EDB POSTGRES AI

Summer Academy

Stop Managing Silos: A Modern Analytics Architecture Built on Postgres

01 02 **03** 04

Malik Bilal - Senior Sales Engineer
Miles Richardson - Senior Director, Product Management

AGENDA

What we'll cover today

01 · The Silo Problem

Why teams keep bolting on a new analytics system every time query times climb.

02 · A Converged Architecture

Postgres, WarehousePG, and ClickHouse, all reading and writing one open catalog.

03 · Demonstration

Real-time anomaly detection across two engines and one shared catalog.

04 · The Trade-offs

What converging buys you, and where a standalone engine still makes sense.

05 · Architecture

How PGAA, WarehousePG, and ClickHouse work together under the hood.

06 · Getting Started

A staged path to converge, one workload at a time.

Bring your questions — we'll leave time for Q&A at the end.

EDB POSTGRES AI

First Poll

Please answer on your screen

1- What's driving the pressure to add (or replace) an analytics system in your organization?

2- What does your current technology stack include?

3- Is your current data infrastructure able to cope with the rapid adoption of coding agents?

4- Do you build "workarounds" (read-replicas, new clusters, special privileges, etc) for teams who do not trust the data warehouse or the data is not available quick enough?

The Cost of Bolting On Another System

EDB POSTGRES AI

WHY TEAMS REACH FOR A NEW SYSTEM

The pressures pushing organizations toward another platform

New workloads & use cases

Real-time dashboards, AI features, and fast-growing data volumes stretch what the existing database was built for.

Performance issues

Query times climb as data grows, and reporting starts competing with production traffic for the same capacity.

Licensing & infrastructure costs

Consumption-based pricing and proprietary appliances get harder to predict, budget, and scale.

Every pressure point looks like a reason to add a system. It's actually a reason to converge.

THE SILO PROBLEM

Each new system becomes another silo to feed, govern, and sync

More systems

Each new analytics engine is one more thing to run, patch, secure, and staff.

More pipelines

ETL jobs move data between systems and become their own point of failure.

More copies of the truth

The same data lives in multiple places, governed a little differently in each one.

The database was never the bottleneck. The silos were.

A Converged Architecture

ONE OPEN POSTGRES FOUNDATION

Transactional, analytical, and AI workloads, without moving data

01

OLTP

System of record

Mission-critical Postgres for core operations, the trusted source of truth.

02

ANALYTICS

Insight at scale

Analyze enterprise data at scale on the same platform. No separate warehouse, no copies.

03

AI

Intelligence in place

Models, search, and agents running inside the data through SQL, with governance built in.

ONE PLATFORM **No different products to integrate, data never moves between silos to create value.**

