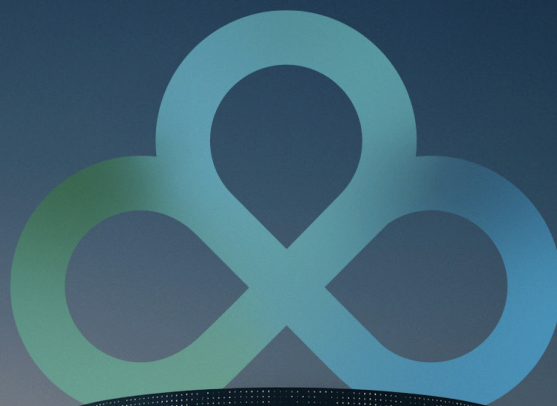




EDB PostgreSQL® AI: AI Factory

**The fastest way to securely build, test,
and launch sovereign AI applications**



The challenge: The GenAI integration complexity trap slows time to value

Organizations face mounting pressure to accelerate GenAI initiatives, which drives them to piece together off-the-shelf AI components. However, this creates data silos, governance issues, and security risks as sensitive data flows to numerous external vendors. This fragmented approach is also technically complex, requiring scarce AI developer talent to spend 50%–80% of their time on integration and driving 6–12 month timelines for building, testing, and deployment. This leaves business users left waiting for AI value and ultimately forces many companies to abandon GenAI projects before they reach production.

The solution: From concept to production in days

EDB Postgres AI Factory delivers a secure, unified GenAI platform that eliminates implementation complexity and accelerates time to value. By integrating vector capabilities, automated AI data preparation, and model serving with an application builder in a single, cohesive system, organizations can deploy production-ready GenAI inferencing applications in days instead of months. Flexible interface options include a low-code environment for developers to craft sophisticated GenAI and agentic applications, alongside a no-code option for business users to meet domain-specific needs. This comprehensive approach maintains complete data sovereignty and governance within your trusted Postgres environment while providing the flexibility to adapt as technology evolves.

- **Simplified architecture:** EDB Postgres AI provides a single GenAI solution that reins in data sprawl and reduces implementation cycles from months to days, allowing your teams to focus on business innovation rather than technical plumbing.
- **One platform for all users:** With both low-code and no-code interfaces, AI Factory empowers developers and business users alike to build GenAI inferencing applications without resource bottlenecks.
- **Enhanced data security:** By keeping all data in a single environment rather than sending it to external vendors, AI Factory enables you to confidently leverage sensitive data with complete data sovereignty and governance.

AI Factory: A single GenAI inferencing and agentic platform

Accelerate time to market and streamline your entire AI workflow with a turnkey GenAI application builder that boosts productivity. Leverage your existing data while maintaining complete data sovereignty and governance. Confidently scale AI initiatives across departments without GenAI experts or complex vendor relationships.

- **GenAI Builder:** Become an AI-first enterprise by empowering every department to build custom GenAI with a robust low-code SDK for developers and a point-and-click interface for business users. Ensure security and accuracy with custom knowledge bases, rules, and sovereign deployment, protecting sensitive data.
- **Agent Studio:** Enhance productivity and decision-making with AI agents that perform tasks autonomously. Start with open source agent templates or build custom agent workflows tailored to your business needs.
- **AI Pipeline:** Accelerate GenAI time-to-value with set-and-forget AI data management. Use just five lines of code to set up an AI Pipeline that automatically syncs embeddings with source data, ensuring always-up-to-date AI knowledge bases without costly infrastructure maintenance.
- **Vector Engine:** Bring together AI and business data in a single, secure location in Postgres, eliminating data movement to external vendors. Gain complete data sovereignty with rapid semantic search powered by open source pgvector and an intelligent retriever that handles vector similarity calculations for you.
- **Model Serving:** Eliminate vendor lock-in with flexible on-premises model deployment. Seamlessly swap between models as business needs evolve, without infrastructure changes. Maximize ROI on hardware such as NVIDIA and Supermicro with intelligent scaling that optimizes resources automatically.

Key benefits

- **Accelerate time to market:** Transform GenAI projects from concept to production in days instead of months with a complete end-to-end platform. Start building right away with an intuitive application builder, and use just five lines of code to create an AI pipeline for always-up-to-date knowledge bases, eliminating complex integration work.
- **Equip every team to build custom, sovereign AI:** Enable every skill level and department to create sovereign GenAI solutions that deliver immediate business impact. Equip technical teams with a Python SDK for complex projects, while business users build powerful applications through an intuitive no-code interface.
- **Design only once with secure, automated data management:** Confidently pursue scalable AI initiatives with sovereign data control and governance in secure-by-default Postgres. Adapt seamlessly to new use cases with reusable architecture, automated AI data pipelines, and swappable models.
- **Deliver more value from your data and systems with agentic AI:** Turn data into action with intelligent AI agents – autonomous systems that unlock new insights and efficiencies by solving problems, performing repetitive tasks, and analyzing vast quantities of data. Leverage multi-agent workflows to bring end-of-life systems in line with modern infrastructure, reclaiming value from past investments.

Key use cases

- **Sovereign AI:** Your controlled, adaptable AI platform. Secure, flexible, and cost-effective where your data lives. [Learn More »](#)
- **Cognitive AI:** Create GenAI applications that mimic human understanding by seamlessly combining text, images, and other data modalities. [Learn More »](#)
- **Virtual Expert:** Build custom chatbots, copilots, translators, and more with automated GenAI data orchestration in enterprise-hardened Postgres—all using an intuitive, familiar interface. [Learn More »](#)

Hear from our customers

“EDB Postgres AI Factory took off like wildfire within EDB. What used to take us a year, we now deliver in weeks. It’s fundamentally transformed how we bring GenAI to life—faster, more securely, and at scale.”

– Dan Merzlyak, SVP, Global Head of Data, Analytics, and AI at EDB

Real-world results

“McKnight Consulting Group compared the integrated EDB Postgres AI platform against a usual DIY approach using AWS. Across all the eight critical components of an enterprise AI factory, EDB Postgres AI reduced overall complexity by 67%, delivered a three-times faster design-to-delivery capability (from seven months to nine weeks), and reduced maintenance complexity and costs by 38%.”

– McKnight Consulting Group

Learn more

[Web page:](#) AI Factory

[Web page:](#) GenAI Builder

[Web page:](#) Agent Studio

[Web page:](#) GenAI Orchestration

[On-demand webinar:](#) Beyond Bots: How AI Can Transform

Contact Centers in the Financial Services Industry

[On-demand webinar:](#) Sovereign AI: Achieve Generative AI Innovation with Postgres Control, Security, and Speed

[Blog:](#) Building a Future-Proof AI Foundation with EDB Postgres AI

About EDB Postgres AI

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 it.