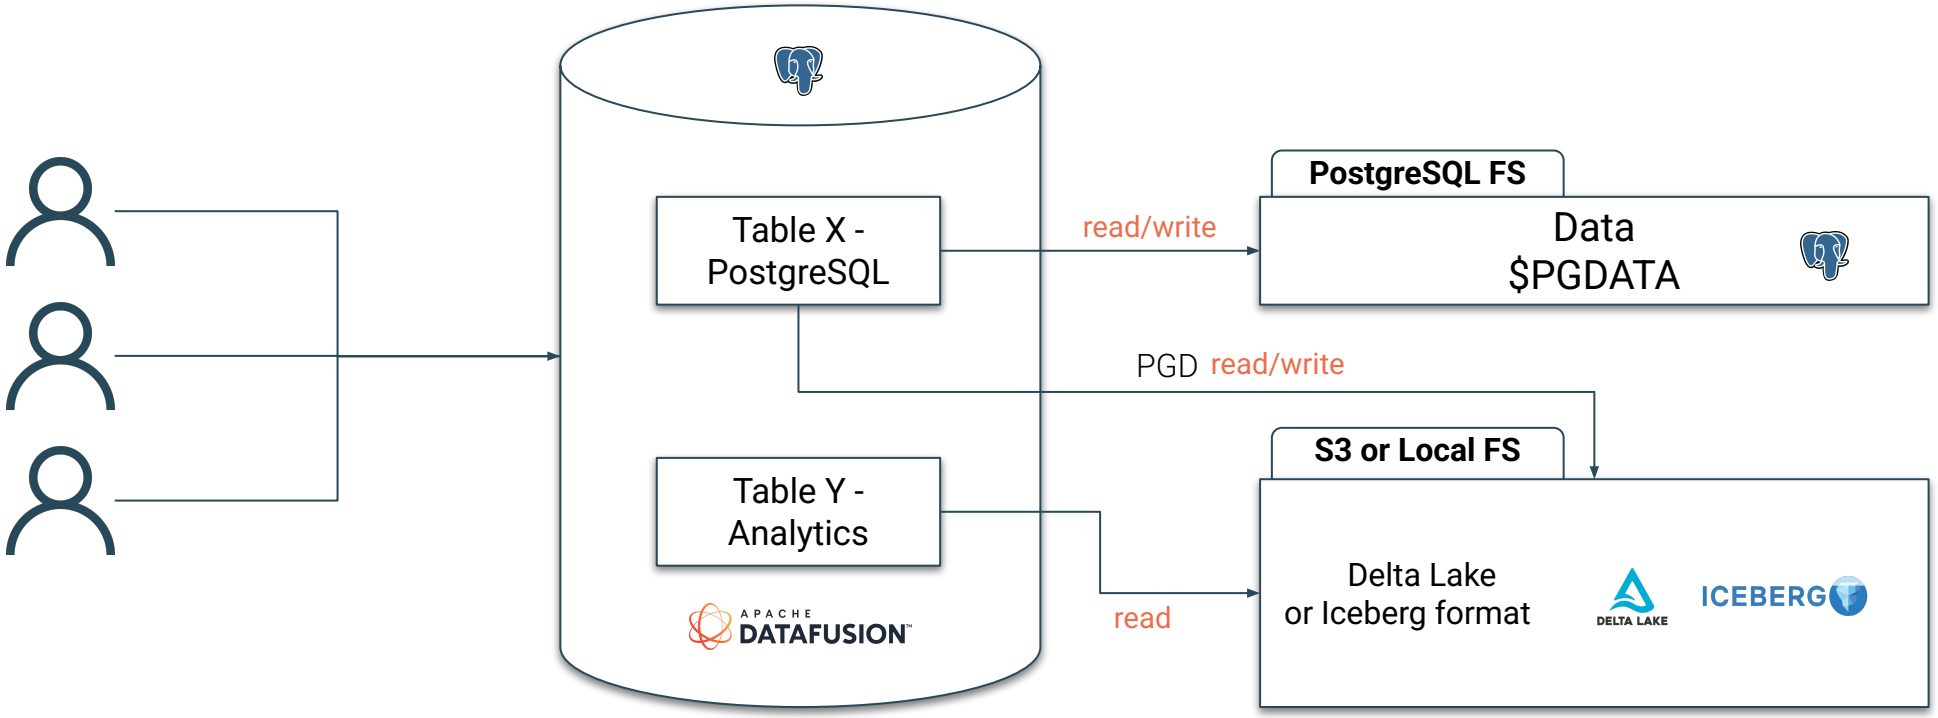

EDB POSTGRES AI

Analytics Engine

GET 30X FASTER INSIGHTS DIRECTLY FROM OPERATIONAL DATA

Challenges

- Slow insights
- Business disruptions and downtime
- High costs from bloated storage

EDB Postgres AI Solution »

- **Optimized for Columnar Data:** 30x faster analytical queries
- **Tiered Tables:** Improved cost efficiency (5x smaller formats, 18x more economical storage), Up to 99.999% availability (unaffected by analytics)