THE FULL DATA LIFECYCLE

THE ANALYTICS ESTATE TODAY

THE COST OF FRAGMENTATION

PIPELINE SPRAWL
Copies and ETL between every system

STALE INSIGHTS
Analytics always lag reality

RUNAWAY COST
Consumption bills with no ceiling

GOVERNANCE GAPS
Rebuilt for every new engine

VENDOR LOCK-IN
Formats and regions you don't control

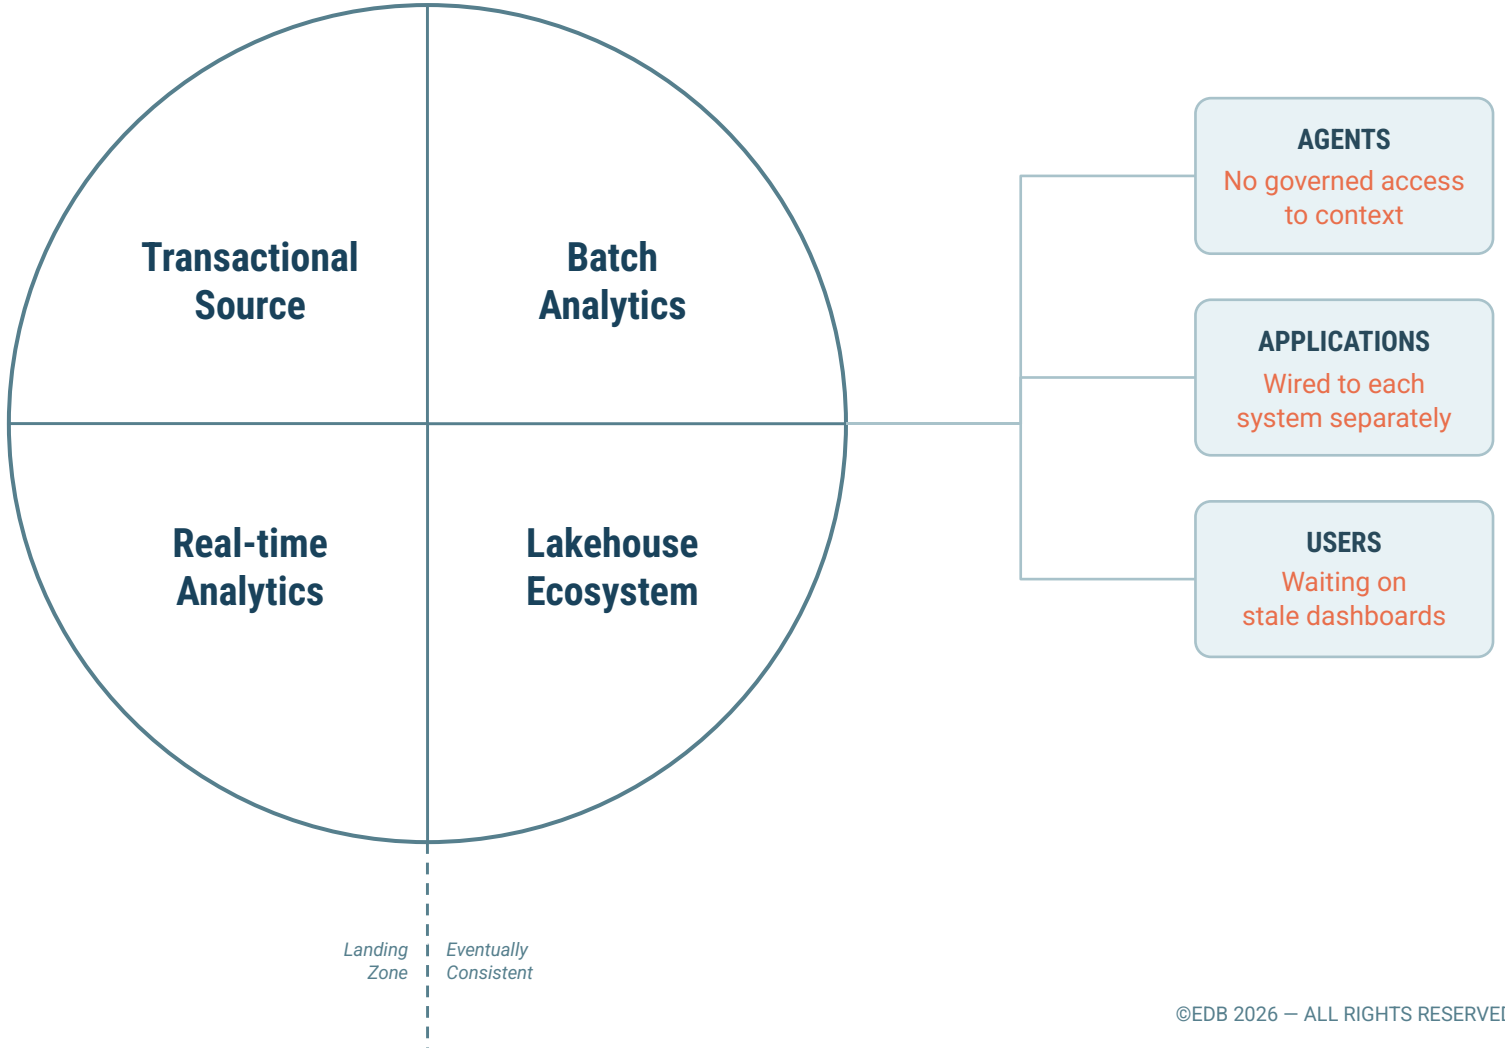

EDB POSTGRES AI

THE FULL DATA LIFECYCLE

CONVERGED ANALYTICS ON EDB POSTGRES AI

STEP 1 · TRANSACTIONAL

Postgres

With PGAA Bridge to Lakehouse

- Postgres as the transactional core
- PGAA replicates changes into Iceberg
- Open formats, shared with every engine

What's is happening?

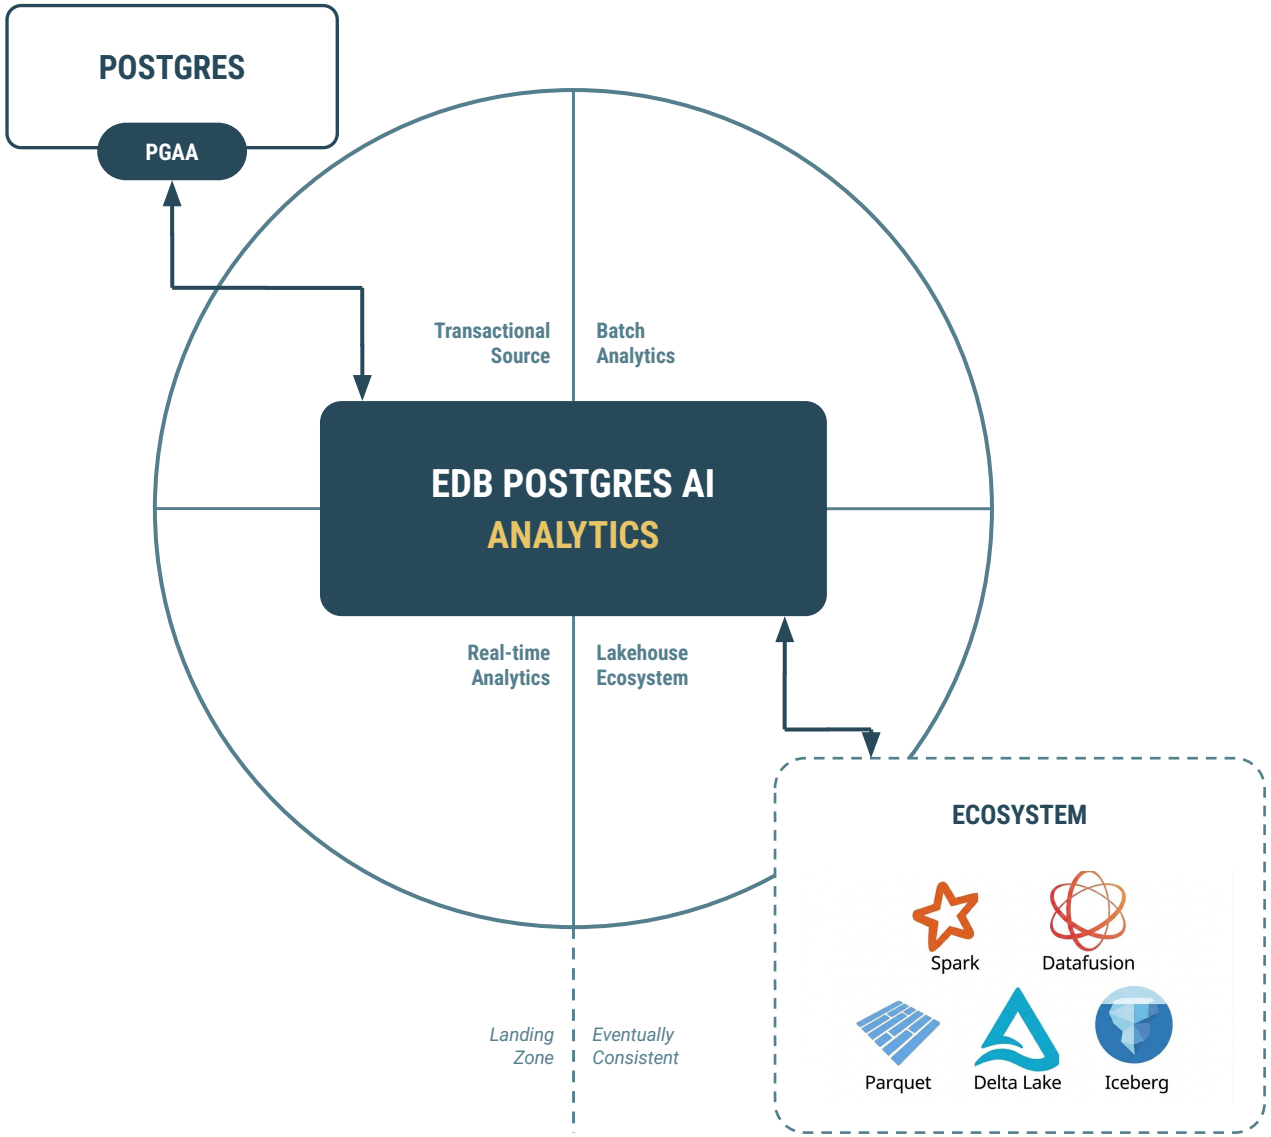

THE FULL DATA LIFECYCLE

CONVERGED ANALYTICS ON EDB POSTGRES AI

STEP 2 · BATCH

WarehousePG

Petabyte Analytics with Postgres

- Postgres-native MPP, shared-nothing scale-out
- Deep history, complex joins, in-database ML
- Reads + writes the shared Iceberg tables

What happened?

