


TECHNICAL BRIEF

Agentic Database

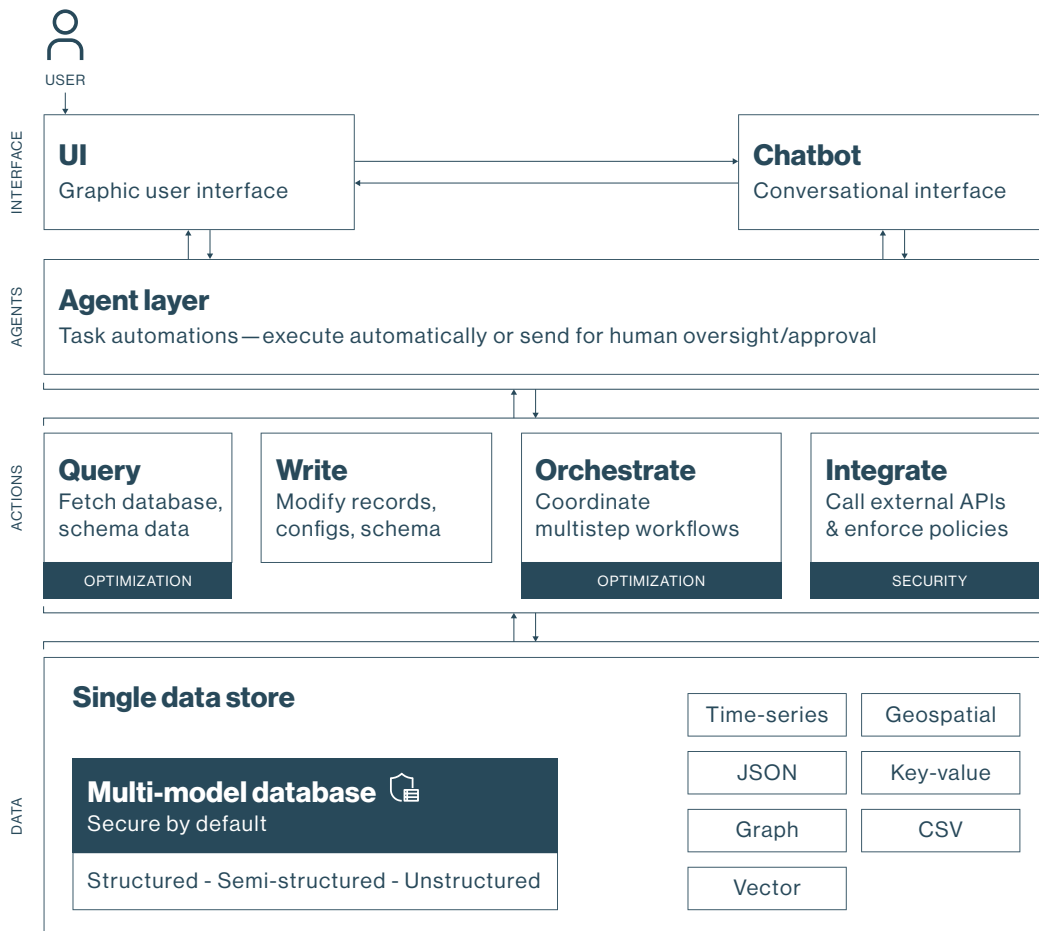
Autonomous Postgres® platform for AI-scale, multi-model, distributed systems



Solution overview

Modern data platforms are under unprecedented pressure to support increasingly diverse workloads while maintaining performance, reliability, and security. Traditional PostgreSQL deployments, while powerful and flexible, still require significant manual intervention for tuning, scaling, and lifecycle management. At the same time, organizations are introducing specialized databases to support AI, time-series, and document workloads—leading to fragmented architectures and operational complexity.

EDB Postgres AI (EDB PG AI) extends Postgres into a self-optimizing, multi-model data platform, embedding operational intelligence directly into the engine. Rather than relying on external observability tools or manual tuning workflows, the system continuously monitors, analyzes, and optimizes itself in real time. This enables platform teams to operate at scale with significantly reduced operational overhead while supporting modern AI and distributed application requirements.



