
Your Databricks Estate Is Costing More Than You Think

The hidden risks of unmanaged platform growth

A leadership briefing on cost leakage, estate redundancy, governance drift, security exposure and AI-era privacy risk across maturing Databricks environments.

\$723B

forecast worldwide public cloud end-user spend in 2025, per Gartner.¹

29%

estimated wasted cloud spend in Flexera's 2026 State of the Cloud research.²

\$4.44M

global average cost of a data breach in IBM's 2025 study.³

90%

of organisations say privacy programmes expanded because of AI, per Cisco.⁴

The practical question for Databricks customers is no longer whether the platform can create value. It is whether leadership can see where cost, risk, duplication and data exposure have accumulated before they become operating liabilities.

1. The problem after adoption: Databricks becomes an estate.

Most organisations focus intensely on getting Databricks live: migrations, lakehouse architecture, pipelines, dashboards, ML workflows and AI use cases. After adoption, the quieter problem begins. The platform grows faster than the organisation's ability to govern, optimise, secure and explain it.

FROM USE CASE TO ESTATE

Databricks rarely remains a single workspace or a single delivery programme. Its product surface spans data engineering, warehousing, governance, BI, application development, ML and AI.⁵ Azure Databricks is described as an enterprise-grade platform for building, deploying, sharing and maintaining data, analytics and AI solutions at scale, with cloud infrastructure managed and deployed in the customer's cloud account.⁶

That breadth is valuable. It is also why a platform can quietly become an estate: a living system of workspaces, clusters, jobs, SQL warehouses, dashboards, notebooks, tables, service principals, external locations, model endpoints, embeddings, vector indexes and LLM-powered workflows.

The leadership question changes from: **Can we build on Databricks?** to: **Can we prove the estate is efficient, secure, governed and proportionate?**

WHAT UNMANAGED GROWTH CREATES

- Compute and SQL warehouses created by many teams with uneven standards.
- Dashboards, jobs and tables that continue to run after ownership or value becomes unclear.
- AI workloads that spread through APIs, SDKs, notebooks, agents and applications.
- Permissions and service principals that accumulate faster than reviews occur.
- Business logic duplicated across pipelines, marts, dashboards and feature stores.
- Leadership reporting that blends production value, experimentation, waste and risk into one opaque number.

This is not just a DBU problem. The full cost of an unmanaged estate includes cloud infrastructure, duplicated workloads, unnecessary refreshes, hidden AI usage, security exposure, privacy evidence gaps, operational drag and poor executive visibility.

AI is accelerating the estate problem. McKinsey reported that 78% of organisations used AI in at least one business function in its 2025 survey.⁷ Flexera reported 81% of organisations using generative AI in 2026.⁸ Gartner predicts 50% of cloud compute resources will be devoted to AI workloads by 2029, up from less than 10% today.⁹

\$4.8B

Databricks revenue run-rate reported in Q3 2025.¹⁰

20k+

global Databricks customers reported by Reuters.¹⁰

78%

organisations using AI in at least one function, per McKinsey.⁷

25%

Gartner prediction for significant cloud adoption dissatisfaction by 2028.⁹

2. The hidden cost problem is bigger than DBUs.

Databricks cost issues are rarely caused by one visible mistake. They come from many small choices that become normal: oversized clusters, loose policies, high-frequency jobs, SQL warehouse sprawl, dashboards refreshing by habit, duplicate pipelines and AI services left without clear usage controls.

WHY SPEND BECOMES OPAQUE

Elasticity allows teams to move quickly. It also allows spend to become distributed. A small change in runtime, worker count, instance family, Photon usage, autoscaling range, warehouse size, schedule cadence or model endpoint configuration can materially change cost.

Flexera reported that 85% of organisations cite managing cloud spend as a top challenge, and estimated wasted cloud spend at 29%.⁸ For Databricks customers, this is the macro context: the issue is not that teams do not know the bill, but that they cannot always explain which behaviour produced it.

A useful diagnostic: for every 100 DBUs consumed last month, can the organisation identify the owner, workload, business unit, environment, policy, schedule and reason?

Databricks system tables are designed to support this kind of analysis. They provide operational data for costs, security events, compute, jobs and data/AI workloads; billing tables centralise account usage across regions and workspaces.¹¹

Autoscaling is useful, but it is not a cost strategy by itself. Databricks notes that compute autoscaling has scale-down limitations for Structured Streaming and recommends Lakeflow Spark Declarative Pipelines with enhanced autoscaling for streaming workloads.¹⁶ Some predictable workloads may require workload-specific policies, schedule design and ownership review rather than generic elasticity.