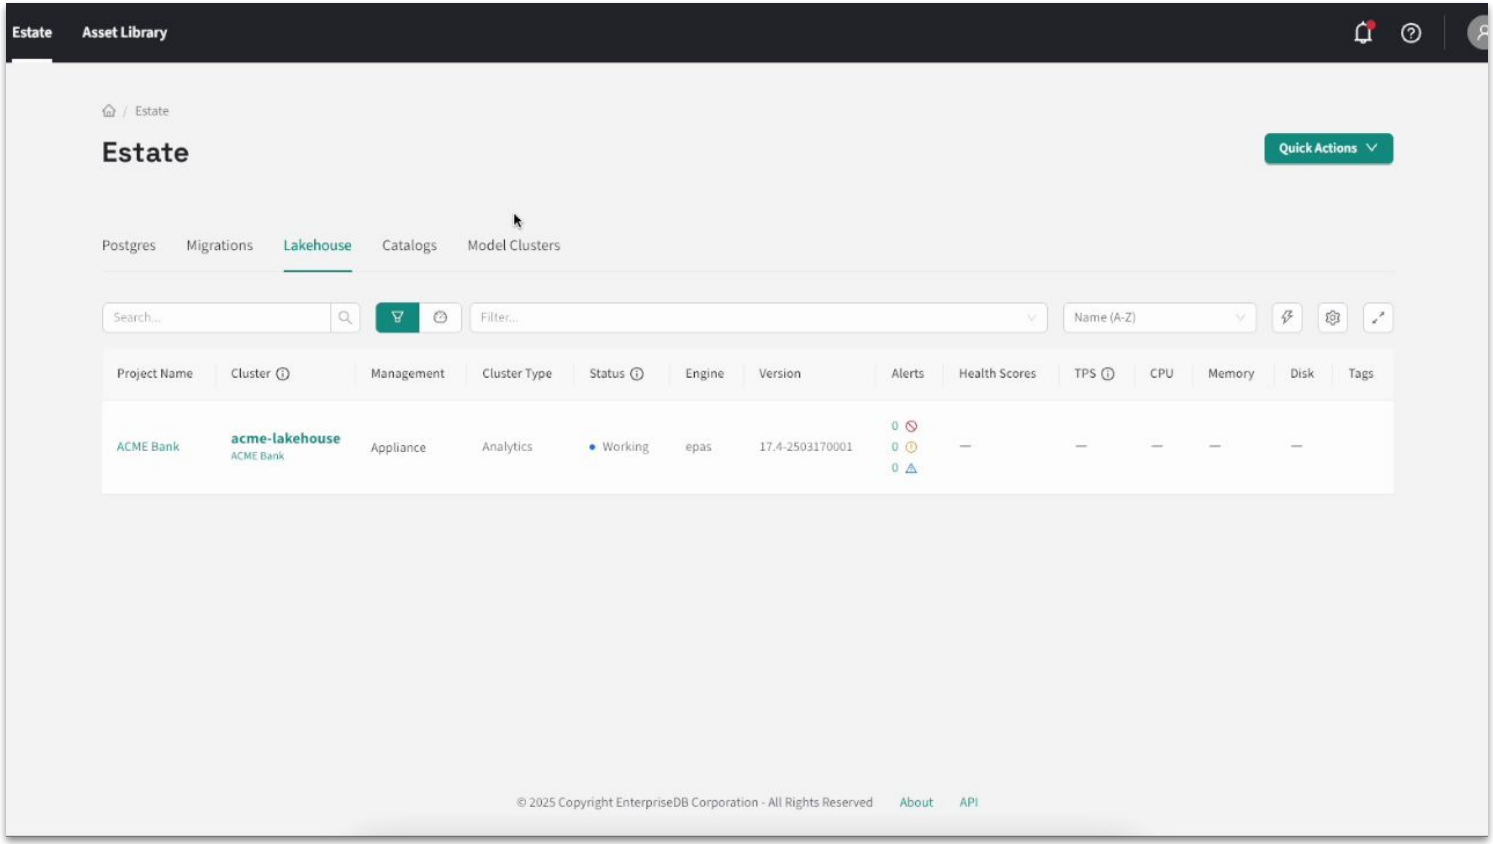

WATCH DEMO >>

Watch a demo showing how Analytics Accelerator improves dashboard performance by improving query speeds vs. the standard Postgres engine.