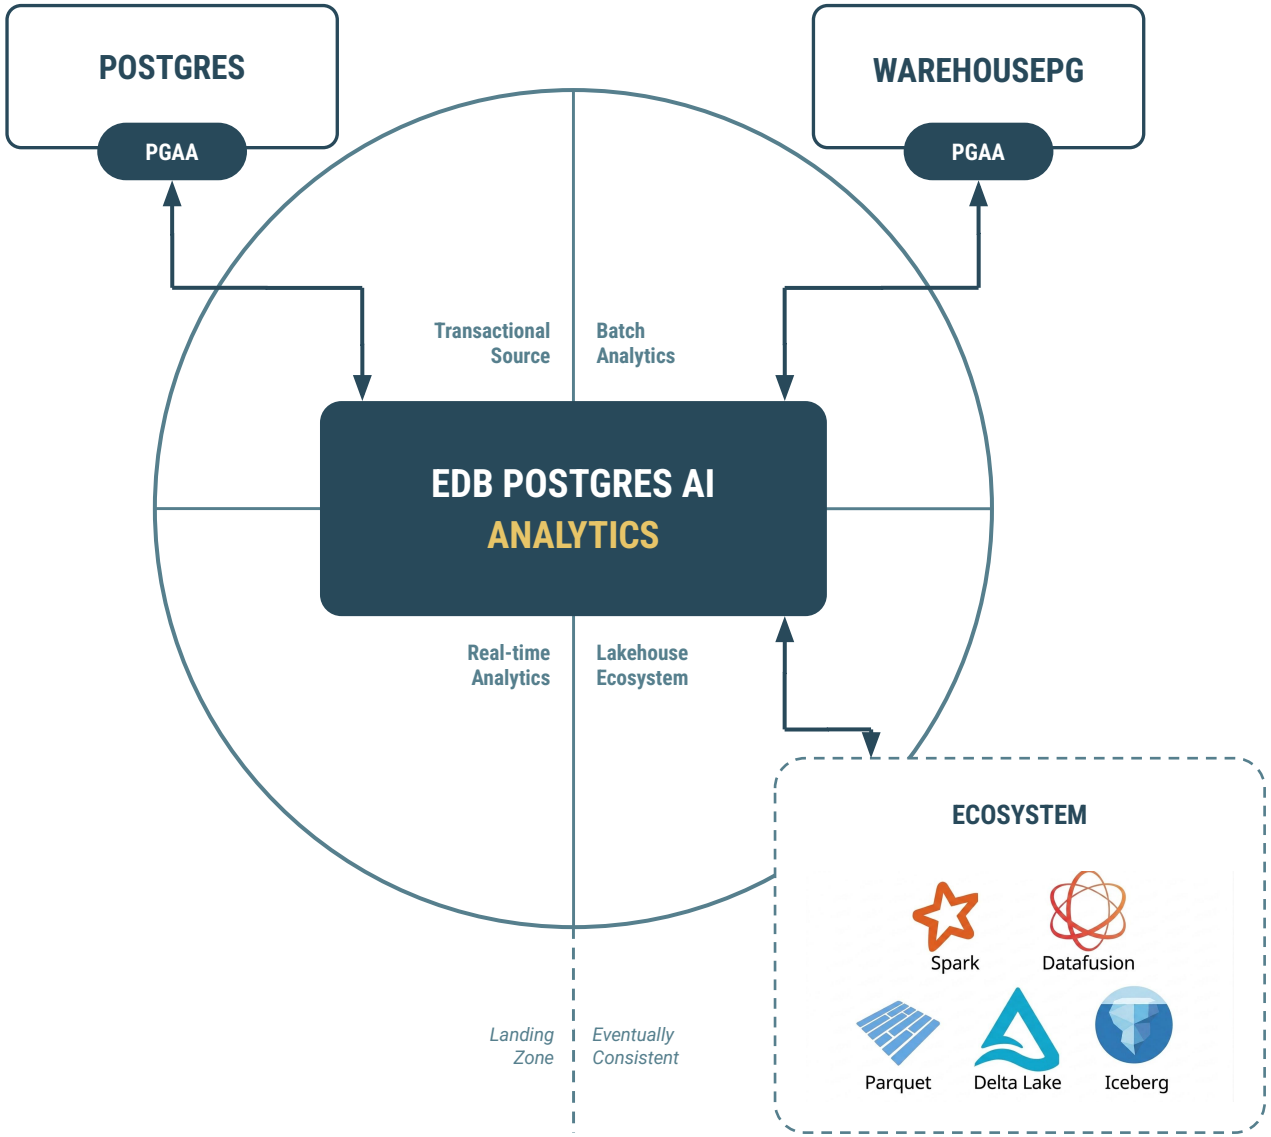

THE FULL DATA LIFECYCLE

CONVERGED ANALYTICS ON EDB POSTGRES AI

STEP 3 · REAL-TIME

ClickHouse

Sub-Second on Live Events

- Extreme-throughput columnar engine
- Sub-second reporting layer (<10 ms)
- Web/app analytics, high-volume telemetry & logs

What's happening right now?

