



EDB Postgres® AI for WarehousePG



Executive summary

Is your data warehouse a strategic asset or a costly liability? Leading enterprises are too often hindered by proprietary, aging data platforms—such as Teradata or cloud-exclusive systems like Snowflake—with stifling vendor constraints, unpredictable costs, and limited flexibility for modernization.

EDB Postgres AI (EDB PG AI) provides WarehousePG, the open source, petabyte-scale data warehouse that empowers you to reclaim ownership of your data strategy and achieve true data sovereignty.

EDB PG AI delivers a seamless, high-performance path to petabyte-scale analytics, enabling superior agility and up to 5x total cost of ownership (TCO) reduction through a predictable, cost-controlled model—without proprietary constraints. WarehousePG transforms your data infrastructure from a cost burden into a decisive competitive advantage. With EDB, you gain ultimate control over your technology roadmap and your future, backed by the world's foremost Postgres experts.

Business challenge

Your data strategy is trapped between strict governance mandates, expensive cloud contracts, and the operational complexity of fragmented infrastructure. Analytics teams suffer the most, forced to work with a complex web of disconnected tools where analysts wait on slow ETL, data scientists fight for access to siloed data, and innovation stalls because data is rarely where it needs to be. This environment creates pervasive financial and operational friction:

- **High costs and low flexibility:** Adhering to strict global and regional mandates (GDPR, HIPAA, etc.) is impossible when cloud-only solutions force data storage outside mandated jurisdictions, exposing the business to major compliance risks.
- **Data sovereignty, security, and compliance challenges:** Cloud-only solutions force data storage outside mandated jurisdictions, exposing your business to major compliance risks.
- **AI/ML infrastructure silos:** Innovating with AI is delayed when data scientists have to move data to separate infrastructure for training, inferencing, and deployment.
- **Risky migrations from legacy platforms:** Modernization projects often fail because they require extensive refactoring, costly reskilling of entire data teams, and months of disruption.

Enterprises demand a solution that gives them the flexibility and open source economics without sacrificing performance, reliability, or sovereignty.

Solution overview

EDB PG AI redefines petabyte-scale analytics by placing sovereignty and financial control at the center with WarehousePG: an open source, petabyte-scale data warehouse that gives enterprises a high-performance escape hatch from restrictive systems. As opposed to Snowflake's cloud-only approach or Teradata's legacy lock-in, WarehousePG delivers freedom—with the advanced analytics capabilities and AI readiness customers expect, but with data they own that they can deploy anywhere, resulting in up to 5x TCO reduction.

Business value

With WarehousePG, EDB PG AI delivers the strategic infrastructure control and operational agility modern enterprises need to grow revenue and accelerate innovation.

- **Uninterrupted business growth:** Achieve up to 5x TCO improvement for your analytical database by eliminating high cloud access costs and restrictive licensing. Reinvest those savings into customer acquisition and market expansion powered by comprehensive insights—all backed by 24x7 support from the world's top Postgres experts.
- **Complete data sovereignty:** Meet strict regulatory mandates and achieve total control of your data without sacrificing performance. WarehousePG's broad data access and deployment flexibility eliminate expensive, risky data movement.
- **Accelerated time-to-value for advanced analytics and AI:** Gain faster, deeper insights with an AI-ready warehouse built for performance and ease of use with familiar SQL. WarehousePG combines in-database machine learning, vector processing, and data lake capabilities, helping your teams move from raw data to actionable intelligence instantly.
- **Seamless modernization:** Stop paying a fortune for end-of-life or proprietary platforms. WarehousePG provides a low-risk, high-SQL-parity path to modernization, letting you reclaim analytics control from rigid systems.

Customer success

Enterprise data leaders trust WarehousePG to deliver the stability, performance, and control required for their mission-critical analytics workloads.



MNTN: Achieving Enterprise Stability for Mission-Critical Ad-Tech

MNTN, a leading connected TV ad-tech platform, required a solution that guaranteed operational uptime and responsive support for mission-critical performance marketing. WarehousePG provided MNTN with the necessary petabyte-scale stability and performance, backed by EDB's 24x7 expert support—crucial for turning insights into competitive advantage.

"The performance is there, the stability is there, the support is responsive as they should be. I'm just happy that there's somebody there that can be with me in the middle of the night and I'm not, quite literally, hacking open source code trying to get the database recovered."