Architectural principles

EDB PG AI's agentic database is built on three foundational architectural principles: embedded intelligence, unified multi-model data, distributed scalability, and sovereign control.

1. The first principle is **embedded intelligence**: Operational knowledge is integrated directly into the Postgres platform. This intelligence is informed by more than two decades of PostgreSQL expertise and enables autonomous tuning, proactive recommendations, and continuous optimization without external tooling.
2. The second principle is **multi-model unification**, allowing a single database engine to support relational, document, time-series, graph, and vector data. All data types are accessible through a standard SQL interface, eliminating the need for multiple specialized systems and reducing integration complexity.
3. The third principle is **sovereign-by-design control**, ensuring that all autonomous actions operate within enterprise-defined guardrails. Organizations retain full control over data residency, access policies, and compliance requirements, regardless of deployment model.

Core technical capabilities

EDB PG AI integrates a broad set of capabilities into a cohesive platform that supports both operational and analytical workloads.

Autonomous optimization and operations

The platform continuously monitors more than 200 performance and operational metrics, enabling it to detect anomalies, recommend optimizations, and automatically apply changes where permitted. This reduces manual tuning effort by up to 80% and allows issues to be resolved significantly faster than by way of traditional approaches.

Unlike external observability tools, which provide insights but require human action, EDB PG AI closes the loop by executing optimizations directly within the engine.

Multi-model data support

EDB PG AI natively supports multiple data types, including:

- Relational (structured data)
- JSON/document (semi-structured data)
- Time-series (event and telemetry data)
- Graph
- Vector (for AI workloads including pgvector)

This unified approach allows developers to work with diverse data types using a single SQL interface, simplifying application design and reducing the need for data movement between systems.

Performance benchmarks indicate significant advantages, including as much as 5x–150x faster JSON processing compared to stand-alone document databases and improved performance for vector search.

EDB PG AI is designed to support transactional, analytical, and AI workloads natively. This architecture allows AI applications to access and operate on data directly within the database. The result is reduced latency, improved governance, and simplified infrastructure.

Security and governance

Security and governance are integrated into the platform rather than layered on externally. Features include:

- Role-based and row-level access control
- Transparent Data Encryption (TDE)
- Security recommendations
- Audit logging and compliance tracking
- Software bill of materials (SBOM) support

Autonomous security mechanisms can identify and remediate critical risks within minutes, significantly reducing exposure windows.

Importantly, all automation operates within defined policies, ensuring that enterprises maintain control over security decisions.

Scalability

The platform also supports workload-aware optimization, ensuring that resources are allocated efficiently, based on demand. EDB PG AI automatically scales up CPU and/or memory resources when usage thresholds are exceeded.

Architectural comparison

Capability	EDB PG AI	Oracle Autonomous Database	Aurora/Azure Postgres	Specialized Database Stack
Architecture model	Agentic, embedded intelligence	Autonomous (Oracle-specific)	Managed service	Fragmented systems
Data model support	Unified multi-model	Limited	Limited	Multiple engines required
Operational model	Self-optimizing	Automated (vendor controlled)	Partially automated	Manual across systems
Deployment	Hybrid, on prem, multi-cloud	Cloud only	Cloud only	Mixed
Control and governance	Full enterprise control	Vendor controlled	Vendor controlled	Distributed across databases

*Competitive comparisons are based on publicly available information and are subject to change as vendor offerings evolve and new information is made available. All product names, trademarks, and registered trademarks are the property of their respective owners.

Operational model comparison

Operational Function	EDB PG AI	Traditional Postgres	Cloud Postgres	Specialized Stack
Performance tuning	Recommendations, automation	Manual	Partial automation	Manual
Issue detection	Proactive	Reactive	Reactive	Reactive
Scaling	Automatic + policy driven	Manual	Semiautomatic	Manual
Monitoring	Built-in intelligence	External tools	External tools	Multiple tools
Incident resolution	Automated + assisted	Manual	Assisted	Manual

*Competitive comparisons are based on publicly available information and are subject to change as vendor offerings evolve and new information is made available. All product names, trademarks, and registered trademarks are the property of their respective owners.

Developer and platform experience

EDB PG AI simplifies both developer and platform workflows by reducing complexity and unifying capabilities.