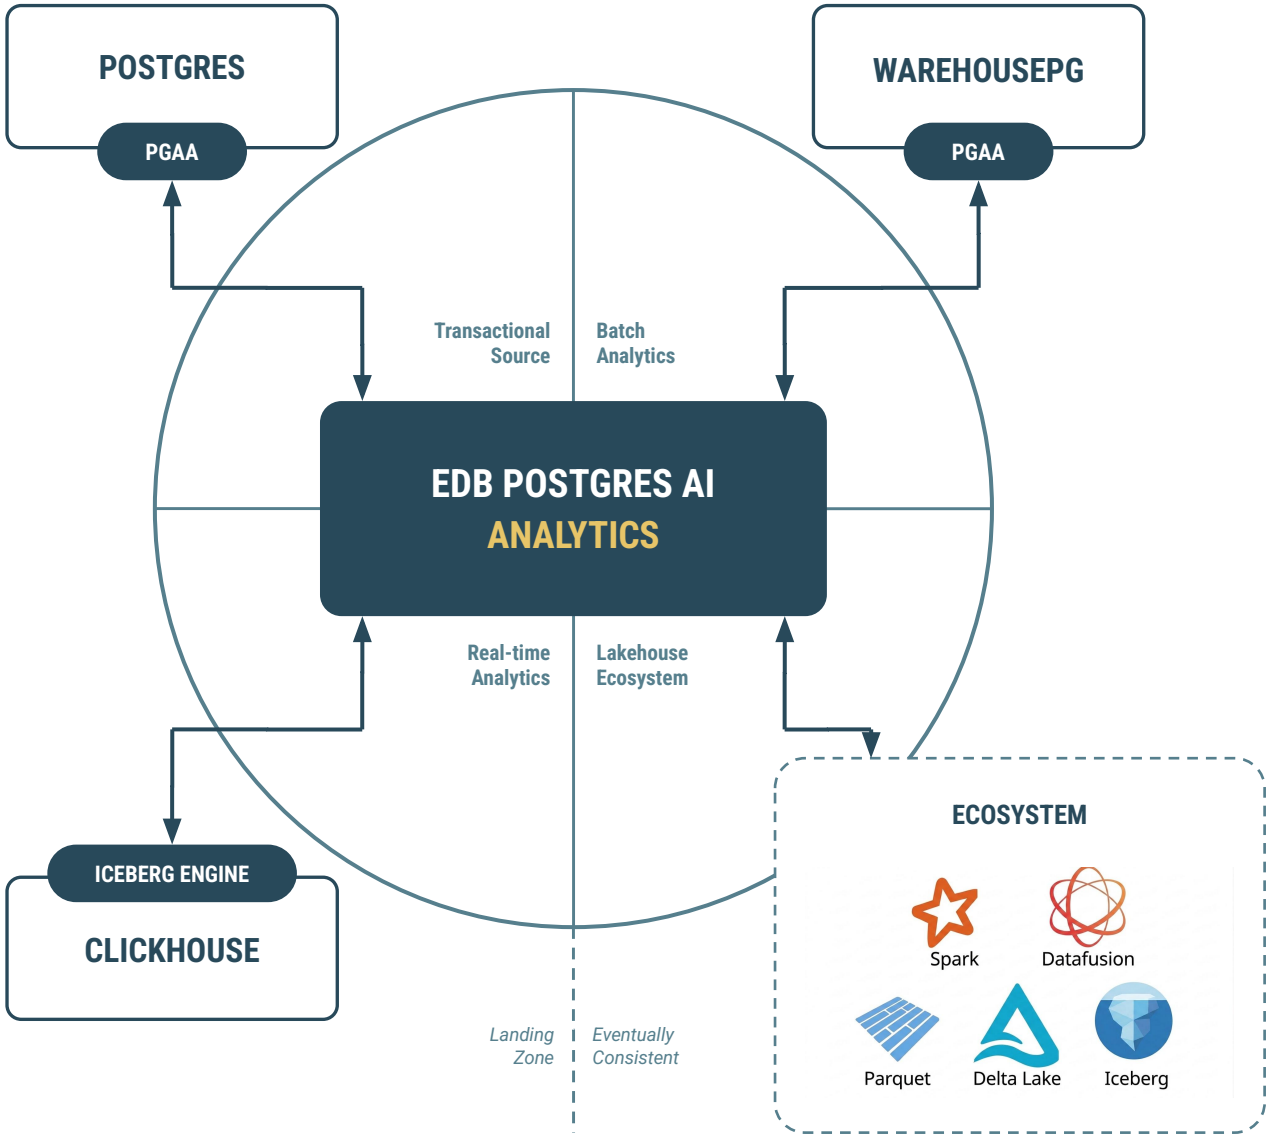

THE FULL DATA LIFECYCLE

CONVERGED ANALYTICS ON EDB POSTGRES AI

MODERNIZE & INNOVATE FASTER

OPEN SOURCE + SOVEREIGN
Own your strategy and your infrastructure

UNIFIED DATA LIFECYCLE
One catalog unifies governance, engines & agents

58% LOWER TCO
Compared to cloud data warehouses

30–100× FASTER
With optional GPU accel for heaviest workloads

ONE STRATEGIC PARTNER
#1 Postgres contributor; proven enterprise OSS

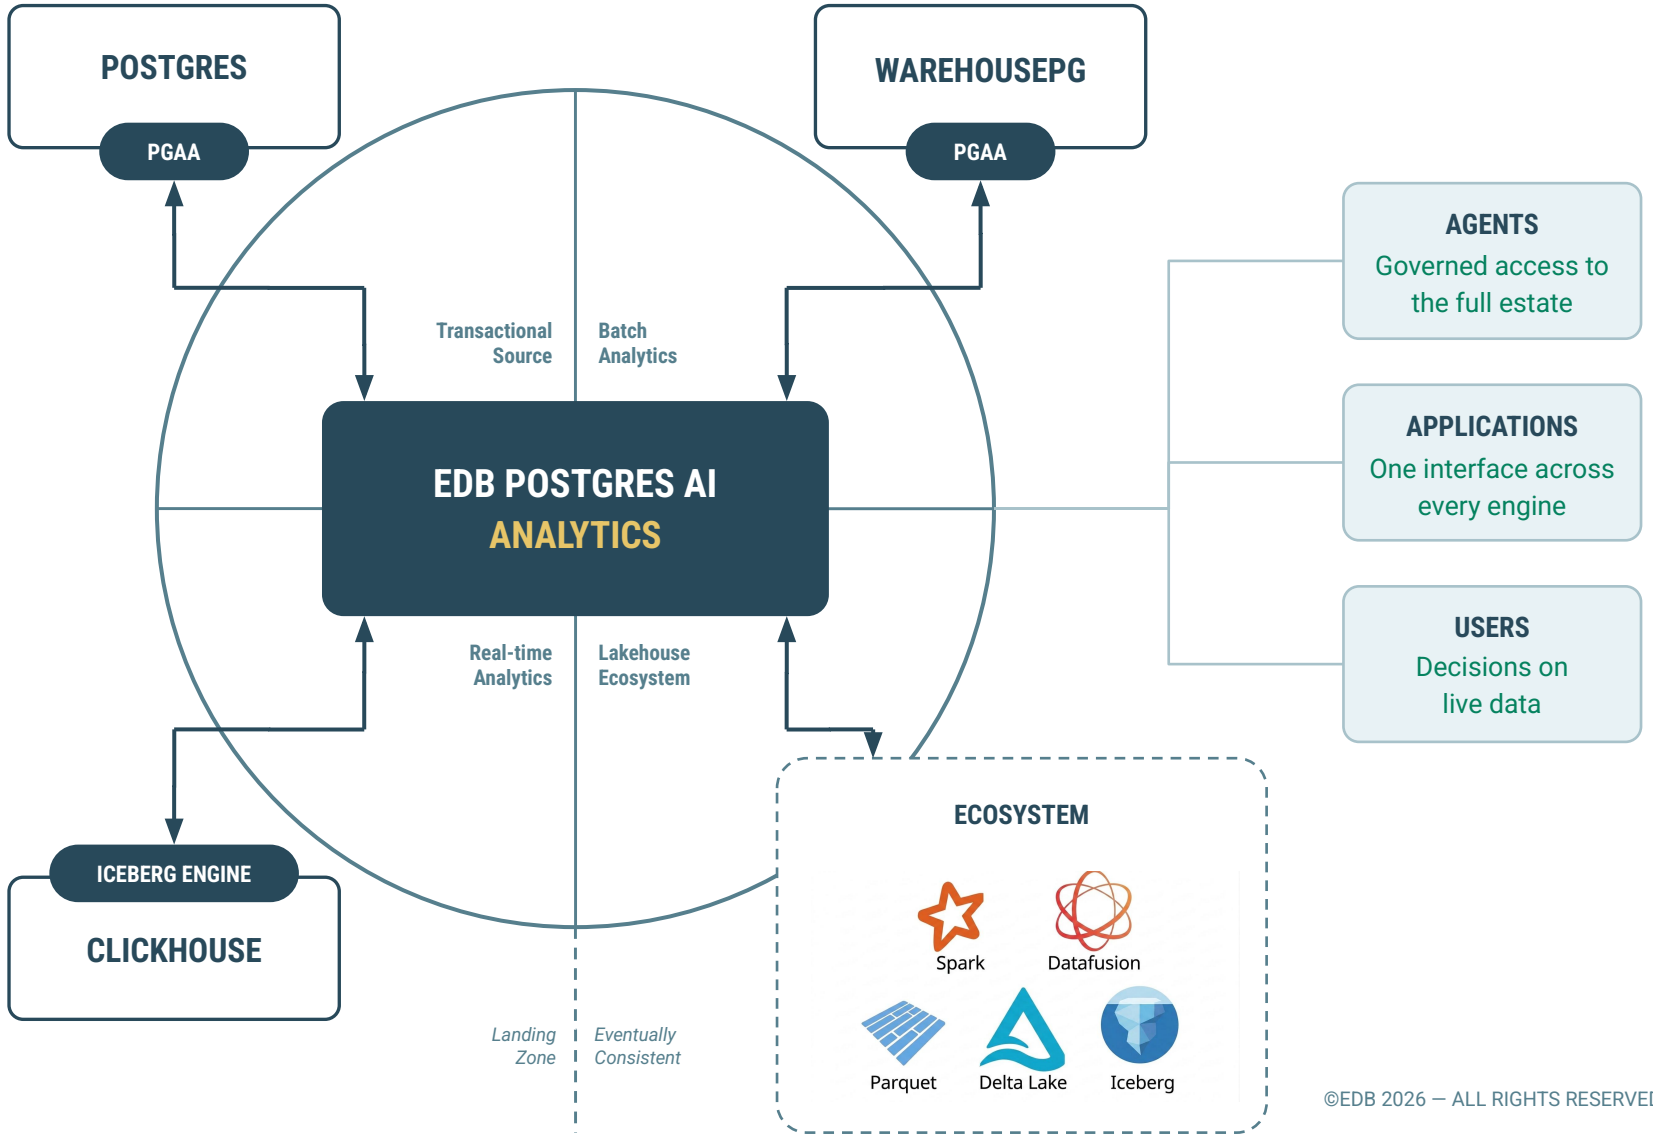

THE CATALOG REPLACES THE PIPELINE

An open Iceberg catalog is the shared source of truth

What changes

Postgres changes stream into the open catalog continuously, instead of running as nightly batch jobs.

What stays put

Object storage holds one copy of the data in an open table format; every engine reads it directly.

What goes away

Point-to-point ETL jobs, and the drift between which system has the latest data.

The catalog is the integration. Nothing has to move to be queried.

HOW THE ENGINES DIVIDE THE WORK

Postgres, WarehousePG, and ClickHouse read and write the same open tables

Postgres · PGD — Transactional source

The system of record. Changes stream into the shared catalog continuously.

WarehousePG — Batch & deep history

Postgres-native MPP for petabyte joins and in-database ML, reading and writing the same tables.

ClickHouse — Real-time serving

Sub-second queries over live events, for dashboards and high-volume telemetry.

Same tables, same catalog — every engine reads what the others write.

EDB POSTGRES AI

Demonstration

Converging vs. Standalone: The Trade-offs

THE CASE FOR CONVERGENCE

Benefits of Convergence

Operational simplicity

One team and one skill set (SQL), instead of a specialist per engine.

Predictable cost

