



SOVEREIGN AI

**Your controlled, adaptable
AI platform.** Secure, flexible,
and cost-effective where
your data lives.



Sovereign AI: Your controlled, adaptable AI platform. Secure, flexible, and cost-effective where your data lives.

THE CHALLENGES

Technology leaders face a complex and expansive landscape of AI tools as they race to operationalize generative AI (GenAI) in a rapidly evolving market. The pressure to gain a competitive edge is intense — but this ambition is tempered by the need to build robust, scalable AI platforms that can adapt to future business transformations. While some tools enable quick proofs of concept, they often fall short of enterprise-grade requirements. The challenge lies in balancing speed-to-market with a methodical approach that recognizes AI as, fundamentally, a data storage and management problem. This is particularly crucial in sectors like finance, healthcare, and government, where data integrity and regulatory compliance are paramount. Many businesses are wary of exposing sensitive information to cloud-based large language models — and instead are seeking hybrid or private AI solutions that offer flexibility without compromising security. Ultimately, the goal is to create a foundation that not only meets current needs but also supports on-demand AI applications as businesses evolve, without repeated architecture overhauls for each new use case.

- **Compliant, well-governed AI solutions:** Ensuring AI deployments maintain data integrity and regulatory compliance are top priorities for 60% of data leaders.
- **Maintaining security and control:** Addressing CEOs' reluctance to expose sensitive data to cloud-based LLMs but bring their data to a hybrid, private solution.
- **Velocity-vigilance tradeoff:** Balancing AI integration with the need for secure, stable, and observable data in business operations.

THE SOLUTION

Sovereign AI, powered by AI Accelerator from EDB Postgres® AI, enables GenAI inferencing while keeping all data in a controlled, private environment — ensuring well-governed AI models on your organization's own terms. By running preferred models alongside existing data in Postgres, sovereign AI eliminates data movement and enhances security through robust encryption, logging, and auditing. The solution adapts to your existing infrastructure with deployment options across on-premises, private cloud, or hybrid environments. Plus, sovereign AI helps with cost optimization by eliminating cloud data access fees and enabling locally hosted models — empowering you to leverage AI securely and efficiently within your organization's own environment.

- **Private AI environment:** Avoid exposing sensitive data to public cloud-based models. Instead, bring a wide range of models close to well-governed data in Postgres — enabling vector search and a wide variety of powerful GenAI inference applications with comprehensive observability and audit logging.
- **Deployment options:** Deploy sovereign AI anywhere across private infrastructure — in private data centers, VPC enclosures in the cloud, or hybrid.
- **Access controls and encryption:** Protect business data with features like role-based access control (RBAC), row-level security (RLS), and transparent data encryption (TDE).
- **Locally hosted open source models:** Open source AI models sit right alongside business data unified in a private Postgres environment, eliminating vendor lock-in and cloud access fees.

KEY RESOURCES

- **Related Products and Solutions**
 - [EDB Postgres AI »](#)
 - [EDB AI Accelerator »](#)
 - [Use Case: Cognitive AI »](#)
 - [Use Case: Virtual Expert »](#)
 - [EDB Postgres Distributed »](#)
 - [EDB Postgres Extended Server »](#)
 - [EDB Postgres Advanced Server »](#)
- **Blogs**
 - [More Personalized Service, Less Waiting: How AI is Transforming Insurance Contact Centers »](#)
 - [Using Postgres for Analytics, Search, and AI »](#)
 - [The Expanding World of AI and Postgres »](#)
- **Whitepapers**
 - [Intelligent Data: Unleashing AI with PostgreSQL »](#)
 - [Security Best Practices for Postgres »](#)
- **Demos**
 - [Bring AI Models to Your Postgres Data »](#)
- **Webinars**
 - [Operationalizing AI for Postgres »](#)

THE BENEFITS

EDB Postgres AI empowers businesses with full control, enhanced flexibility, optimized costs, and advanced security — enabling cutting-edge AI solutions while maintaining strong data governance and deploying AI strategies with maximum efficiency across private or hybrid environments.