29%

estimated wasted cloud spend, Flexera 2026.²

85%

say cloud spend management is a top challenge.⁸

288

daily runs for a five-minute job schedule.

DBU/h

a native policy constraint for Databricks compute creation.¹³

RECURRING LEAKAGE PATTERNS

PATTERN	EXECUTIVE TRANSLATION
Five-minute schedules	A job running every five minutes executes 288 times per day; hourly is 24 times - a 12x cadence difference before compute size is considered.
Dashboard refresh sprawl	Scheduled Databricks dashboard updates run SQL logic and populate result caches; unnecessary schedules turn BI convenience into warehouse load. ¹²
Weak compute policies	Policies can set cluster type, maximum resources per user and maximum DBUs per hour; without them, cost controls become user discretion. ¹³
Idle compute	Automatic termination is a configurable inactivity control, not merely a good habit. ¹⁴
AI endpoints	AI Search endpoints incur charges after an index is created and can keep serving costs even when query traffic is absent. ¹⁵

The most dangerous Databricks costs are not always the largest line items. They are the ones nobody owns.

3. What leadership cannot see cannot be tuned.

As more teams build independently, duplication becomes inevitable. Similar tables, overlapping dashboards, copied notebooks, duplicate transformations and parallel feature pipelines can all look like productivity while creating multiple versions of the truth and multiple cost bases for the same business logic.

THE REDUNDANCY PROBLEM

Redundancy is expensive because it is rarely visible as redundancy. It appears as extra jobs, extra storage, extra dashboards, extra data movement, extra permissions and extra maintenance. It also weakens trust: teams see conflicting metrics and nobody can say which asset is canonical.

Unity Catalog lineage is one native mechanism for recovering visibility. Databricks says lineage can identify downstream tables, jobs and dashboards before changing or deleting a table or column, investigate root causes and track sensitive data flow for audits.¹⁷

Redundancy often looks like delivery velocity. Teams are shipping. Dashboards are appearing. AI experiments are launching. The estate may still be becoming less coherent, less governable and more expensive.

THE OBSERVABILITY GAP

Many organisations monitor pipelines and business dashboards. Fewer continuously monitor the Databricks estate itself: who owns spend, which workloads are critical, which assets are obsolete, which dashboards are used, which warehouses are underutilised and which permissions no longer reflect reality.

The FinOps Foundation says visibility into technology spend should be accessible, timely, accurate and consistent across organisational levels.¹⁸ Databricks system tables provide the raw account-level telemetry, but the discipline of interpreting it still has to be owned.¹¹

QUESTIONS AN ESTATE VIEW SHOULD ANSWER

- Which workloads are driving the top 20% of spend?
- Which jobs, dashboards or tables have no clear owner?
- Which assets have not been used recently but still consume cost or risk?
- Which business concepts exist in more than one implementation?
- Which workspaces drift furthest from policy?

WARNING SIGN	WHAT IT USUALLY MEANS	LEADERSHIP RISK
Blended Databricks bill	Usage is not reliably attributed by owner, product, environment or workload.	Finance can see spend but cannot challenge necessity.
Dashboard uncertainty	Reports exist that nobody can confidently retire or consolidate.	Executives receive metrics without knowing their lineage or value.
Personal production assets	Sandbox notebooks, personal clusters or informal jobs support business operations.	Continuity, security and support boundaries are unclear.
ClickOps configuration	Manual UI changes alter policies, schedules, warehouses or access without durable review.	The estate cannot be reconstructed from documentation or code.

A Databricks estate cannot be properly governed if leadership cannot see what exists, what it costs, who owns it and what risk it introduces.

4. Governance drift and security exposure compound in the AI era.

Governance often looks strong at design time. Naming standards exist. Unity Catalog structures are planned. Access models are approved. Then exceptions, experiments, copied jobs, temporary permissions and new AI workflows accumulate. The gap between the designed platform and the lived platform is governance drift.

SECURITY EXPOSURE IS ESTATE EXPOSURE

Databricks-specific advisories show why estate posture matters. CVE-2024-49194 affected Databricks JDBC Driver 2.x before 2.6.40 and could allow remote code execution via a JDBC URL parameter.¹⁹ Databricks also documents controls for no-isolation shared clusters because data or internal credentials may be accessible to code running in the same shared environment.²⁰

The broader security context is unforgiving. Verizon's 2025 DBIR analysed 22,052 incidents and 12,195 confirmed breaches; credential abuse and exploitation of vulnerabilities were leading initial attack vectors.²¹ GitGuardian reported 28.65 million new hardcoded secrets added to public GitHub commits in 2025.²²

An estate that was secure six months ago may not be secure today. New users, tokens, service principals, runtimes, dashboards, endpoints and integrations change the risk surface continuously.