Capacity-based pricing instead of consumption billing that scales unpredictably with usage.

Unified governance

One access-control and audit surface across every engine, instead of one per silo.

Independent benchmarking found converged, Postgres-based analytics up to 58% lower TCO than leading cloud data warehouses.

THE CASE FOR CONVERGENCE

Trade-offs of Convergence

Consistency model

Real-time analytics layers are typically eventually consistent; workloads needing hard transactional guarantees stay on Postgres · PGD directly.

Specialized tuning

Extreme-scale, single-purpose workloads can still benefit from a narrowly-tuned, dedicated engine.

Migration effort

Existing pipelines and BI connections need a staged cutover, not a rip-and-replace.

The right question isn't "converge everything." It's "which workload needs which guarantee."

MAKING IT PERFORM AT SCALE

Where the speed actually comes from

Shared-nothing scale-out

WarehousePG distributes joins and scans across segments, Postgres-native from the ground up.

In-process acceleration

An embedded analytical query engine speeds up single-node reads without leaving Postgres.

GPU offload for the heaviest jobs

Scale-out execution with GPU acceleration for the largest analytical workloads.

Benchmarked at up to 50–100× faster on the heaviest workloads, with GPU acceleration.

EDB POSTGRES AI

Converged Analytics Architecture

EDB POSTGRES AI

PGAA inside WarehousePG — Parallel Lakehouse Access

Every segment runs PGAA with DataFusion and reads Iceberg directly; Spark executors scale out one per segment.

HOW IT RUNS

- Iceberg is the shared source of truth
- Coordinator plans; segments execute
- Each segment reads Iceberg in parallel
- Optional Spark tier, one executor per segment