EDB POSTGRES AI

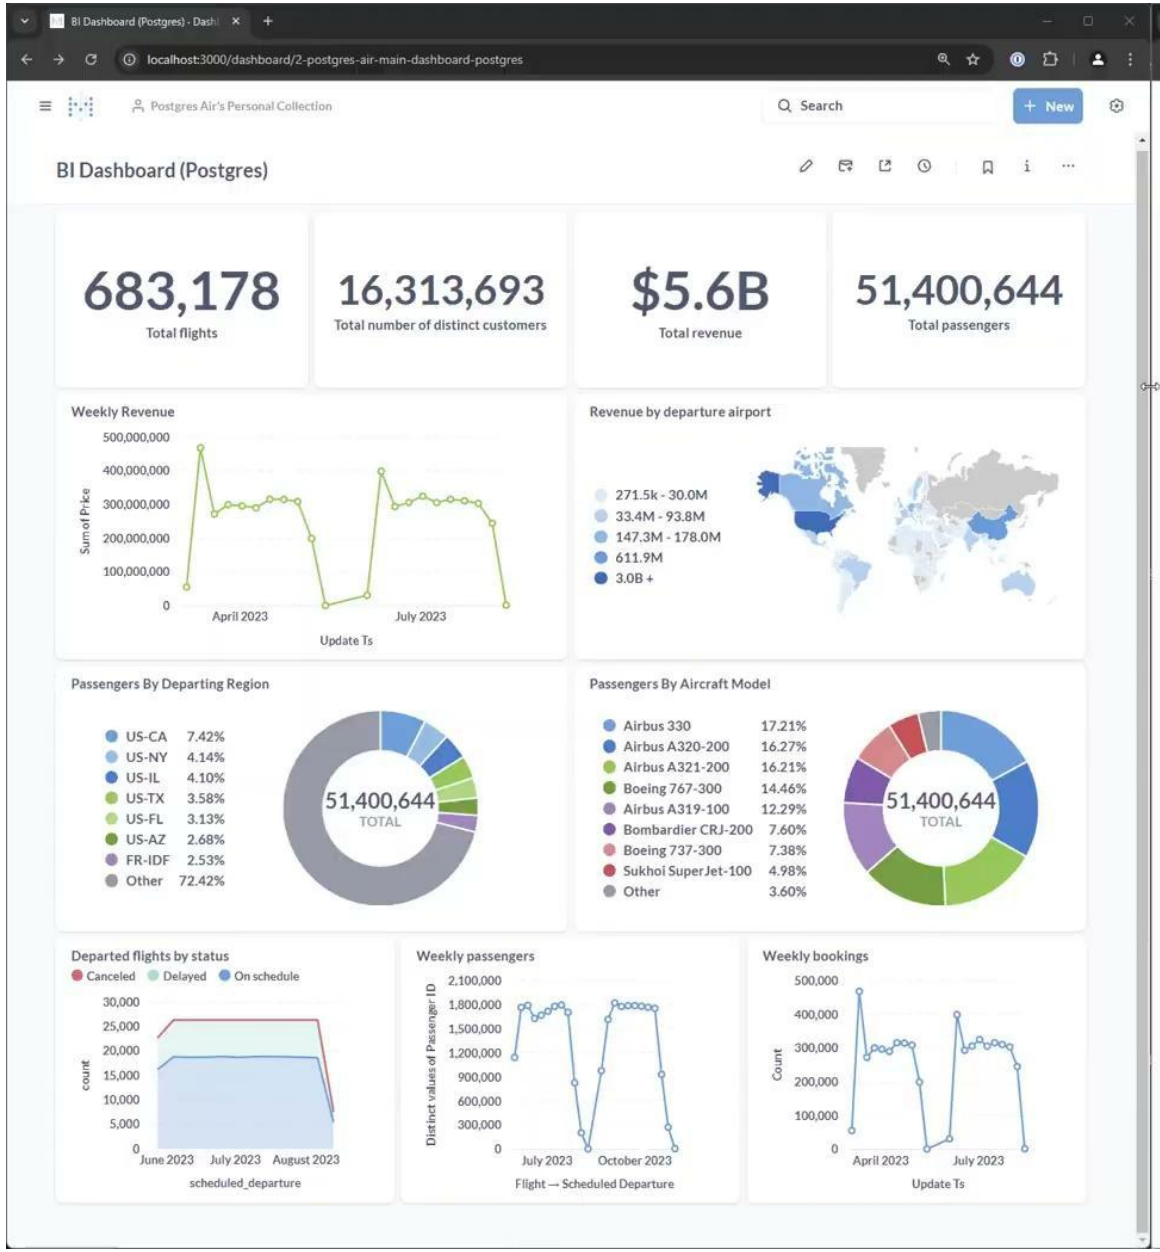

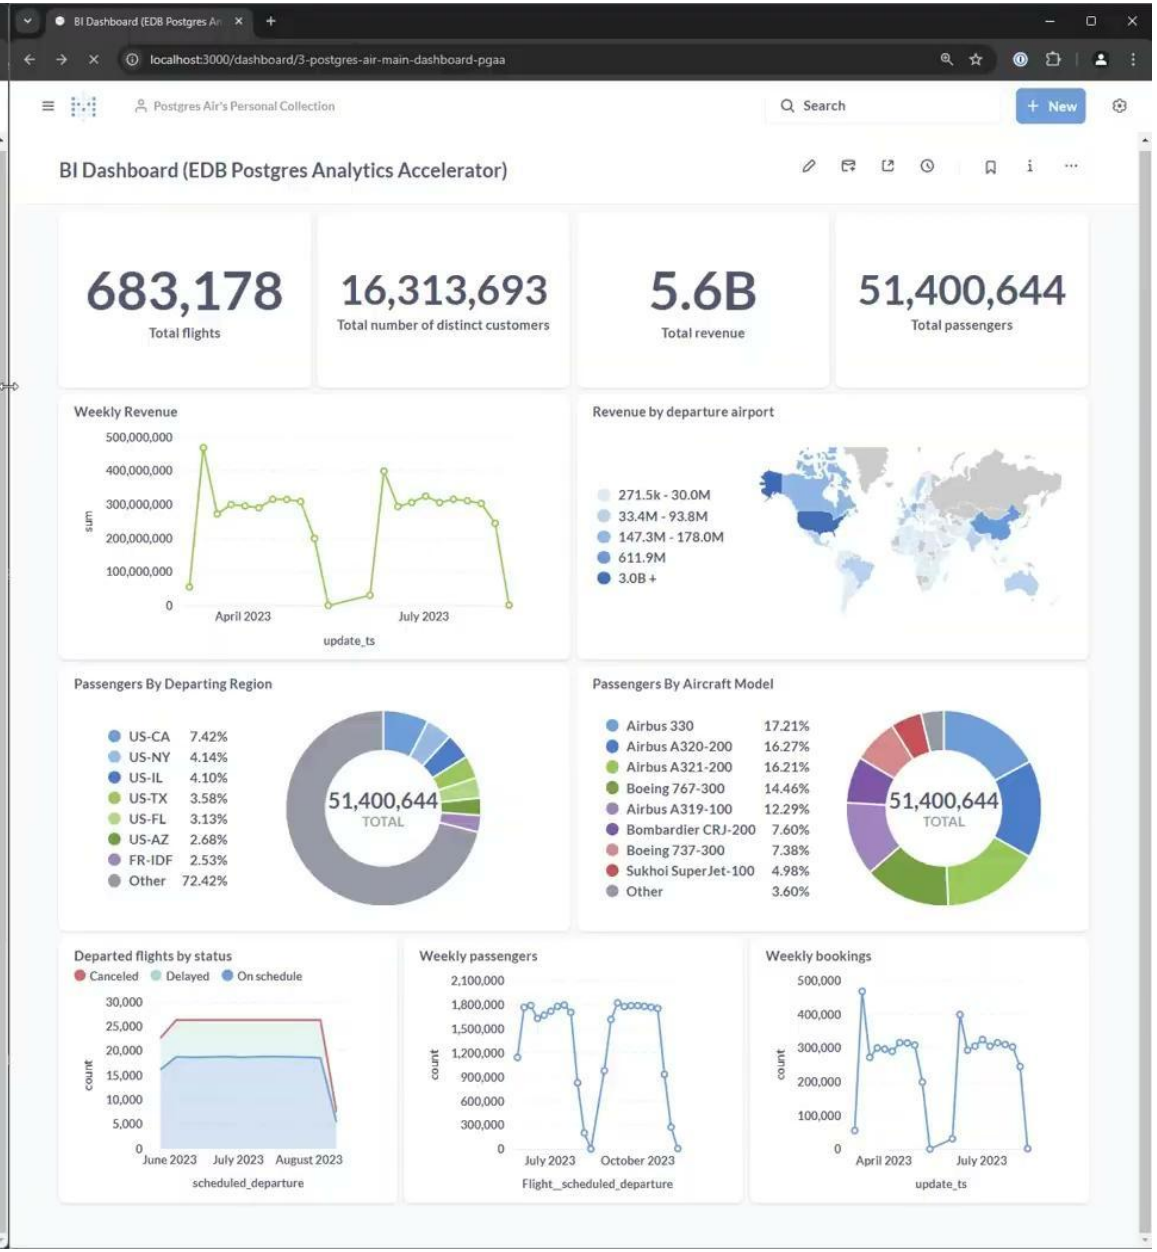

Thank you

EDB POSTGRES AI

CUFEP 2025

Part 2: EDB Postgres AI & Strategy
From Transactional to AI, Through Analytics at Scale

Sergio Romera, **Senior Manager, Sales Engineering South Team**

Agenda

Secure OSS

Hybrid
Control Plane

Roadmap

EDB POSTGRES AI

EDB Postgres® AI

Database

- ENTERPRISE POSTGRES SERVER
- ORACLE COMPATIBLE SERVER
- COMMUNITY POSTGRESQL SERVER
- MULTI-MODEL EXTENSIONS
- MANAGEMENT & OBSERVABILITY
- KUBERNETES OPERATORS
- SUPPLY CHAIN SECURITY
- MIGRATION TOOLS
- HIGH AVAILABILITY

Analytics Accelerator

- ANALYTICS ENGINE
- LAKEHOUSE CONNECTOR
- SUPPORT FOR GREENPLUM WORKLOADS

AI Factory

- VECTOR ENGINE
- AI PIPELINE
- GENAI BUILDER
- AGENT STUDIO
- MODEL SERVING

Hybrid Management

- HYBRID OBSERVABILITY
- HYBRID DBAAS
- DISTRIBUTED HA (99.999%)
- MIGRATIONS

Deploy Anywhere

- SOVEREIGN DATA AND AI FACTORY
- HYBRID SOFTWARE
- MANAGED PLATFORM
- INTEGRATION PARTNERS: AWS, GCP, AZURE, RED HAT OPENSHIFT, IBM, SUPERMICRO, NVIDIA, GRIPTAPE

