

The Agentic Shift

Where the market is going, what it costs to ignore it, and how EDB Postgres AI platform answers both.

CDAO, Kuala Lumpur 2026

Kuala Lumpur, Malaysia · October 6 2026

Dr Sala Muthukrishnan

Field CTO APJ — Analytics & AI

Roadmap 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 ideas and timelines 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 September 3, 2026.

We undertake no obligation to revise or update publicly any forward-looking statement for any reason.

YOUR SPEAKER

Dr Sala Muthukrishnan

Field CTO APJ – Analytics & AI

EDB · APJ

25+ YEARS in Data & Technology

EXPERIENCE HP · IBM · Oracle · VMware · DXC · EDB

FOCUS Analytics · AI · Data Platforms



Helping APJ enterprises turn Postgres into their AI data foundation

EDB POSTGRES AI

Four Movements, Forty Minutes

From market pressure to platform — with the numbers at each step.



01

AI Everywhere

The market shift, the cost pressure, and Malaysia's own numbers.



02

One Platform

The Agentic Lakehouse — transactions, analytics, and AI together.



03

Inside the Platform

Sharding, converged analytics, and the WarehousePG next announcement.



04

AI for Productivity

Faster developers, faster teams, better-served customers — and what's next.

TODAY Every claim in this session traces back to a named source or a working diagram.

PART 1 OF 4

AI Everywhere

The market shift, the cost pressure, and what it means for Malaysia.

Three Questions Every CDO Is Now Asking Me

Across 25+ years and every conversation this year, the questions have converged.



Trust

"Can I trust what the agent did with my data?"

Not whether the model is smart — whether its actions on your data are provable.



Accountability

"Who's accountable when the agent gets it wrong?"

The org chart hasn't caught up to what agents can now do.



Optics

"How do I show the board this is under control?"

Speed impresses executives. Governance is what keeps your job.

THE PATTERN Speed without governance doesn't survive a board review.

AI at Scale Requires Control — and an Open Foundation

Malaysia's AI, privacy, open-source, and cloud direction point toward the same architecture.

CONTROL & SOVEREIGNTY

Where is my data?

Who can access it?

Where does my AI run?

Where does my agent memory live?

OPEN FOUNDATION

Digital government leadership

Malaysia's MyDIGITAL blueprint drives a highly data-driven public sector.

Open-source promotion


MDEC actively drives AI and cloud adoption nationwide.

Cloud-first modernization

THE DIRECTION Open · Sovereign · Portable · Consolidated — a natural fit for EDB Postgres AI.

Every AI Conversation Re-Bills From Zero

Without persistent memory, cost climbs before ROI is ever proven.



The full chat history resends every turn.

Multiply that across every user, every rollout — and the bill compounds long before value is proven.

THE FIX Persistent agent memory — covered in Part 4 — is built to solve exactly this.

Fewer Engines, Not More

Every specialized analytics engine is another copy, another bill, another seam to secure.



OLTP Engine

A separate engine for transactions.



Data Warehouse

A separate warehouse for historical analytics.



Vector Database

A bolt-on vector database for AI.



ETL & Sync

Custom pipelines holding it all together.

4+ vendors. 4+ contracts. 4+ security models to keep in sync.

THE ALTERNATIVE One converged platform — WarehousePG, ClickHouse, and PGAA, on the same governed Iceberg store.


PART 2 OF 4

One Platform

The Agentic Lakehouse — transactions, analytics, and AI, together.

The Agentic Lakehouse, at a Glance

Agents and apps sit on top — everything below is one platform.



One governed data foundation — fewer copies and fewer integration seams, whether the consumer is a dashboard, a report, or an agent.

PART 3 OF 4

Inside the Platform

How sharding, converged analytics, and Agent Factory actually work.

Three Engines — Each Answers a Question

One integrated stack, three entry points into the conversation.



WarehousePG

“What happened — and what will happen?”

Open, petabyte-scale MPP warehouse. Per-core replacement for Snowflake and Greenplum, with in-database ML.