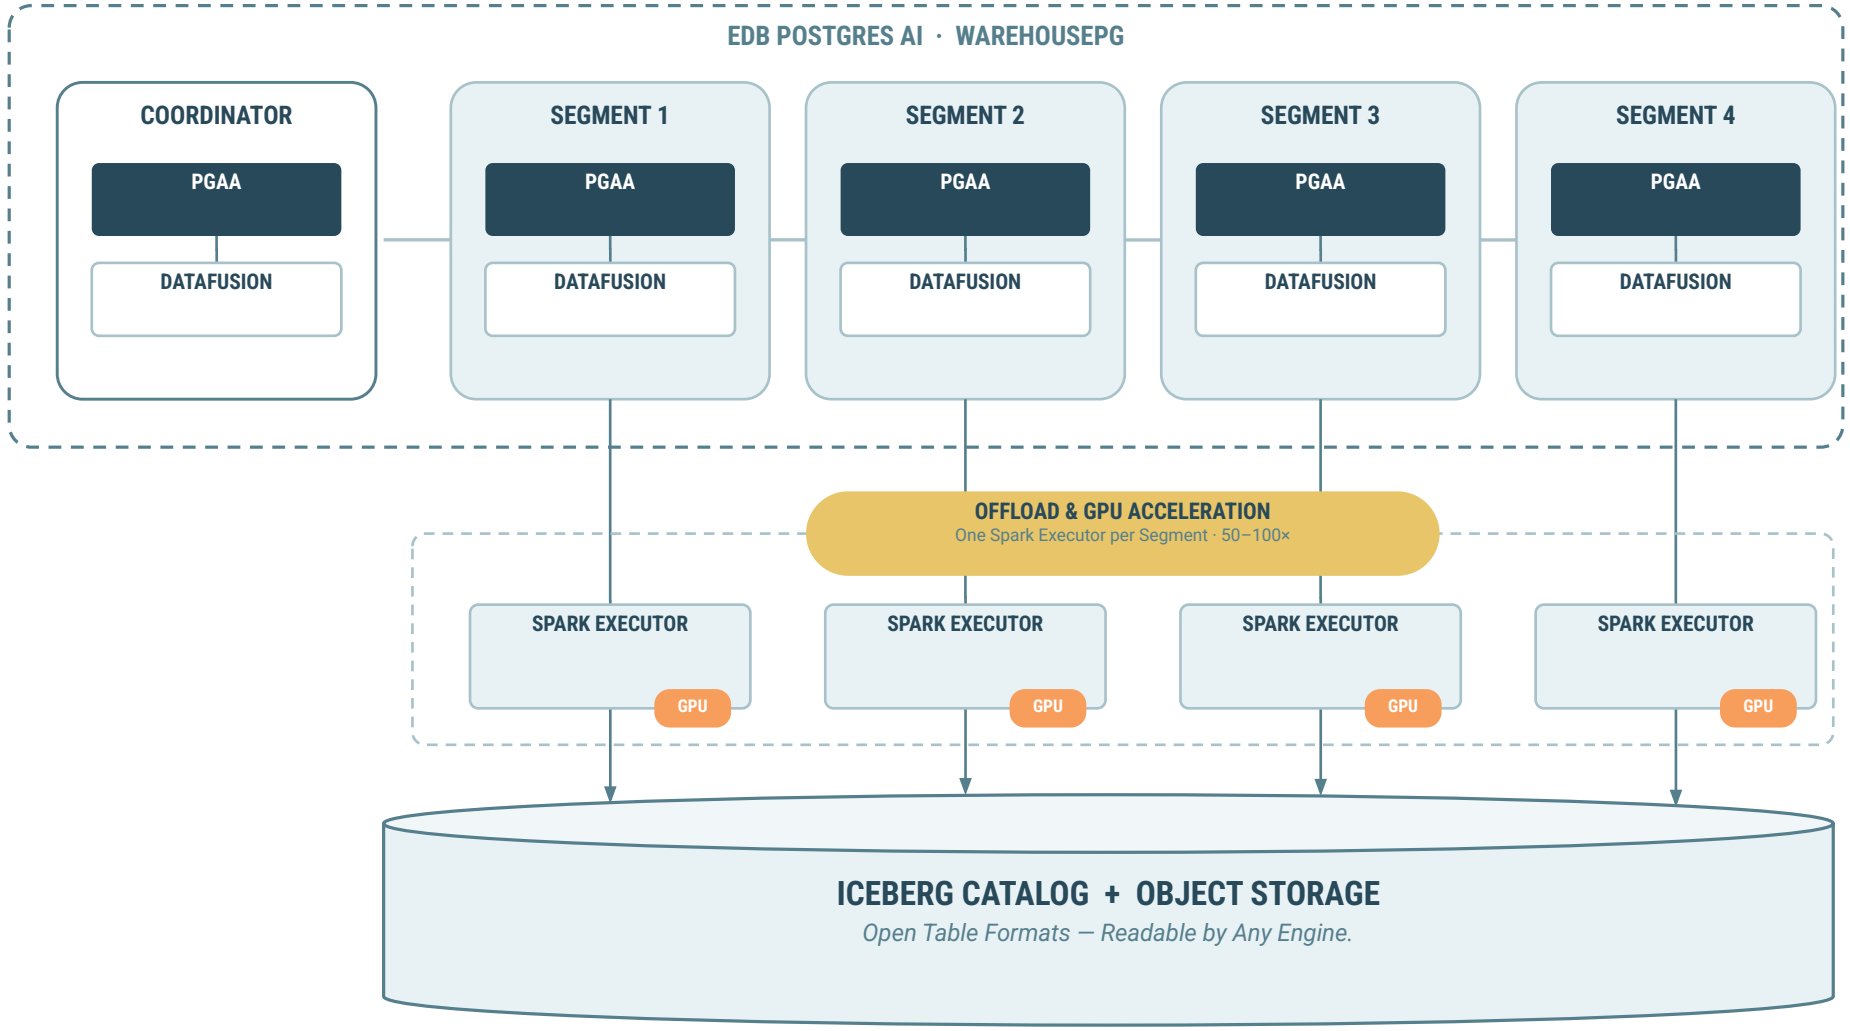

Adding the source — Postgres Replicates into the Lakehouse

Transactions stream into Iceberg; the warehouse reads them at full parallelism.

HOW IT RUNS

- Iceberg is the shared source of truth
- **Postgres replicates into Iceberg (CDC + Bulk)**
- **Postgres offloads analytics to the Spark tier**