Sovereign Assurance

- MANAGED SERVICES
- PILOT TO PRODUCTION SLAS
- OUTCOME EXECUTION

EDB POSTGRES AI

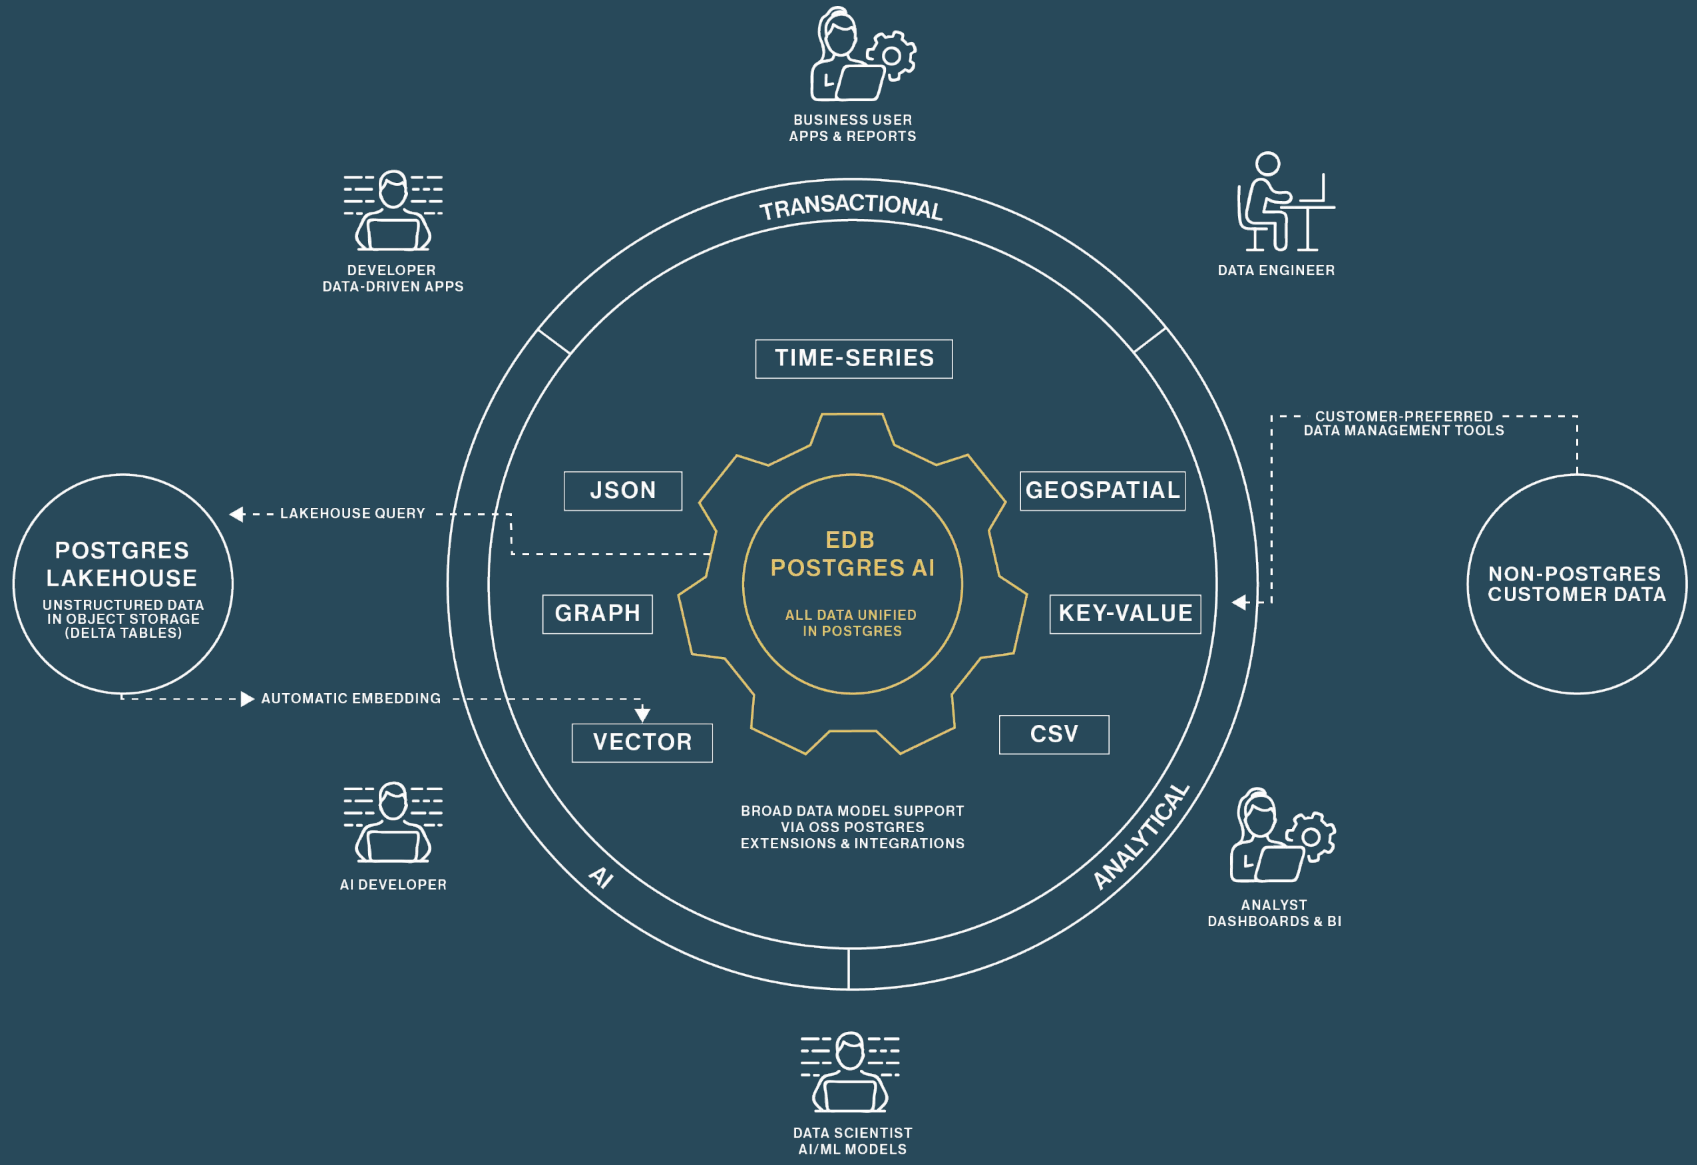

EDB POSTGRES AI

EDB POSTGRES AI HYBRID MANAGEMENT

UNLOCK HYBRID CONTROL AND UNIFIED OBSERVABILITY IN A SOVEREIGN CONTAINER

Turnkey data and AI sovereignty

Operate efficiently with data platform automation

Take action faster with deep observability

Securely modernize for AI-ready data

KEY FEATURES

- **Secure, end-to-end control:** Encrypt and contain data to prevent unwanted access.
- **Cloud-native automation:** Automate your hybrid data infrastructure on demand, on your terms.
- **Unified observability:** 200+ metrics for optimal performance and HA for EDB and external Postgres.
- **Intelligent recommendations:** Identify problems 5x faster and boost app performance by up to 8x.
- **Sovereign migrations:** Seamless migrations within your secure environment.
- **AI-ready data**: Sync existing data and create a private knowledge base to modernize legacy apps faster.

