CIO BRIEF

Reclaiming Control of the Enterprise Data Warehouse

Modernize analytics with WarehousePG to reclaim control over costs and data sovereignty.

Enterprise data platforms are at an inflection point. Over the past decade, most organizations moved their analytics infrastructure in one of two directions: They either remained on legacy on-premises data warehouses or migrated aggressively to cloud-only analytics platforms. Both approaches delivered value initially. Today, however, both are creating new strategic challenges for CIOs.

Three forces are converging to reshape the data warehouse landscape: escalating cloud costs, increasing data sovereignty requirements, and the aging of legacy analytics platforms. Together, these pressures are forcing organizations to reconsider how—and where—their most critical data infrastructure operates. For CIOs responsible for balancing innovation, cost discipline, and regulatory compliance, the question is no longer whether change is required but how to modernize without introducing unnecessary risk or disruption.

The emerging cost problem

Cloud data warehouses such as Snowflake, Redshift, and Databricks transformed analytics by providing elastic scale and simplified operations. However, their consumption-based pricing models are now exposing enterprises to unpredictable cost structures.

As analytics workloads expand and organizations begin integrating AI and machine learning pipelines, compute consumption rises sharply. Charges tied to concurrency, query execution, storage growth, and data movement often compound in ways that make forecasting difficult. Many enterprises now report that analytics infrastructure is one of the fastest-growing components of their cloud spend.

This “cloud cost shock” has led many CIOs to reconsider whether all workloads should remain in cloud-only platforms. Instead, organizations are increasingly evaluating hybrid strategies that retain cloud flexibility while restoring economic predictability.

Data sovereignty is becoming nonnegotiable

At the same time, regulatory and geopolitical dynamics are placing new constraints on how data can be stored and processed. Regulations such as GDPR, the U.S. CLOUD Act, and emerging national data residency policies are forcing organizations—particularly in regulated industries—to maintain tighter control over their data environments.

For financial institutions, public sector agencies, healthcare organizations, and global enterprises operating in multiple jurisdictions, the ability to determine where data resides and who governs access is now a board-level issue. Cloud-only data warehouse architectures often conflict with these requirements. CIOs increasingly need platforms that can operate across multiple environments, including on-premises data centers, sovereign cloud regions, and hybrid deployments.

Legacy platforms are reaching end of life

While cloud economics and governance pressures are rising, many enterprises are still running critical analytics workloads on legacy data warehouse systems such as Teradata, Netezza, Exadata, or Greenplum.

These platforms are reaching natural renewal cycles. Maintenance costs are increasing, performance scaling is limited, and integration with modern AI and analytics workflows can be complex. Vendor consolidation in parts of the data platform market has also introduced uncertainty around long-term support and product direction.

For CIOs, this convergence of cost pressure, regulatory requirements, and aging infrastructure creates a clear modernization trigger. The challenge is finding a platform that delivers the performance and scale of modern analytics while maintaining control over cost and governance.

A new model: Open, sovereign data warehousing

EDB Postgres AI for WarehousePG represents an emerging approach to this challenge. Built as an open source massively parallel processing (MPP) data warehouse based on Postgres, it enables enterprises to run large-scale analytics workloads without the constraints of cloud-only architectures or proprietary legacy systems.

The platform combines the scalability associated with modern cloud data warehouses with the flexibility of open infrastructure. Organizations can deploy WarehousePG across multiple environments, including on-premises infrastructure; public cloud platforms such as AWS, Azure, or Google Cloud; or hybrid architectures that span both.

This flexibility allows CIOs to align analytics infrastructure with governance requirements, economic priorities, and operational strategy rather than being constrained by vendor-specific deployment models.

Predictable economics

One of the most significant advantages of this model is cost predictability. WarehousePG uses capacity-based licensing tied to compute resources rather than consumption-based pricing tied to query activity or data movement.