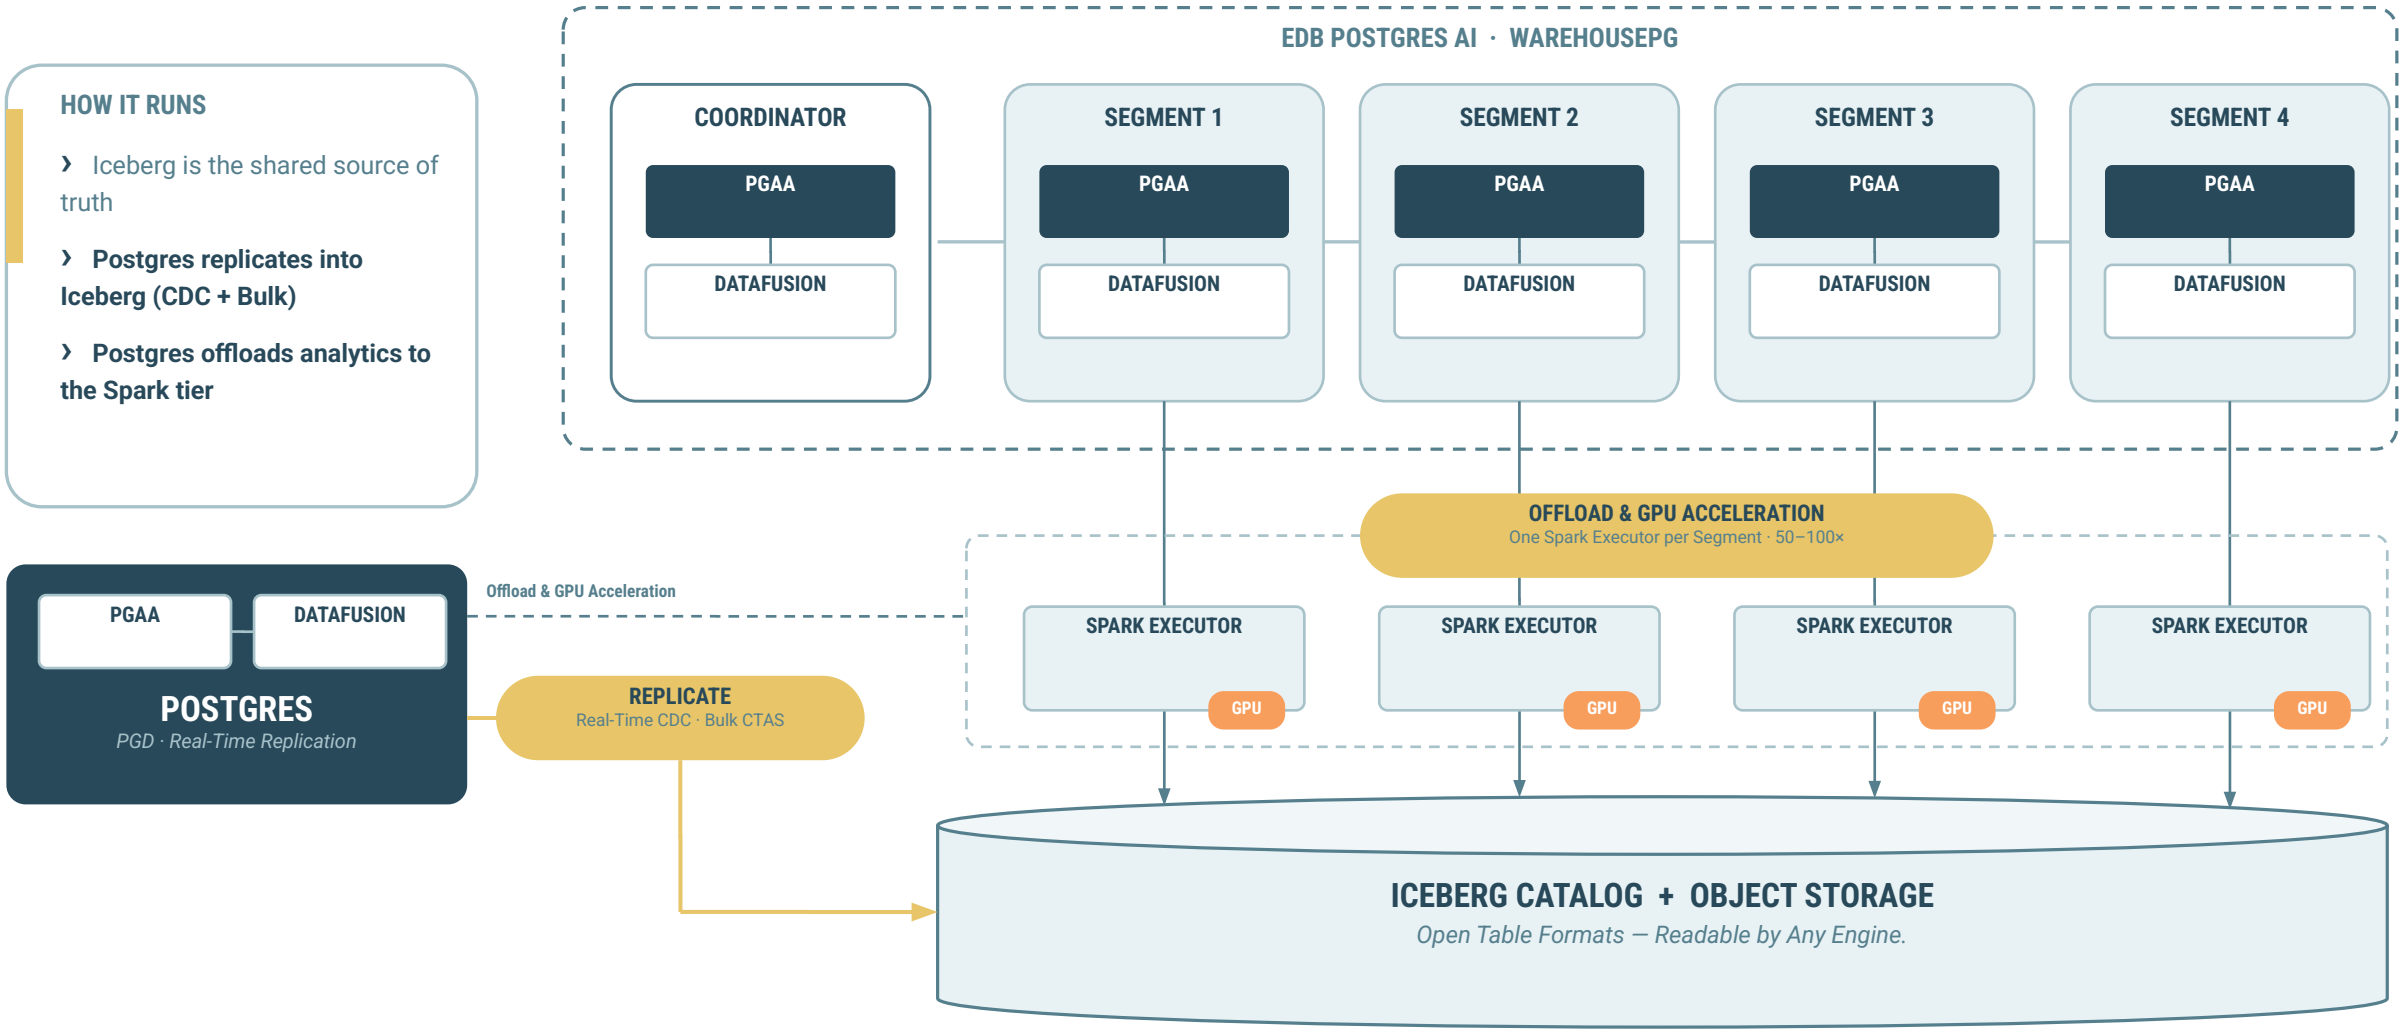

Adding real-time — ClickHouse reads the same tables

ClickHouse serves the lakehouse sub-second; everything already on screen stays put.