- **Full control:** Build cutting-edge GenAI inference applications while retaining strong data governance practices with all data and models retained in a controlled, highly observable Postgres environment.
- **Greater flexibility:** Comprehensive private deployment options offer your business the versatility to pivot AI strategies in response to market, regulatory, or operational changes.
- **Enhanced data security and compliance:** Protect sensitive data while pursuing advanced GenAI functionality, ensuring regulatory compliance and minimizing security risks.
- **Cost optimization:** Reduce total cost of ownership (TCO) by voiding unpredictable cloud fees and optimizing AI spending with storage formats that are 18x cost-effective versus standard Postgres — unified in a single data platform built on open source technologies.

THE OUTCOMES

- Leverage strong data control, end-to-end observability, and enhanced security to reduce the risk of data breaches and compliance violations, protecting company reputation and avoiding costly penalties.
- Rapidly deploy AI solutions using proprietary data, accelerating time-to-market. Swap models effortlessly, adapt to business needs instantly, and implement GenAI across business units.
- Maximize ROI on AI with end-to-end control over AI applications. You can fine-tune costs, scalability, and deployment to ensure optimal cost-performance.

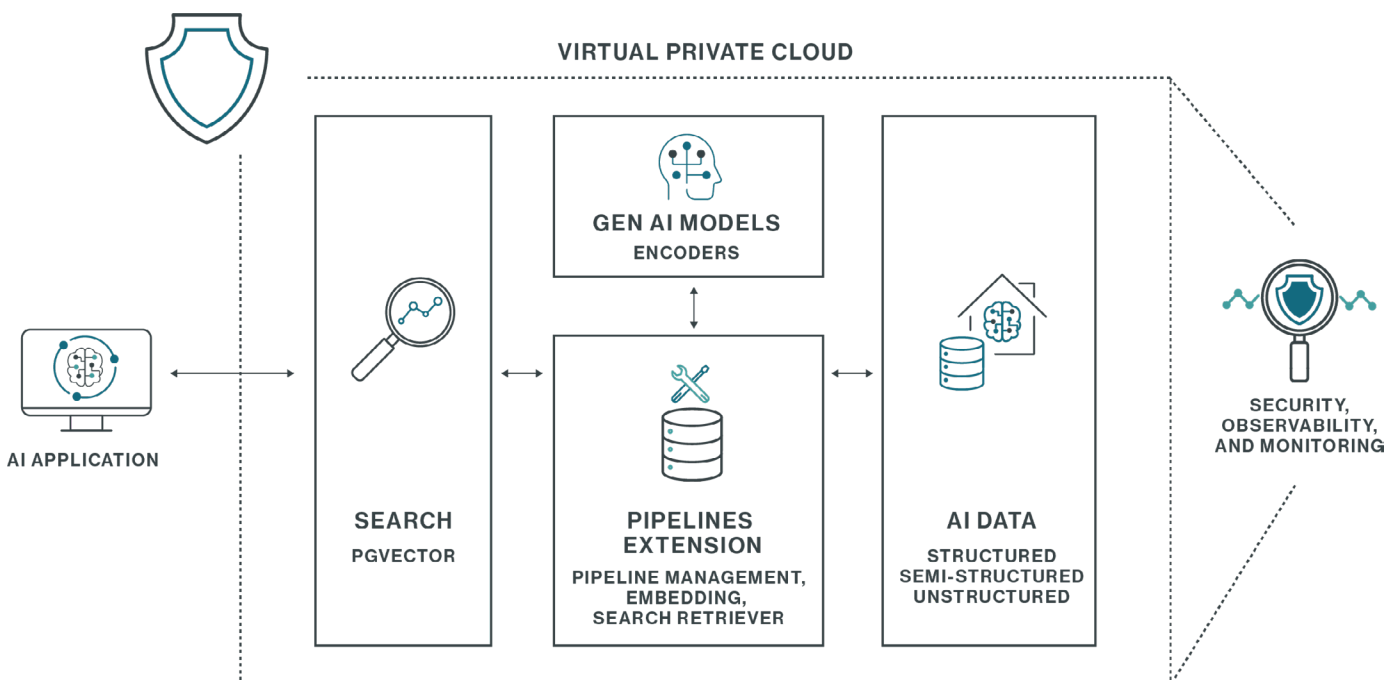


Figure 1. EDB Postgres AI enables sovereign AI: Maintain control of data in a private environment – ensuring well-governed AI models on the organization's own terms. Build GenAI inference applications faster, with fewer lines of code.

FREQUENTLY ASKED QUESTIONS

Q: What is sovereign AI?

A: Sovereign AI, powered by AI Accelerator from EDB Postgres AI, is a private AI solution that allows businesses to maintain full control over their data and AI operations by bringing GenAI models to a secure Postgres environment — providing flexible deployment options, enhanced security, and cost efficiency without relying on cloud-based services.

Q: What is AI Accelerator from EDB Postgres AI?

A: AI Accelerator provides the fastest way to test and launch enterprise GenAI inference applications with the powerful EDB Pipelines extension that comes preloaded with pgvector — providing automated vector embedding, storage, and retrieval workflows directly in Postgres. This allows developers to achieve complex GenAI functionality with minimal code. With swappable configurations, businesses can easily adapt to various needs by switching models and integrating multiple data modalities, while automated pipelines ensure always-up-to-date data for efficient query and retrieval.

Q: How does sovereign AI address data security concerns?