EDB POSTGRES AI

Unified Observability

UNLOCK DEEP OBSERVABILITY FROM A SINGLE PANE OF GLASS

Challenges

- Fragmented deployments
- Problematic queries
- Slow app performance

EDB Postgres AI Solution »

- **Single pane of glass:** View 100s of Postgres clusters
- **Query diagnostics, alerts, notifications:** Resolve bottlenecks 5x faster
- **Intelligent recommendations:** 8x faster app performance

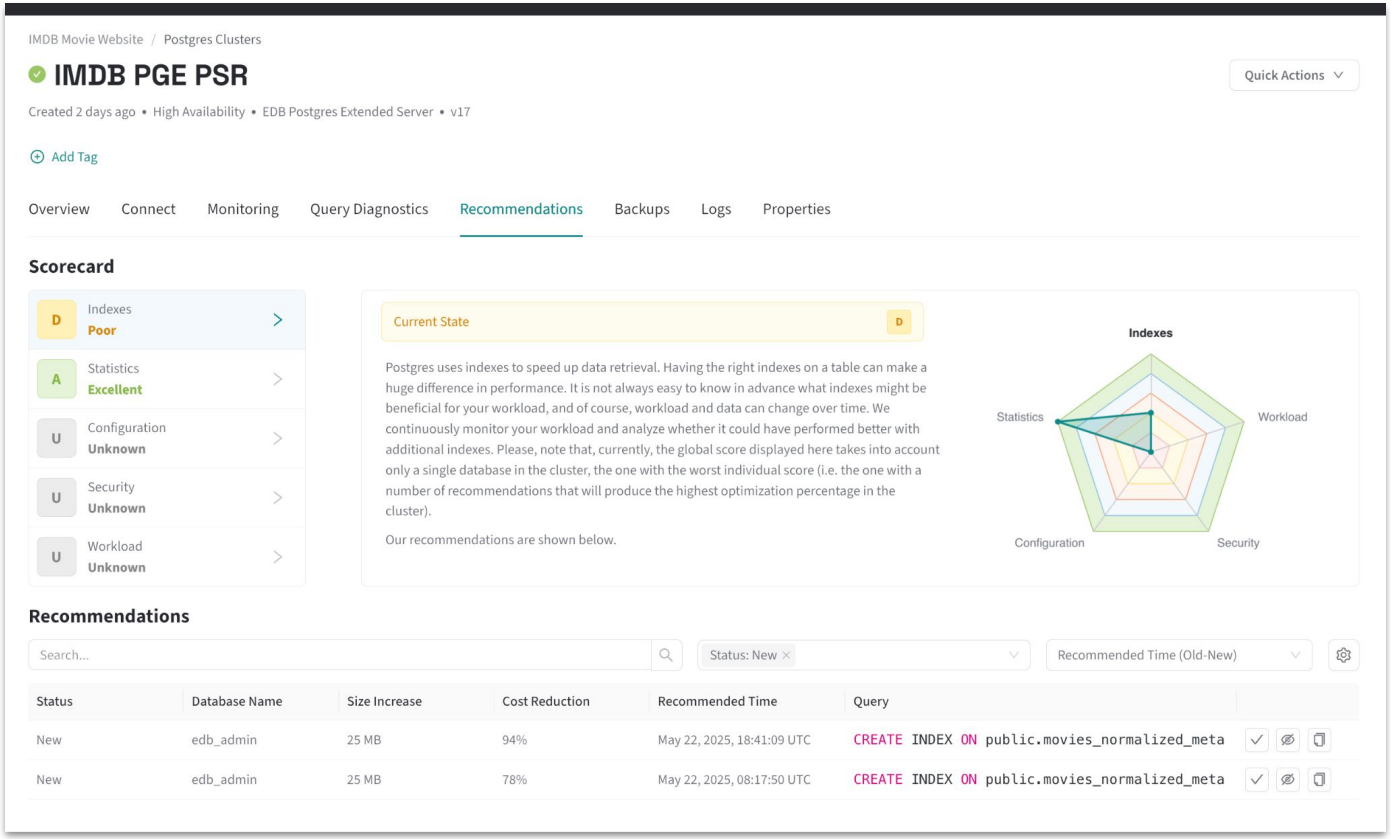

WATCH DEMO >>

Watch a demo showing how to see and apply recommendations from EDB Postgres AI to resolve problematic queries and improve performance.

EDB POSTGRES AI

Hybrid DBaaS

AGILITY WITH SOVEREIGN CONTROL IN HYBRID ENVIRONMENTS

Challenges

- Multiple, disparate tools
- DBAs spend 30-50% of time on undifferentiated admin

EDB Postgres AI Solution »

- **Console:** Manage 100s of clusters with a simple interface
- **Automated provisioning, backups, recovery, upgrades:** Boost productivity 30%
- **Access management:** security on your own terms

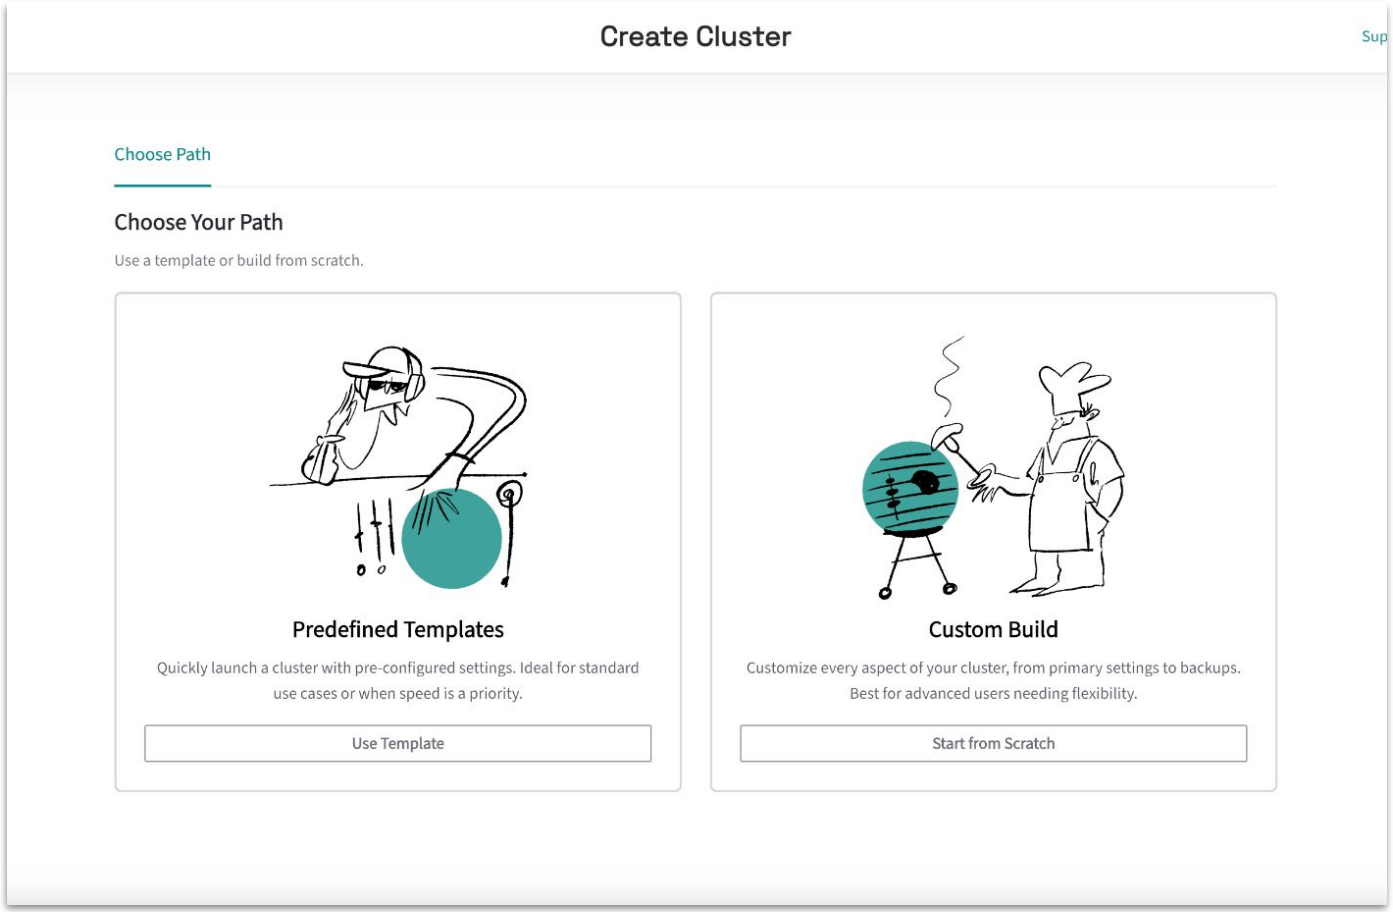

WATCH DEMO >>

Watch a demo showing how to create a cluster template using just a few clicks in an easy-to-use interface.

