CIO BRIEF

Enforcing Data and AI: Governance at the Execution Layer

The architectural challenge: Governance is upstream

Most enterprises suffer from a structural misalignment between where governance policies are defined and where data access decisions are actually executed. Policies are defined in catalogs, gateway tools, and documents, but agent requests are executed in the database. This separation introduces three major risks:

- The first is **latency**. Policies enforced upstream prevent full visibility into the query being executed, while downstream controls engage too late, after data has already been accessed.
- The second is **loss of context**. External enforcement layers cannot interpret the intent behind a query. They record who accessed data and when, but not why it was authorized nor whether the purpose matched the action.
- The third is **bypass risk**. Any agent connecting to the database through a misconfigured service account, a malicious prompt injection, or a direct connection circumvents every upstream control. Traditionally, those violations surface in an audit log after the data has already been read.

In this environment, governance exists in theory but not in practice. It can document violations but cannot prevent them.

The EDB Postgres® AI approach: Governance at the point of execution

EDB Postgres AI (EDB PG AI) addresses these challenges by bringing data and agent governance into Postgres as an inherent part of query processing. This turns the database from a passive record into an active enforcement boundary, where every query is evaluated before it executes.

EDB PG AI achieves this by ingraining identity, declared agent purpose, permissions, and policies in Postgres natively. This means you know not only who is making the request and when, but also the intent behind the request and the reason why it was authorized.

Governance is enforced at the only point in the system where both the full intent of the request and the actual data access path converge: inside the database itself.

The core model

EDB PG AI's agent governance model is built around three capabilities:

- The first is **purpose-bound access control**. Every agent session declares an explicit business purpose at the point of connection. EDB maps that purpose to a PostgreSQL role, which determines what data the agent can access. Row-level security policies enforce those boundaries dynamically, ensuring that agents can only reach data consistent with their declared purpose.
- The second is **accountability for agent actions**. When an agent acts, it acts on behalf of a human or organization. EDB PG AI preserves and records that relationship, so every data action—whether a human or an agent executed it—can be traced back to the person or entity responsible. Accountability does not disappear when an agent is in the loop.
- The third is **session-scoped auditability**. EDB PG AI binds every query within a session to a single session ID, creating a complete, ordered record of everything an agent did within a given context. This gives compliance and audit teams a coherent, end-to-end narrative rather than disconnected log entries.

Intent-aware session context

Identity tracking alone is not sufficient for governance. EDB PG AI connects an agent's declared purpose to each session, giving organizations the ability to understand why an agent request was made. This means that two different agents can query the same table, but they will not have access to the same data if it does not align with their approved intent.

Pre-execution enforcement eliminates risk

Other governance tools, such as catalogs, gateways, and policy engines, operate above the database. They can provide a certain level of visibility and produce a record of completed actions. EDB PG AI operates inside the database and can approve or block queries depending on the agent's access, purpose, and intent. Because the decision is made before the agent touches any data, there is no exposure and nothing to remediate after the fact. Unauthorized queries never execute.

Performance and operational viability

Governance at the execution layer does not come at the cost of performance. Because all governance logic runs inside Postgres, EDB PG AI adds no network calls or external dependencies. Enforcement happens at the speed of your database and without latency, supporting critical workloads that depend on access to real-time data such as fraud detection. Governance remains intact if a failover occurs, without adding additional operational overhead.

Dual provenance and accountability

EDB PG AI captures two records for every query: what data was accessed, and why that access was authorized. This gives enterprises the complete picture that regulators, auditors, and compliance teams require: proof that access occurred, and also proof that it was justified under specific conditions at the moment it happened. For CIOs, that means the audit answer is pre-built. Security, compliance, and finance get what they need to sign off on agents moving to production.

Comparative analysis

When evaluating governance approaches for agentic workloads, the most important factors to consider are where enforcement actually runs, and whether it can be bypassed.

Databricks, for example, has invested meaningfully in governance, extending catalog and gateway layers to cover agent interactions alongside data assets. These improve governance outcomes but are not complete solutions, as these controls operate outside the database. The guarantee of governance depends on every agent using the approved connection path—an assumption that does not hold at scale. Cloud-only architectures such as Snowflake carry an additional constraint: They limit deployment options for organizations with on-premises or air-gapped deployments.