A: Sovereign AI uses a model-to-data approach, enabling GenAI inference from well-governed data in Postgres without exposing sensitive information to cloud-based models. EDB Postgres AI also comes with robust security features like RBAC and encryption at rest and in transit. This eliminates data movement and ensures compliance with data privacy regulations.

Q: What are the key features of sovereign AI?

A: Key features include:

- A private AI environment where Gen AI models run close to well-governed data in Postgres with full observability and audit logging.
- Deployment options across on-prem, private cloud, or hybrid infrastructure to suit your business needs.
- Access controls and encryption features like RBAC, RLS, TDE, and more.
- Locally hosted open source models that eliminate vendor lock-in and cloud fees.

Q: Can sovereign AI work with existing data infrastructure?

A: YEDB Postgres AI's integration with open lakehouse tables in object storage allows for seamless incorporation of data into GenAI applications without requiring multiple databases or performance-degrading data pipelines.

Q: What benefits does sovereign AI offer to businesses?

A: Sovereign AI offers:

- Full control of data for stronger governance.
- Greater flexibility to meet shifting market, regulatory, and operational demands.
- Enhanced security and compliance for minimized risk.
- Cost optimization via 5x cost-efficient formats and eliminating vendor lock-in.

Q: Is sovereign AI suitable for specific industries?

A: While beneficial for many industries, sovereign AI is particularly valuable in sectors like finance, healthcare, and government, where data security and compliance are critical and heavily regulated.

Q: How can sovereign AI accelerate time-to-market for AI solutions?

A: EDB Postgres AI provides a unified sovereign AI platform for storing, managing, and interacting with your data, serving as an ideal AI base layer that won't require future changes. This approach keeps AI simple and under your control. It enables rapid deployment of customizable AI by allowing businesses to swap models with minimal code, instantly adapt to changing business needs, and implement AI across different business units seamlessly.

Q: How does sovereign AI help optimize costs while scaling AI operations?

A: With full control over deployment and resource management, sovereign AI reduces costs by avoiding cloud data access fees, optimizing storage, and enabling scalable AI solutions — ensuring maximum return on investment.

Q: Can sovereign AI help balance AI integration with data security?

A: Yes, sovereign AI addresses the velocity-vigilance tradeoff by enabling rapid AI integration while maintaining secure, stable, and observable data in business operations.