EDB POSTGRES AI

EDB POSTGRES AI DATABASE

SINGLE DATABASE PLATFORM FOR ANY WORKLOAD, READY FOR SOVEREIGN AI.

Build without traditional boundaries

Deliver always-on experiences

Secure your data assets for sovereign AI

Boost productivity for busy teams

KEY FEATURES

- **Multi-model database:** Relational, document, time series, columnar, vector, and more.
- **Distributed HA:** 99.999% HA, built on a geo-distributed, active-active architecture.
- **Secure by default:** Transparent Data Encryption and data, database, and app protection eliminate threats.
- **Flexible deployment**: Automate management across hybrid and multi-cloud environments.

EDB POSTGRES AI

High Availability

DELIVER ALWAYS-ON EXPERIENCES WITH UP TO 99.999% HA

Challenges

- HA requirements for mission-critical apps
- Difficulty setting up HA clusters

EDB Postgres AI Solution »

- **High availability options (HA, Advanced HA, Distributed HA):** Up to 99.999% HA
- **Automated provisioning:** Few clicks to deploy HA clusters

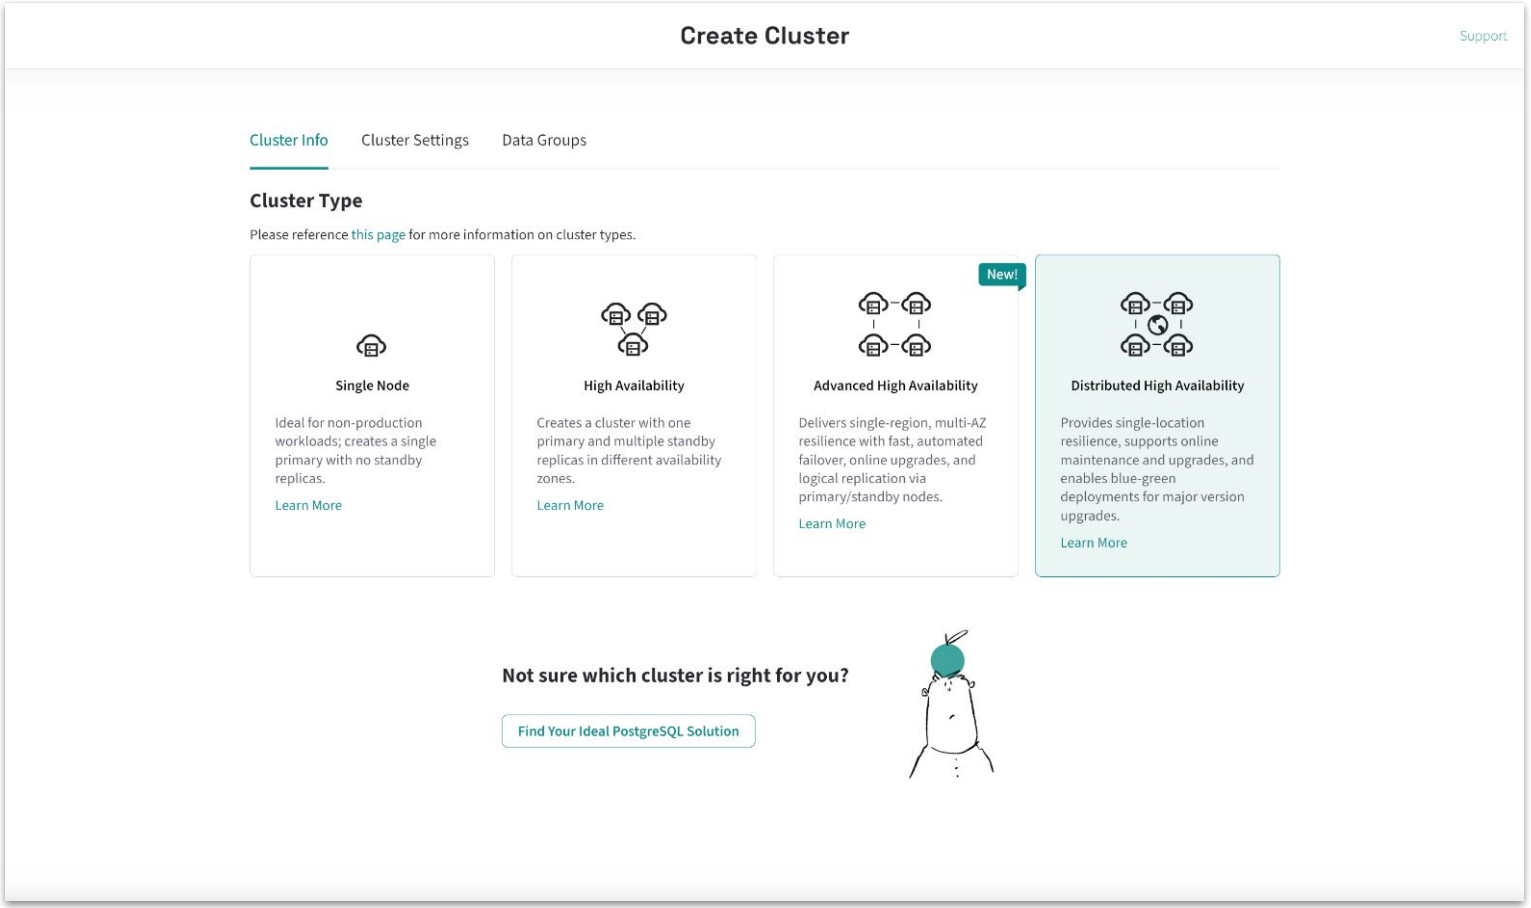

WATCH DEMO >>

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

EDB POSTGRES AI

EDB POSTGRES AI FACTORY

SECURELY BUILD, TEST, AND LAUNCH SOVEREIGN AI APPLICATIONS

Accelerate time to market for GenAI apps

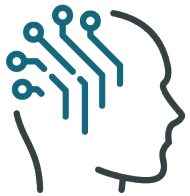

Equip every team to build custom, sovereign AI

Design only once with secure, automated data management

Deliver more value from your data with agentic AI

KEY FEATURES

- **GenAI builder:** Accurate GenAI with robust low-code SDK for devs and no-code interface for others.
- **Agent studio:** Deploy out-of-the-box agents or build custom agent workflows to enhance productivity.
- **AI pipeline:** Auto-embedding in a single system for always-current knowledge bases without the complexity.
- **Model serving:** Run customized AI models on a secure vector database without exposing data to the cloud.
- **Hybrid deployment:** Leverage your own infrastructure or a pre-configured solution with Nvidia and Supermicro.

