TECHNICAL BRIEF

EDB Postgres® AI Governance

How EDB Postgres AI embeds governance directly into the Postgres database—and why it matters

With agents waiting to go to production, governance is not just a cataloging or monitoring problem. It is an execution problem: Systems must evaluate data permissions, action purpose, and intent simultaneously, at machine speed, and with full contextual awareness.

This means that enterprises will be required to connect the internal mechanics of their source databases with the broader pressures from AI agents, high-frequency decisioning, and regulatory expectations.

EDB Postgres AI (EDB PG AI) addresses this by bringing governance into the data layer, where every query is intercepted and evaluated before execution. Extending session context with intent metadata enables decisions based on purpose as well as on identity and access. Pre-execution control prevents violations before they occur. With dual provenance, the platform also provides full auditability of what data was accessed and why it was authorized. This means there is a complete record at the moment agents take action, not a reconstruction devised after the fact. Because all access paths converge in the database, policy enforcement remains consistent.

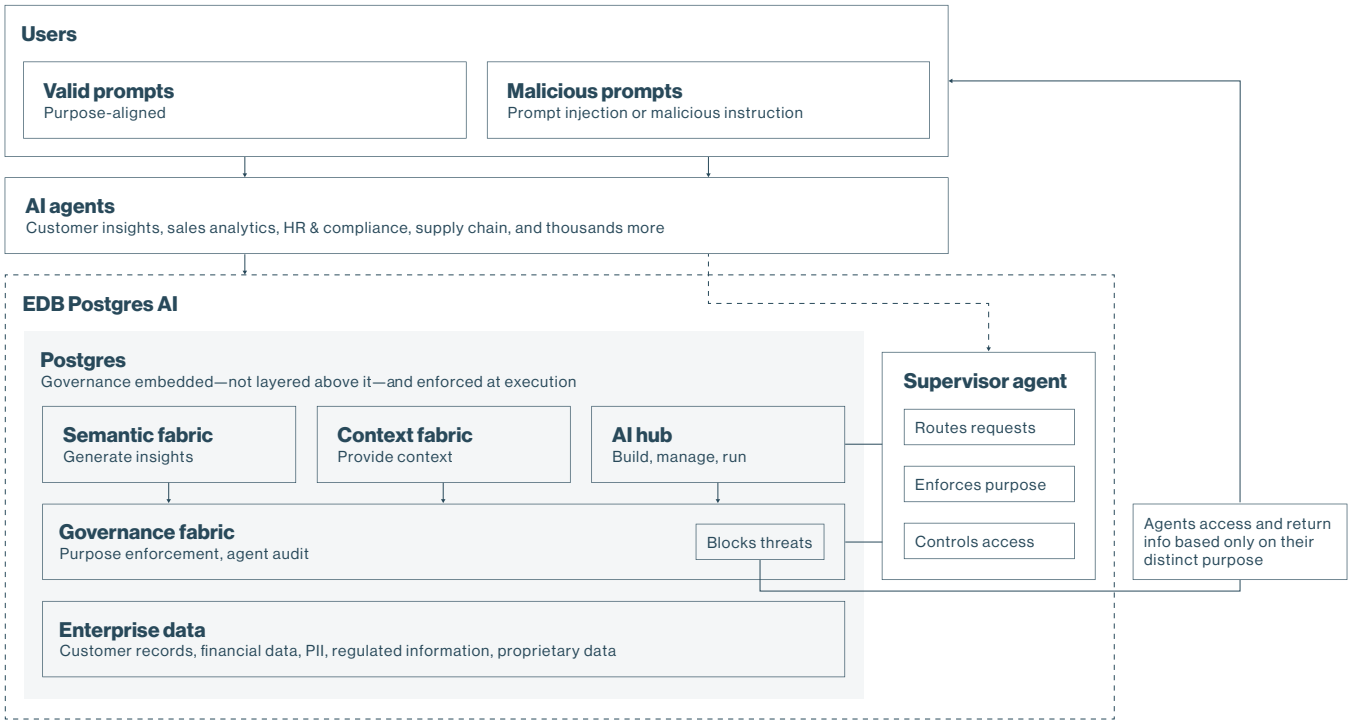

The EDB PG AI approach

Active enforcement boundary

EDB PG AI extends the role of Postgres from a passive database into an active enforcement boundary into an active enforcement boundary. Governance decisions are made at the exact moment a query executes.

The distinction matters. Catalog tools such as Collibra already provide robust policy definition and documentation to support governance. Failures are instead caused by the mismatch between where policy is defined and where decisions actually occur. Catalog tools define what governance should look like, and EDB PG AI enforces it at the only point where governance is physically reliable: the database.