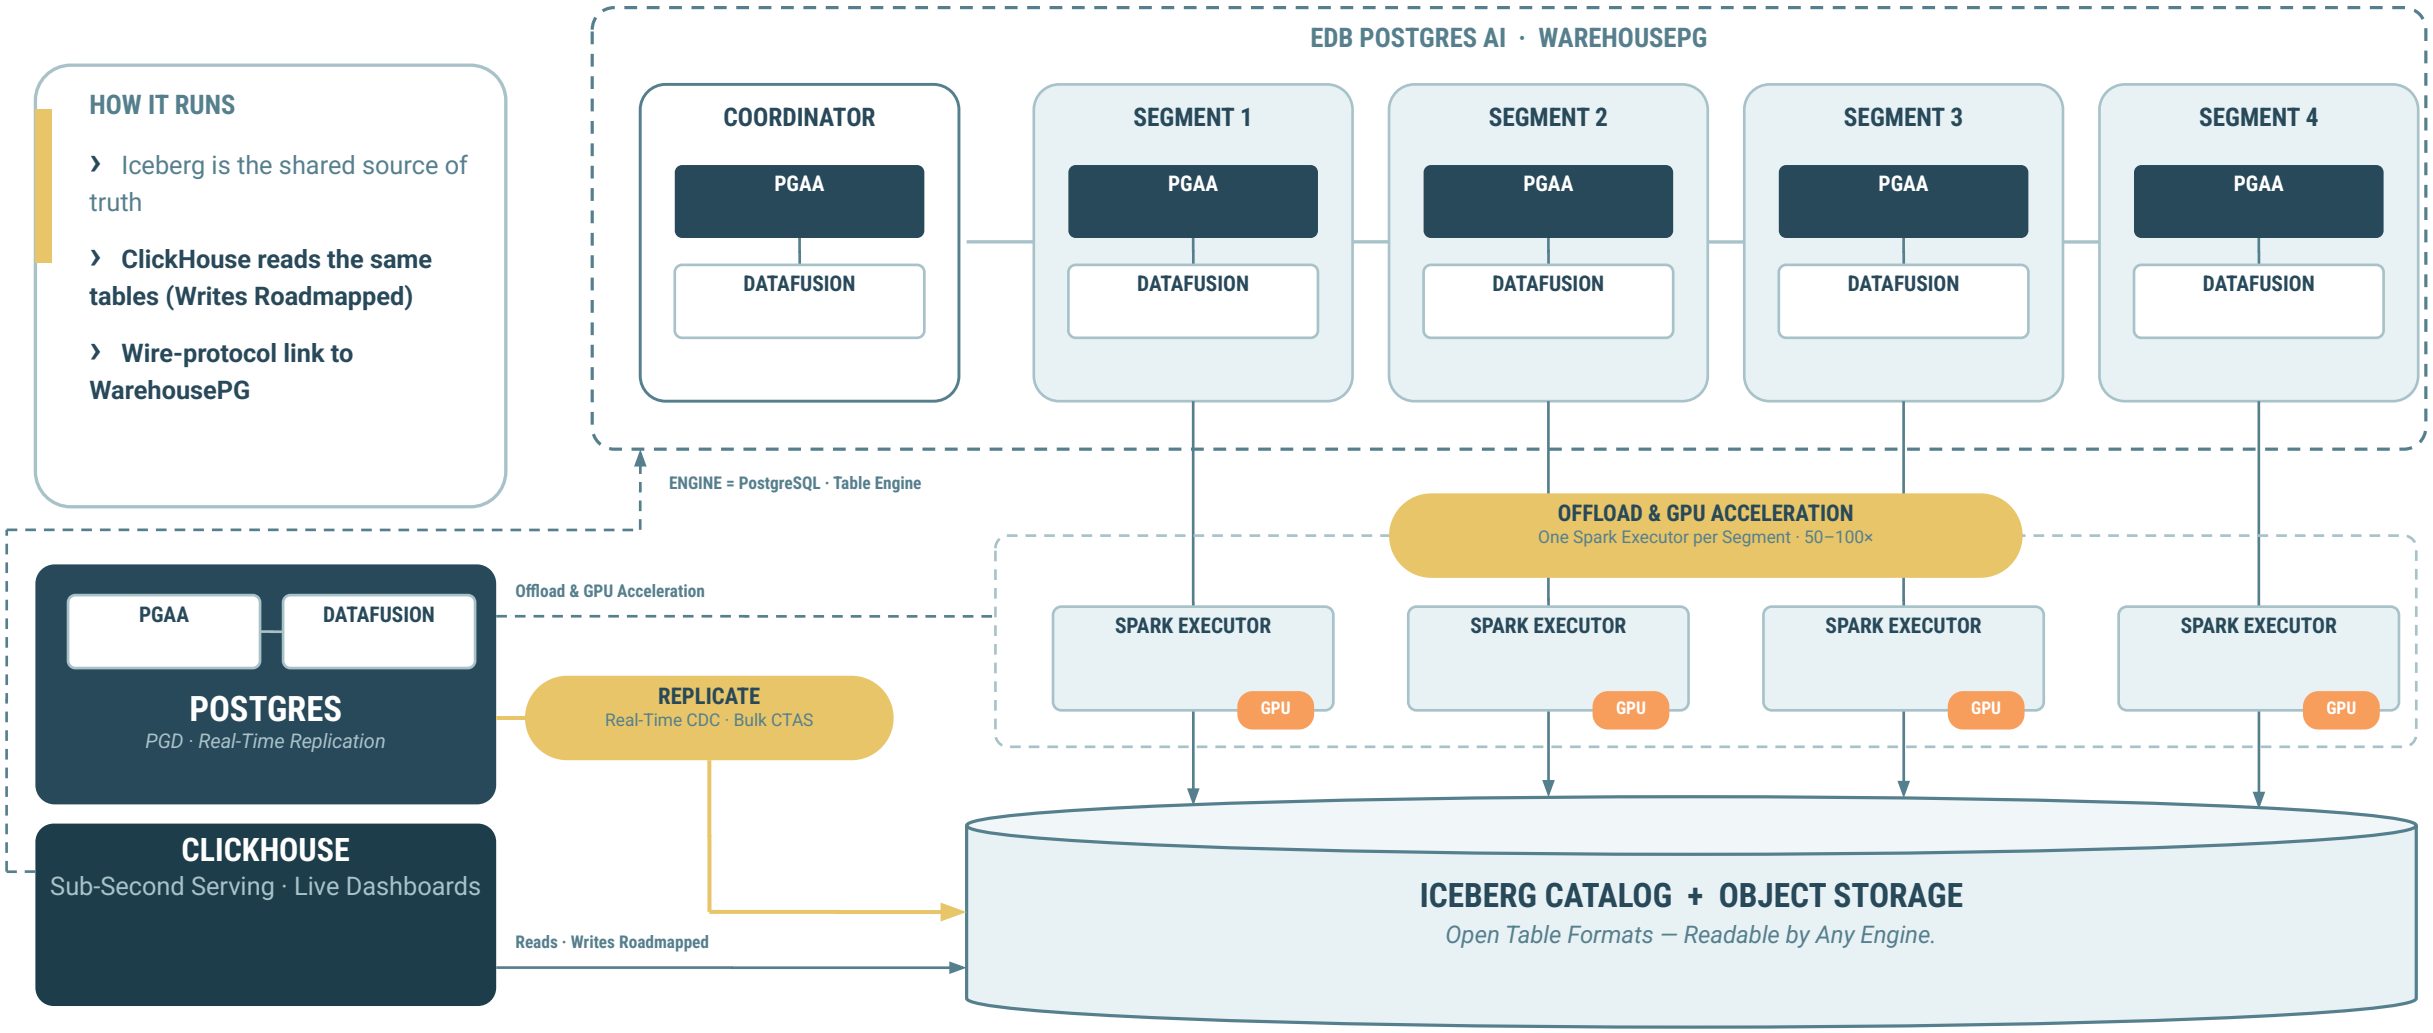

CONVERGENCE IS INEVITABLE. CONTROL IS A CHOICE.

PROPRIETARY UNIFIED STACK	OPEN FOUNDATIONS
Open format on top, closed compute	**Open engines: Postgres, WarehousePG, ClickHouse**
Runs only in the vendor's cloud	**On-prem, sovereign, or hybrid, your call**
Per-query billing that climbs with use	**Predictable per-core, up to 58% lower TCO**

Converge and lock in at once.

Converge and keep control.

EDB POSTGRES AI

Second Poll

Please answer on your screen

1- Have you worked with any of these before?

- Postgres
- Apache Iceberg
- ClickHouse
- WarehousePG / MPP data warehouses
- None of these yet

2- Which area do you have most interest in to learn more about"

- Iceberg bridge foundation layer (PGAA)
- Petabyte analytics (WarehousePG)
- Real-time/sub-second reporting (ClickHouse)
- Not ready to start yet — just researching

How to Get Started

HOW TEAMS GET STARTED

A staged path, not a migration project - Prove it on one workload, then repeat.

Start with one workload

Prove the pattern on a single reporting or analytics use case first.

Keep existing connections

Standard Postgres wire protocol means BI tools and drivers don’t change.

Expand engine by engine

Add WarehousePG or ClickHouse only where a specific workload calls for it.

Learn more and reach out to us at https://www.enterprisedb.com/products/converged-analytics

EDB POSTGRES AI

Thank you

Q&A