EDB POSTGRES AI

GenAI Builder

EMPOWER EVERY TEAM TO BUILD TAILORED, SOVEREIGN GENAI APPLICATIONS

Challenges

- AI development skill gaps
- 6-12 month GenAI time to value
- Budget overruns of 40-60%

EDB Postgres AI Solution »

- **Build your way:** Any skill level can build production GenAI in days, not months
- **Off-prompt:** Ensure data sovereignty and reduce costs by storing sensitive data in-memory instead of exposing to 3P LLMs

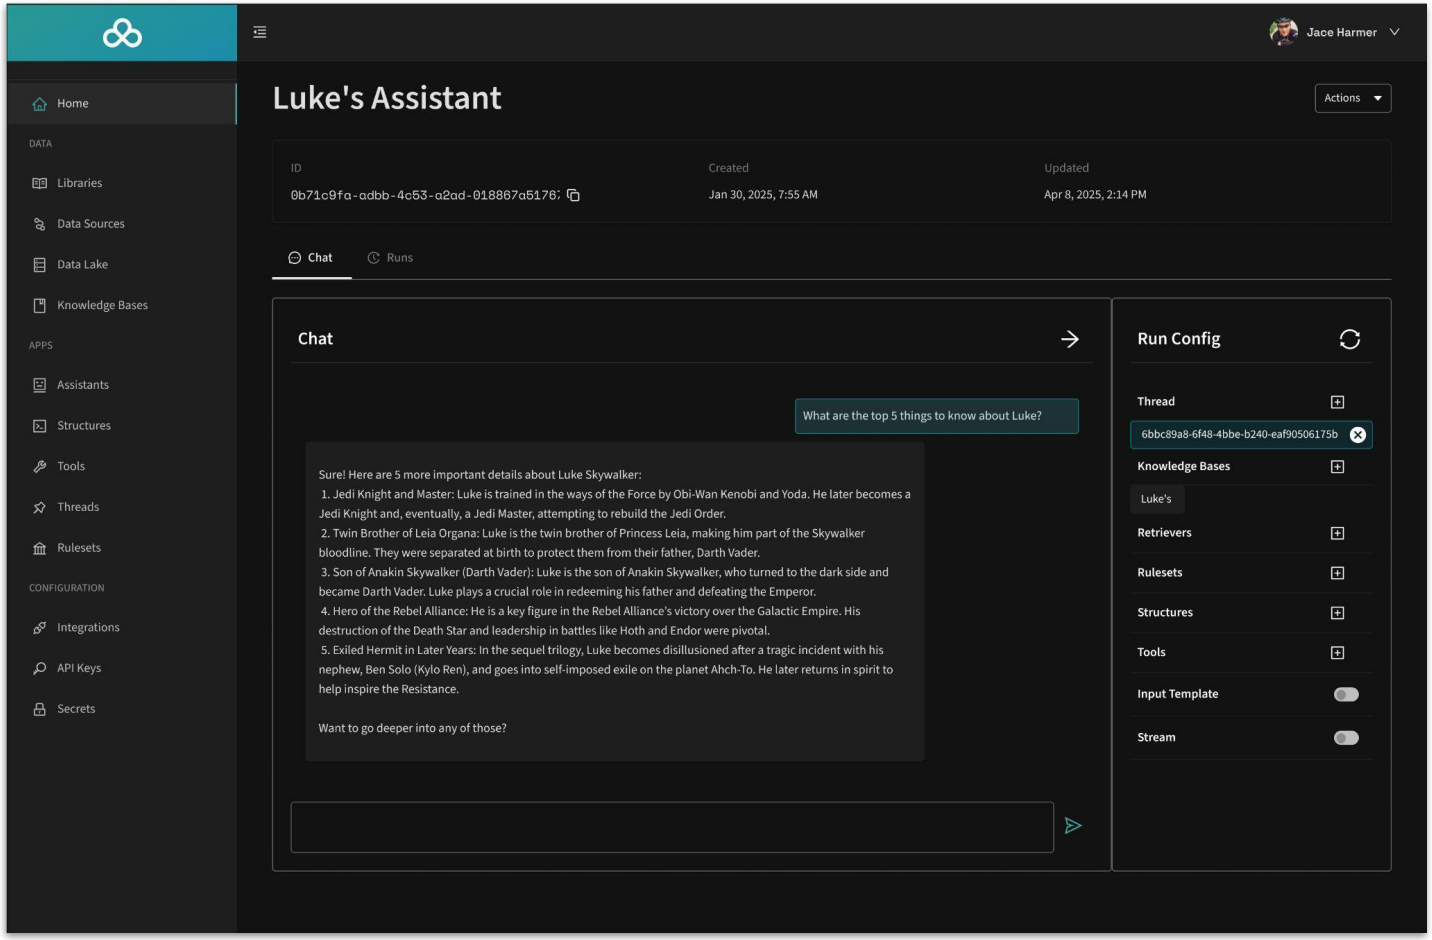

WATCH DEMO >>

Watch an end-to-end demo of how to build a GenAI application with AI Factory.

EDB POSTGRES AI

Analytics Engine

GET 30X FASTER INSIGHTS DIRECTLY FROM OPERATIONAL DATA

Challenges

- Slow insights
- Business disruptions and downtime
- High costs from bloated storage

EDB Postgres AI Solution »

- **Optimized for Columnar Data:** 30x faster analytical queries
- **Tiered Tables:** Improved cost efficiency (5x smaller formats, 18x more economical storage), Up to 99.999% availability (unaffected by analytics)