Intercepting the query lifecycle

EDB PG AI embeds governance directly into its Postgres database. Every query, regardless of origin, is intercepted and converted into a normalized form that EDB PG AI evaluates before execution begins. This ensures that governance is applied when full query intent is known but before any data is accessed. Unlike upstream API gateways, which lack full query semantics, or downstream audit systems, which act too late, EDB PG AI ensures that no unauthorized or inappropriate action can physically execute inside the database engine.

Attaching intent to queries

Traditional governance answers, “Who are you?” Agentic governance requires answering, “What are you trying to do right now?”

To handle AI agents, EDB PG AI extends the Postgres session model by introducing structured context, augmenting identity (user/role) with intent metadata such as task, agent type, and sensitivity level. This metadata is injected into the session via configuration variables or protocol extensions, making it available during planning and execution.

By embedding intent into the execution context, EDB PG AI can distinguish between similar queries issued for different purposes, then enforce different outcomes. This prevents valid data access from turning into inappropriate use, which application-layer guardrails, prompt injection protection, and policy catalogs can’t reliably stop.

Embedding a policy engine in execution

Rather than delegating decisions to an external service, EDB PG AI implements and enforces it inside Postgres. Policies governing data access and intent are evaluated together during query execution.

The key advantage is determinism and performance. Because policies are precompiled, EDB PG AI evaluates governance from a stable foundation. The same query under the same context produces the same outcome. By eliminating network calls and external dependencies, EDB PG AI executes governance decisions at database speed. More important, because data access and intent are evaluated together inside the same engine, EDB PG AI resolves the long-standing fragmentation between data governance and AI governance.

Pre-execution decisioning

Once the query and its context are evaluated, EDB PG AI performs a pre-execution decision: Allow the read, or deny it. Governance enforcement occurs before the database touches storage buffers or returns any rows. This eliminates entire classes of risk, including accidental leakage through partial results, side channels, or logging artifacts. It also enables fine-grained control, meaning queries are not simply blocked but can be safely transformed to comply with policy. This preserves usability without sacrificing control.

Maintaining real-time performance

Production workloads such as fraud detection operate under millisecond latency constraints. Governance that introduces external dependencies cannot meet that bar.

EDB PG AI is designed to meet this requirement. Because policy and intent enforcement runs in the database, precompiled and executed without external calls, governance decisions do not introduce network hops or integration overhead. Embedding governance inside Postgres means enforcement scales with query throughput rather than becoming a bottleneck. In effect, governance inherits the performance characteristics of the database itself.

Capturing dual provenance

EDB PG AI adds the observability that enterprises need to operate Postgres at scale. It now extends that observability by recording two records for every query: what data was accessed, and why that access was authorized. This includes which policies were evaluated and what purpose was declared at the moment the decision was made. Governance transforms from a control mechanism into an accountability framework.

In regulated environments, it is no longer sufficient to prove that access occurred. Organizations must demonstrate that access was justified under specific conditions. By capturing dual provenance (both records) at execution time, EDB PG AI enables complete auditability, supports forensic analysis, and provides the foundation for explainable AI systems.

These audit logs are also the standardized trace records that enterprises require to populate their lineage and compliance reporting surfaces. This makes EDB PG AI a necessary complement to existing governance catalog and gateway solutions.

Ensuring consistency across humans and agents

Because all access ultimately resolves to SQL executed within the database, EDB PG AI enforces governance uniformly across all actors: applications, human analysts, and AI agents alike. There is no distinction in the enforcement path, which eliminates the inconsistencies that arise in layered architectures in which policies are implemented differently across APIs, services, and tools. This prevents governance drift, a major risk in environments where policy definition in catalog tools is not matched by reliable enforcement at the data layer. Policy intent and actual enforcement remain aligned over time, regardless of how access patterns evolve.

Structured summary of key advantages

Capability	EDB PG AI	Technical implementation detail	Why it makes a difference
Enforcement point	Database	Intercepts every query before execution begins	Governance is applied exactly where execution happens—the only point no agent can bypass
Identity and intent	Session context extension	Structured metadata attached to the session	Enables intent-aware decisions, not just identity-based access
Data governance	Row-level security and masking	Applied per declared purpose	Agents can only reach data consistent with their declared purpose, even if credentials permit more
Pre-execution control	Allow or block	Execution is reviewed, then allowed, modified, or blocked before data is touched	Unauthorized queries never execute—no exposure, no partial results, nothing to remediate
Performance	In-database evaluation	Policy evaluation runs inside PostgreSQL without external calls	Governance does not introduce network hops or integration overhead
Dual provenance	Extended audit tables	Records what data was accessed and why that access was authorized at execution time	Complete accountability at the moment it happens, not reconstructed after the fact
Consistency	Single enforcement path	All access resolves through the database	Eliminates policy drift across 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.

These capabilities make up a unified enforcement system for agents accessing core data, one that operates at the only point in the stack where governance cannot be bypassed.

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 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.