For enterprises running high-concurrency analytics environments, this shift can significantly improve cost transparency and long-term planning. Independent benchmarking has demonstrated that organizations can achieve substantial total cost of ownership reductions compared with cloud-only platforms, particularly over multiyear operating horizons.

For CIOs managing increasingly complex cloud budgets, predictable infrastructure economics are becoming an essential part of data strategy.

Reducing migration risk

Historically, the biggest barrier to data warehouse modernization has been migration complexity. Rewriting queries, refactoring applications, and retraining teams can make platform transitions slow and expensive.

WarehousePG addresses this challenge by leveraging SQL compatibility and existing Postgres ecosystem tools. For many legacy warehouse platforms, migration can occur with minimal application changes. In environments running Greenplum, WarehousePG offers binary compatibility, enabling migrations with little or no query modification.

This dramatically reduces the operational risk associated with modernization projects and allows organizations to validate new architectures through targeted pilot workloads before broader adoption.

Proven enterprise deployments

Several enterprise deployments illustrate how this approach can operate at scale. Organizations in financial services, retail, and digital advertising are already using WarehousePG to process high-volume analytics workloads ranging from tens of terabytes to multiple petabytes of data.

These deployments demonstrate that open, Postgres-based analytics platforms can deliver the performance and reliability required for mission-critical environments while maintaining governance flexibility and cost control.

The strategic opportunity

For CIOs, the evolution of the data warehouse market presents an opportunity to reset the balance between innovation and control. Modern analytics platforms must support AI-driven workloads, massive concurrency, and global data distribution. At the same time, they must operate within tighter cost, regulatory, and governance constraints. WarehousePG offers a path to modernize analytics infrastructure while reclaiming control over where data lives, how it is managed, and how costs evolve over time.

The shift toward sovereign, open data platforms is accelerating. Organizations that begin this transition now can reduce long-term infrastructure costs, mitigate regulatory risk, and build analytics foundations designed for the next generation of data and AI workloads.

The 5 questions CIOs should ask their teams and vendors

As organizations evaluate the next generation of data warehouse platforms, CIOs should challenge both internal teams and technology providers with five critical questions:

1 What is our true three-year cost model, and are we underestimating it?

Many analytics platforms appear inexpensive initially but scale unpredictably as workloads grow. CIOs should require clear answers to:

- What is the projected cost of our analytics platform over three years?
- How does pricing scale with concurrency and AI workloads?
- What portion of costs are consumption-based versus predictable capacity?

The goal is to ensure the platform supports sustainable economics as analytics demand expands.

2 How difficult will migration really be, given where we want to go?

Migration complexity often becomes the largest barrier to modernization. CIOs should probe:

- Will applications require refactoring?
- Will SQL queries remain compatible?
- How long do typical migrations take for organizations like ours?

The most viable platforms minimize disruption and allow organizations to prove value quickly through targeted migrations.

3 Does the platform we are on prevent vendor lock-in, or are we inflexible facing the future?

Vendor lock-in can limit long-term strategic flexibility. Key questions include:

- Is the platform based on open standards or proprietary architecture?
- Can workloads move between environments?
- Are we dependent on a single vendor ecosystem?

Open architectures reduce long-term risk and improve negotiating leverage.

4 Is the platform ready for the next generation of data and AI workloads, and if not, what do we do?

Analytics platforms must increasingly support AI-driven workloads and large-scale data processing. CIOs should ask:

- Can the platform scale to petabyte-level analytics?
- Does it support high concurrency and mixed workloads?
- How easily can AI pipelines integrate with the data platform?

The right architecture must support both today's analytics needs and tomorrow's AI workloads.

5 How difficult will migration really be, given where we want to go?

Migration complexity often becomes the largest barrier to modernization. CIOs should probe:

- Will applications require refactoring?
- Will SQL queries remain compatible?
- How long do typical migrations take for organizations like ours?

The most viable platforms minimize disruption and allow organizations to prove value quickly through targeted migrations.

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.