PII IS NOW A RUNTIME CONTROL PROBLEM

PII can flow through notebooks, dashboards, pipelines, BI exports, features, prompts, embeddings, vector search and agents. Databricks' GDPR/CCPA guidance notes that deletion obligations can extend beyond Delta Lake to upstream sources such as Kafka, files, databases, queues and cloud storage.²³

Unity Catalog row filters and column masks, including ABAC policies using governed tags, are designed for consistent access controls across many tables and columns.²⁴ But controls only help when classification, ownership, masking, lineage and deletion evidence are continuously reviewed.

AI RAISES THE STAKES

- IBM reported 13% of organisations experienced breaches involving AI models or applications; 97% of those lacked proper AI access controls.²⁵
- Cisco found only 12% of AI governance committees were mature and proactive.⁴
- IBM put the global average breach cost at \$4.44 million.³

CONTROL AREA	WHAT SHOULD BE EVIDENCED
Classification	Where PII, financial data, health data, IP and confidential columns exist across catalogs and schemas.
Access	Which users, groups, service principals, jobs, dashboards and external assets can query sensitive data.
Lineage	Which downstream tables, dashboards, notebooks, models, extracts and reports consume regulated data.
Deletion	Whether right-to-erasure workflows reach Delta tables, derived tables, materialized views, streaming tables and upstream systems.

In the AI era, unmanaged data estates do not just waste money. They expand the blast radius of sensitive data.

5. The estate will not tune itself.

Databricks estates do not become inefficient overnight. They drift: a cluster is oversized, a dashboard refreshes too often, a job runs more frequently than the business needs, a temporary permission becomes permanent, a vector index remains live, a service principal is forgotten, and a table containing PII appears in a workflow nobody has reviewed.

A Audit

Inventory the estate as it actually exists: workspaces, runtimes, policies, SQL warehouses, jobs, clusters, dashboards, tables, tags, tokens, PII controls, data layout, AI workloads and cost attribution.

R Report

Translate technical evidence into leadership-grade visibility: cost leakage, redundancy, ownership gaps, governance drift, security exposure, privacy evidence gaps and affected business services.

T Tune

Prioritise and implement targeted improvements: rightsizing, schedules, policies, tagging, warehouse rationalisation, permission cleanup, lineage checks, runtime upgrades and retirement of stale assets.

The risk of not auditing is not that nothing is configured. It is that nobody can prove which configurations matter, which have drifted and which are costing the business today.

BOARD-LEVEL QUESTIONS

- What percentage of Databricks spend is owner-attributed and policy-governed?
- Which workloads carry the highest avoidable cost or latency?
- Can we trace sensitive data from source to dashboard, model or extract?
- Which production dependencies sit in personal, sandbox or legacy areas?
- What would we show an auditor within 48 hours?

THE DELTATUNE POSITION

Deltatune helps Databricks customers turn platform entropy into a quantified improvement backlog: lower DBU waste, faster workloads, clearer ownership, cleaner access controls, stronger privacy evidence and fewer surprises for finance, security, data governance and leadership.

This paper is not affiliated with or endorsed by Databricks. Databricks product names are used descriptively.

SOURCES

1. Gartner, public cloud spending forecast
2. Flexera, 2026 State of the Cloud
3. IBM, Cost of a Data Breach Report 2025
4. Cisco, 2026 Data and Privacy Benchmark Study
5. Databricks platform overview
6. Microsoft Learn, Azure Databricks introduction
7. McKinsey, State of AI 2025
8. Flexera, cloud value and AI waste release
9. Gartner, top trends shaping cloud
10. Reuters, Databricks valuation and revenue run-rate
11. Databricks system tables
12. Databricks dashboard scheduling
13. Databricks compute policies
14. Databricks compute configuration
15. Databricks AI Search cost management
16. Databricks Structured Streaming production guidance
17. Databricks Unity Catalog lineage
18. FinOps Foundation principles
19. Databricks CVE-2024-49194 advisory
20. Databricks admin isolation on shared clusters
21. Verizon, 2025 DBIR
22. GitGuardian, State of Secrets Sprawl 2026
23. Databricks GDPR/CCPA preparation for Delta Lake
24. Databricks row filters, masks and ABAC
25. IBM, AI model/app breach findings