Do-it-yourself architectures, which combine multiple tools for identity management, policy enforcement, and data access, offer flexibility at the cost of complexity. These systems inevitably suffer from synchronization challenges, increased latency, and a proliferation of failure modes. Over time, they become difficult to maintain and prone to governance drift.

Governance catalog tools such as Collibra, Informatica, and Alation define, document, and communicate data policy. EDB PG AI enforces that policy at the layer where data actually lives. The two are complementary: Catalog platforms provide the policy framework EDB PG AI enforces, and EDB PG AI provides the standardized audit and trace logs those tools require. Enterprises do not choose between them. They use both for a complete governance approach.

Comparative platform analysis: EDB PG AI vs. alternatives

While many platforms address elements of data governance, they differ fundamentally in where and how governance is enforced. The distinction is most evident when evaluating execution-time control.

EDB PG AI is differentiated by embedding governance directly into the database execution layer. In contrast, Databricks, Snowflake, MongoDB, and DIY architectures rely on layered or externalized approaches that introduce latency, fragmentation, or operational complexity. These differences become especially pronounced in environments requiring real-time decisioning, regulatory compliance, and AI accountability.

Capability	EDB PG AI	Databricks	Snowflake	DIY architecture
Governance enforcement point	Embedded in the Postgres database	External layers above the database, bypassable	External layers above the database, bypassable	Distributed across services and APIs, inconsistent
Execution-time control	Enforced at execution—allow, block	Gateway capability	Gateway capability	Inconsistent
Intent-aware governance	Native, session-level purpose and intent	Requires external tooling	Not native	Custom built, if any
Governance ecosystem complement	Focus on database layer, complementary to gateway and catalog tools	Partial, depends on orchestration	Policy defined in external catalog, limited enforcement depth	Integration required
Data sovereignty	Full control: self-managed, hybrid, and air-gapped deployments	Cloud-centric	Cloud-centric	Depends on architecture
Auditability	Dual provenance—what data was accessed and why it was authorized	Strong lineage, limited decision explainability	Strong lineage, limited decision explainability	Fragmented

*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.

Strategic implications for CIOs

Governance becomes a runtime capability rather than a compliance overlay. That shift has a direct and immediate consequence for CIOs: The blockers that stall agent deployments, such as security sign-off, compliance approval, audit defensibility, and cost attribution, are addressed at the infrastructure level, before they become organizational negotiation.

Security teams get enforcement they can verify, not policy they have to trust. Compliance teams get a complete, pre-built audit record at the moment of execution, instead of a reconstruction after the fact. Finance teams get purpose-bound attribution they can forecast against. And CIOs get agents in production.

EDB PG AI does not ask enterprises to introduce new governance processes or new vendor relationships. It embeds enforcement into the infrastructure they already run. It eliminates the need for multiple enforcement layers and extends, to every agent that acts on it, the same rigorous controls already applied to the most sensitive operational data.

EDB PG AI already powers some of the most critical and regulated workloads on the planet, from U.S. defense networks, for which data sovereignty is nonnegotiable, to global payment infrastructure and national energy grids, for which downtime and data exposure are not options. When organizations such as these are ready to deploy agents on that same data, the governance foundation is already in place. EDB PG AI extends that same trust to the agentic workforce.

Unified capability model

Capability	How EDB delivers	Why it matters
Enforcement point	Governance runs inside the PostgreSQL execution engine.	Every query is evaluated at the only point no agent can bypass.
Identity and intent	Declared purpose and identity are attached to every session.	The database evaluates not just who is asking but why.
Purpose-bound access	Row-level security policies enforce boundaries dynamically.	Agents can only reach data consistent with their declared purpose.
Pre-execution control	Queries are allowed, blocked, or modified before execution.	Unauthorized queries never execute. No exposure, nothing to remediate.
Performance	Governance logic doesn't leave the database	No network calls, no external dependencies, no latency cost.
Provenance	What data was accessed and why it was authorized.	Audit evidence is generated at execution, not reconstructed after.
Resilience	Governance stays intact alongside data replication in EDB PG AI's distributed architecture.	Enforcement remains intact across global events.

*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 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.