ClickHouse OLAP

“What's happening now?”

Sub-second analytics on billions of live telemetry events, CDRs, and logs — EDB-supported.



Analytics Accelerator

“PGAA — Postgres fronts the lakehouse.”

Query operational data and Iceberg/Parquet storage in one SQL surface, 3.5–5x faster.

ONE INTEGRATED STACK All three read and write one set of open Iceberg tables. The portfolio is the moat.

Transactional Scale: What Ships When

A clear path to horizontal transactional scale — starting with PostgreSQL 19 using EDB PGD in Q1 2027.

AVAILABLE TODAY PGD 6.4 Quorum Commit · majority agreement before commit for resilient multi-region deployments

PG19/PGD · Q1 2027

Native Sharding

Native sharded tables distribute transactional writes horizontally across nodes.



NEXT

Online Rebalancing

Move and rebalance shards live while applications remain online.

ROADMAP Sharding begins with PGD/PG19 in Q1 2027; online rebalancing follows.

PART 4 OF 4

AI for Productivity

Faster developers, faster teams, better-served customers.

ONE PLATFORM, TWO LAYERS

AIDB runs inside Postgres. AI Factory builds and serves around it, under one control plane.

AI FACTORY · IN HYBRID MANAGER

Agent Studio

Langflow-based visual builder for agents, flows and MCP servers

Pipeline Designer

No-code AIDB pipelines and knowledge bases

Model Library & Serving

Curated model images served on KServe with GPU scheduling

Governance & Observability

Audit trails, purpose-scoped access, metrics

↓
OPENAI-COMPATIBLE ENDPOINTS · SQL · MCP

AIDB · INSIDE POSTGRES

Model Hub

Local CPU models plus OpenAI, Anthropic, Gemini, NVIDIA NIM

Data Pipelines

Parse, chunk, OCR, embed and index; auto-sync knowledge bases

Agent Hub

SQL-native agents, tool catalog, delegation, MCP

Semantic KB

Schema indexed by meaning for grounded text-to-SQL

Agent Memory

Conversations, sessions and action logs kept in Postgres

Standalone or managed: AIDB installs as a Postgres extension on PG 14–18 (EPAS, Extended, community) with no dependency on Hybrid Manager. AI Factory adds the platform when you want it.

One AI Foundation. Productivity Across Every Team.

From developer velocity to enterprise search — AI works close to governed Postgres data.

BUILD · OPTIMIZE · SHIP

For developers & DBAs

- **Query & Schema Design**
Generate and validate SQL against the live schema.
- **Development Assistance**
Generate SQL, explain code, and accelerate database development.
- **Performance Tuning**
Explain slow plans and suggest changes — with human approval.

SEARCH · UNDERSTAND · ACT

For every team

- **Enterprise Search** Find years of enterprise data in plain English.
- **DBA Copilot** Explain and tune while DBAs stay in control.
- **Semantic Knowledge Base** Understand your database schema by meaning — enabling natural-language discovery and text-to-SQL.
- **Text-to-Insights** Let non-technical users ask questions without SQL.

THE POINT One governed Postgres foundation — less friction for builders, faster answers for everyone.

EDB Vector Search Extension

Production-grade vector search inside EDB Postgres AI — for PostgreSQL and WarehousePG.

edb_vectorplus

PG14–19 · EPAS · PGE · WHPG

SCALE

100M+ vectors / node

Path toward 1B+ without manual app sharding.

SPEED

>80% faster index builds

Targeting ~10× higher concurrent insert throughput.

QUALITY

>95% Recall@10

Maintain retrieval quality while QPS scales under peak writes.

EFFICIENCY

~70% lower footprint

RaBitQ quantization targets dramatically lower RAM and disk.

WHY IT'S DIFFERENT

Groups similar vectors so we search only where the answer is likely to be · compresses them to use much less memory · pgvector-compatible

TAKEAWAY Keep vectors beside enterprise data, SQL joins, permissions, governance, and AI memory — no external vector DB.

For Any questions...