For platform engineers, the reduction in manual tuning and operational overhead allows teams to focus on system design and optimization rather than maintenance. The inclusion of a conversational interface further simplifies diagnostics, enabling engineers to interact with the system using natural language to query system state, request diagnostics, and manage the platform without scripting.

For developers, the ability to work with multiple data types within a single database reduces architectural complexity and accelerates development cycles. Applications no longer need to integrate with multiple data stores, simplifying both development and deployment.

For AI engineers, the integration of vector search and MCP support provides a direct path to building AI-native applications. Data can be accessed and processed within a governed environment, reducing the need for external pipelines and improving performance.

Performance and efficiency gains

EDB PG AI delivers measurable improvements in performance and efficiency across multiple dimensions.

The platform can process JSON workloads significantly faster than stand-alone document databases and provides substantial improvements in hybrid and vector search performance. Vector query performance is enhanced through optimized implementations, enabling faster AI inference and retrieval operations.

Operational efficiency is also improved through automation, with reductions in manual tuning and faster issue resolution contributing to overall system reliability and performance.

Technical comparative analysis: Agentic database vs. alternatives

Technical Dimension	EDB PG AI	Oracle Autonomous Database	Cloud Postgres (Aurora/Azure)	Specialized Database Stack
Core architecture	Postgres-native with embedded intelligence	Proprietary Oracle engine	Managed PostgreSQL service	Multiple purpose-built engines
Operational model	Proactive, automated operations, scaling, tuning	Autonomous but Oracle specific	Semiautomated, requires manual tuning	Manual across systems
Observability and optimization	Built-in intelligence (200+ metrics, closed-loop optimization)	Built in but opaque	External tools required	Fragmented observability
Multi-model support	Native (relational, JSON, vector, time-series, graph)	Limited multi-model	Extension based, limited	Requires multiple databases
SQL interface consistency	Single unified interface for all data types	Oracle SQL variants	Standard PostgreSQL	Multiple query languages
AI/vector integration	Native pgvector + MCP (agent-native integration)	Add-on capabilities	Extensions and tooling	Separate vector database required
Data movement requirements	None (all workloads unified)	Moderate	Moderate	High (ETL across systems)
Performance optimization	Workload-aware optimization	Automated but Oracle bound	Limited, workload tuning required	Optimized per engine only
JSON/document performance	High-performance native JSON (benchmarks: 5x-150x faster performance)	Moderate	Moderate	Strong (MongoDB, etc.) but isolated
Hybrid/multi-cloud deployment	Full support (on prem, hybrid, multi-cloud)	Cloud only	Cloud only	Complex, fragmented
Data sovereignty control	Full enterprise control (containerization, policies)	Limited	Limited	Complex across systems
Security model	Embedded, policy-driven, autonomous remediation	Strong but vendor controlled	Shared responsibility	Fragmented security layers
Extensibility	700+ Postgres ecosystem integrations	Limited to Oracle ecosystem	Moderate	Depends on each system
Vendor lock-in risk	Low (open Postgres standard)	High	Medium	High (multi-vendor dependencies)
Operational overhead	Low (boosts efficiency by up to 30%)	Medium	Medium	High (multiple systems)

*Competitive comparisons are based on publicly available information and are subject to change as vendor offerings evolve and new information is made available. All product names, trademarks, and registered trademarks are the property of their respective owners.

EDB Postgres AI: The sovereign data and AI platform for the agentic enterprise

EDB PG AI brings together a unified data layer, governance, sovereign control and orchestration, and an agent runtime environment, giving enterprises a trusted foundation for AI on infrastructure they own and control. The platform unifies transactional, analytical, and AI workloads in a single Postgres-based architecture—eliminating ETL, data movement, and operational fragmentation. And you choose where and how to deploy: on-premises, cloud, managed, or certified appliance.

The outcome: production-ready sovereign AI in days or weeks, not months.



EDB Postgres® AI (EDB PG AI) is the sovereign data and AI platform for the agentic enterprise. Built on Postgres, the world's leading open source database, EDB PG AI unifies transactional, analytical, and AI workloads in a single governed architecture, on-premises and across clouds. To learn more, visit www.enterprisedb.com.

© EnterpriseDB Corporation 2026. All rights reserved.