Kyobo Book Centre: Reclaiming Predictable Cost and Sovereignty

Kyobo Book Centre, the largest bookstore chain in South Korea, was seeking a strategic escape from unpredictable and soaring compute costs on their 50 TB cloud data warehouse. Kyobo adopted WarehousePG to establish cost control, gain superior performance, and meet strict data residency mandates.

"We have been plagued by runaway costs for querying our 50 TB cloud data warehouse. EDB Postgres AI for WarehousePG will give us a way to rein in costs with superior performance—and we can do it with total data sovereignty."



Euronext FX: Eliminating Vendor Risk and Reclaiming Open Source Control

Euronext FX, a leading pan-European market infrastructure, needed to eliminate vendor lock-in and technical debt from their existing Greenplum system. WarehousePG delivered a zero-migration binary swap that immediately provided superior enterprise support and open source control across their four global data centers. This seamless, lift-and-shift path ensured future-proof stability for their high-volume workloads.

"We're excited to be working with EDB Postgres AI. Its Support for Greenplum Workloads is helping us maintain control of where and how we deploy open source software."

Key capabilities

EDB PG AI for WarehousePG delivers the power of a modern, open source data warehouse with the enterprise reliability, migration expertise, and Postgres pedigree that only EDB can provide:

- **Massively parallel processing (MPP) architecture:** Ensure consistent scalability and long-term stability with a distributed architecture purpose-built to execute complex transformations on petabyte-scale datasets with high efficiency.
- **Zero lock-in, predictable cost model:** Maximize ROI with open source economics that deliver immediate TCO relief from proprietary software licensing fees. Plus, EDB PG AI is priced per core, eliminating the unpredictability of consumption-based cloud costs.
- **Real-time streaming ingestion (Flow Server):** Enable rapid innovation and up-to-the-second monitoring by handling high-volume event data from sources like Kafka and RabbitMQ with Flow Server.
- **In-database AI/ML and vector capabilities:** Accelerate time-to-value for advanced analytics operations like fraud detection and churn prediction by training models directly on your massive datasets—with MADlib for SQL users and robust in-database Python ML frameworks. Unlock new revenue streams by powering RAG applications and semantic search across unstructured data with the pgvector extension.
- **Hybrid storage (PXF):** Improve cost efficiency by keeping warmer data in WarehousePG while querying cold data in formats like JSON, Parquet, and AVRO directly from cost-effective external storage (like Amazon S3 or HDFS) using the Platform Extension Framework (PXF).
- **Seamless modernization:** Evolve your data infrastructure without disruption. Modernize in hours, not months, with a simple binary swap for Greenplum workloads and high SQL parity for other legacy systems.
- **Deployment flexibility:** Meet sovereign data mandates and reduce vendor lock-in by deploying and operating across any environment—cloud, on-prem, or hybrid—for full data residency compliance.

By adopting EDB PG AI for WarehousePG, you gain full control over your data strategy, backed by global 24x7 support from EDB experts.

Technical architecture

WarehousePG is built upon the standard Postgres engine but achieves petabyte-scale analytics by utilizing an MPP architecture. This design allows WarehousePG to scale horizontally by distributing data and queries across many Postgres nodes (segments) that work together in parallel, eliminating the performance limits of a single server. The workload is managed by a single Coordinator node, which optimizes and distributes the tasks to individual Segments. EDB PG AI for WarehousePG offers deployment flexibility for full data sovereignty and supports hybrid storage by using the PXF to query external data lakes like S3 or HDFS using standard SQL, eliminating complex ETL pipelines. Furthermore, the architecture is AI-ready, supporting in-database machine learning (via the MADlib extension and Python ML frameworks) and native vector capabilities (via the pgvector extension) to enable advanced analytics directly where data resides. For real-time data needs, the dedicated Flow Server component handles high-volume streaming ingestion from sources like Kafka and RabbitMQ. The system ensures enterprise control with predictable workload management and provides high availability through a Standby Coordinator and Mirror Segments.



Disclaimer: Greenplum® is a registered trademark of Broadcom Inc. EDB and EDB Postgres AI are not affiliated with, endorsed by, or sponsored by Broadcom Inc. Any references to Greenplum are for comparative, educational, and interoperability purposes only.

EDB Postgres AI is the first open, enterprise-grade sovereign data and AI platform, with a secure, compliant, and fully scalable environment, on premises and across clouds. Supported by a global partner network, EDB Postgres AI unifies transactional, analytical, and AI workloads, enabling organizations to operationalize their data and LLMs where, when, and how they need them. For more information, visit www.enterprisedb.com

© EnterpriseDB Corporation 2025. All rights reserved.