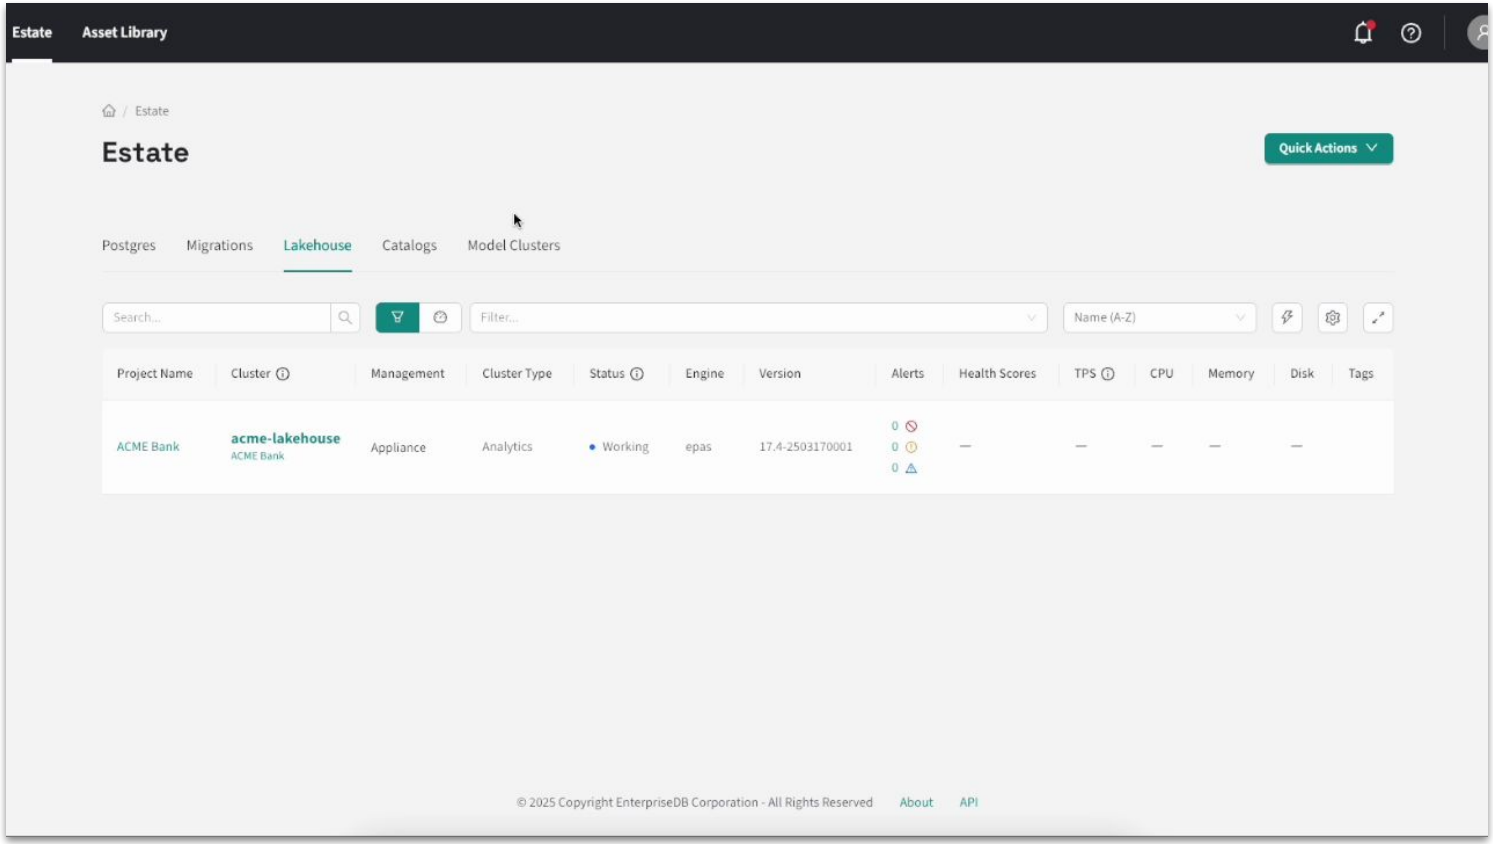

WATCH DEMO >>

Watch a demo showing how Analytics Accelerator improves dashboard performance by improving query speeds vs. the standard Postgres engine.

EDB POSTGRES AI

DEPLOY EDB POSTGRES AI ANYWHERE

ACHIEVE DATA AND AI SOVEREIGNTY IN YOUR ENVIRONMENT OF CHOICE.

HYBRID SOFTWARE

A single, sovereign software installation delivers a consistent experience and cloud agility across hybrid and multi-cloud environments.

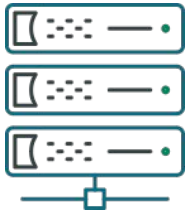

SOVEREIGN DATA AND AI FACTORY

Pre-configured solution integrates EDB Postgres AI with Supermicro servers to bring production-ready Postgres with 99.999% HA to your data center.

MANAGED PLATFORM

Leverage unmatched Postgres expertise and 24x7 support, delivering mission-critical performance and availability for hybrid deployments.

EDB POSTGRES AI

Sovereign Data and AI Factory

THE FAST TRACK TO DATA AND AI SOVEREIGNTY

Challenges

- Extreme performance and HA requirements
- DIY is slow and difficult
- Slow time-to-value for AI

EDB Postgres AI Solution »

- **Pre-tuned system:** Up to 6x performance, up to 5 9s HA
- **Turnkey AI Factory:** 5x faster sovereign AI development
- **Built-in hybrid management:** 90% better value than cloud DBs

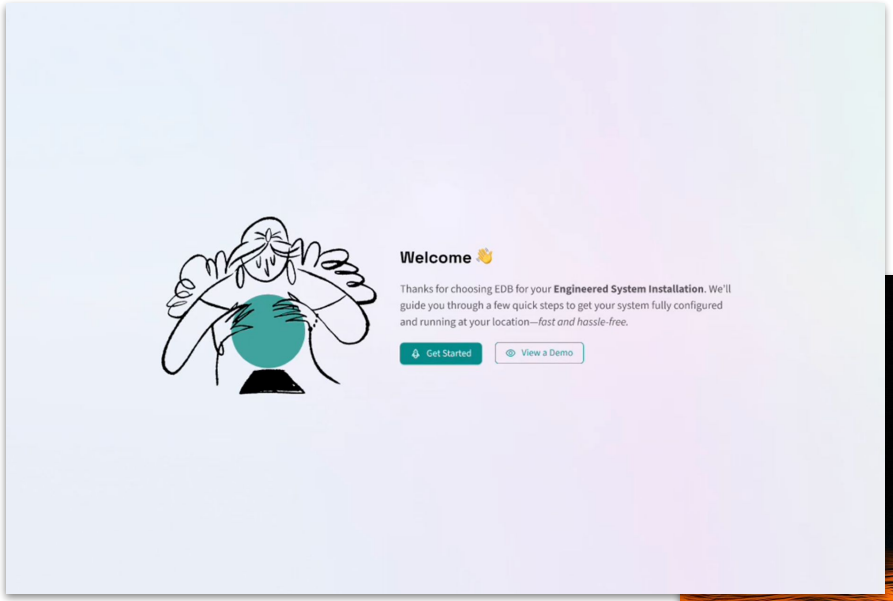

WATCH DEMO >>

See how we've transformed complex hardware installation into a simple, guided experience that enables production-ready Postgres in a few clicks.

EDB POSTGRES AI

Demo

EDB POSTGRES AI

Demo

01 State view

02 Create new DB

03 Sovereign AI

04 Observability recommendations

05 Analytics & AI

06 Q&A

EDB POSTGRES AI

Thank you

EDB POSTGRES AI

Disclaimer

Statements included in this presentation, other than statements or characterizations of historical fact, are forward-looking statements. These forward-looking statements are based on our current expectations, estimates, and projections about our industry, management's beliefs, and certain assumptions made by us, all of which are subject to change. These forward-looking statements are not guarantees of future results and are subject to risks, uncertainties, and assumptions that could cause our actual results to differ materially and adversely from those expressed in any forward-looking statement.

By sharing our product roadmap with you, we are not undertaking an obligation to develop the software with the features and functionality discussed herein.

The forward-looking statements in this presentation are made as of May 2025. We undertake no obligation to revise or update publicly any forward-looking statement for any reason.