Computomic

Unified Security & Governance on Databricks for a Large Global Bank

Consolidating fragmented security systems into a single enterprise control plane

Abstract

A large global bank operated a fragmented security and governance landscape spanning identity, privileged access, data classification, entitlements, masking, lineage, monitoring, and regulatory evidence. Computomic consolidated these control signals into a unified architecture on the Databricks Data Intelligence Platform, with Unity Catalog serving as the enterprise control plane for data and AI. The solution connected existing enterprise security systems to centrally enforced role- and attribute-based policies, automated classification, dynamic row filtering and column masking, lineage-aware monitoring, and policy-as-code deployment. Computomic delivered the program through embedded Forward Deployed Engineer pods combining Databricks, security, governance, DevSecOps, data engineering, and banking-domain expertise.

Executive Summary

The bank had invested in numerous security products, but controls operated independently. Classifications did not automatically drive masking; approved entitlements required manual database grants; and security events lacked consistent ownership, policy, and lineage context. Regulatory evidence was assembled manually across disconnected systems.

Computomic established Databricks Unity Catalog as the common governance layer while retaining specialist systems such as the enterprise identity provider, PAM, SIEM, DLP, catalog, and GRC platforms. Their outputs were normalized into a single metadata and policy framework that enforced controls consistently across workspaces, catalogs, data products, machine-learning assets, and AI applications.

Key Outcomes

- Central governance of data and AI assets across business domains and regions.
- Automatic classification and governed tagging of sensitive data.
- Hybrid RBAC and ABAC using jurisdiction, legal entity, business purpose, environment, and risk attributes.
- Dynamic row filtering, column masking, workspace restrictions, and time-bound privileged access.
- Centralized audit, lineage, entitlement reconciliation, and regulatory evidence generation.
- Repeatable deployment through Terraform, Databricks SQL, automated testing, and policy-as-code.

The Bank's Security Challenge

More than a dozen systems participated in protecting enterprise information. Identity and approval platforms, database roles, DLP scanners, catalogs, masking products, cloud IAM, audit repositories, SIEM, and GRC tools each addressed part of the control landscape. The issue was not missing technology; it was the absence of a common control plane connecting identities, approvals, classifications, policies, assets, lineage, and actual usage.

A classification tool could identify customer financial information without creating an enforceable mask. An access system could approve a user while a separate engineering team still had to implement platform-specific grants. A SIEM alert often lacked the asset owner, business purpose, data sensitivity, and downstream impact required for rapid investigation.

Unified Databricks Security & Governance Architecture

Computomic implemented a hub-and-spoke operating model. A central governance layer defined identity patterns, enterprise taxonomies, reusable security functions, policy templates, audit standards, and deployment automation. Business-aligned Databricks workspaces and catalogs retained domain ownership while inheriting mandatory enterprise controls.

Unity Catalog became the point where enterprise security context was converted into enforceable controls. Identity groups and service principals were synchronized at the account level. Production catalogs were bound to approved workspaces. Governed tags expressed confidentiality, jurisdiction, residency, business domain, purpose of use, and retention. RBAC established baseline access while ABAC dynamically applied row filters and column masks.

REPRESENTATIVE CATALOG MODEL

prod_retail_customer | prod_retail_payments | prod_markets_trading | prod_risk_credit | prod_finance_regulatory | prod_security_monitoring

Schemas followed a consistent lifecycle: landing, bronze, silver, gold, restricted, and published. Direct user grants were prohibited; approved access was provided through enterprise groups and continuously reconciled with entitlement records.

Dynamic Data Protection

Automated Classification and Governed Tags
Computomic built a classification normalization service that consolidated metadata from enterprise catalogs, DLP and discovery tools, legacy databases, schema registries, business glossaries, Databricks metadata, lineage analysis, and Computomic classification models. Inconsistent labels were translated into a common banking taxonomy.

Representative tags included classification=restricted, data_category=customer_pii, jurisdiction=eu, residency=eu_only, approved_purpose=servicing_and_fraud, retention_class=seven_year, and masking_policy=partial_reveal. Each finding retained its source, confidence, rule or model version, reviewer, approval status, and review date. High-confidence findings could invoke protective policies automatically; lower-confidence findings entered stewardship workflows.

Hybrid RBAC and ABAC
RBAC established baseline permissions to catalogs, schemas, tables, functions, models, and volumes. ABAC then determined what an approved user could see based on data sensitivity, jurisdiction, legal entity, user region, business purpose, environment, investigation assignment, and privileged-access status.

- EU customer records were limited to approved EU processing groups and workspaces.
- Trading data was restricted by legal entity and market function.
- Fraud investigators received detailed transaction data while other analysts saw masked values.
- Development users received tokenized identifiers even when source data originated in production.
- Support engineers received temporary access linked to an active PAM ticket.

Computomic created reusable row-filter and column-mask functions rather than embedding security logic separately in every table. Policies were performance-tested against interactive SQL, batch and streaming pipelines, BI, regulatory reporting, feature engineering, and machine-learning workloads.

Centralized Monitoring, Lineage, and Compliance

Computomic normalized Databricks audit and system data with identity, PAM, catalog, entitlement, and GRC information. Authentication, administrative activity, policy changes, governed-tag modifications, service-principal usage, external sharing, failed access, privileged operations, and AI endpoint activity were enriched with data owner, classification, jurisdiction, legal entity, access-request ID, PAM ticket, purpose of use, and upstream/downstream lineage.

This allowed the security operations center to distinguish approved activity from bulk reads, cross-border access, expired privileged sessions, policy weakening, unapproved sharing, and sensitive-data access from unauthorized environments. Lineage became an active control: when a source field was classified, downstream tables, reports, extracts, features, models, and AI applications were assessed for unmasked copies and regulatory impact.

Computomic Accelerators and FDE Differentiation

SECURITY ATTRIBUTES MIGRATOR
Extracted legacy users, groups, roles, grants, labels, masks, and ownership; generated proposed Unity Catalog groups, tags, privileges, and exception reports.

POLICY TRANSLATION ENGINE
Converted banking policy statements into structured rules and generated Databricks SQL, Terraform definitions, automated tests, and control documentation.

ENTITLEMENT RECONCILIATION ENGINE
Compared approved access, identity-group membership, Unity Catalog grants, exceptions, and

actual usage to detect stale, excessive, direct, or unauthorized permissions.

POLICY TEST HARNESS
Validated authorized access, regional restrictions, PII masking, workspace boundaries, expired privileged access, and direct-grant violations before deployment.

Computomic's Forward Deployed Engineer Model
Computomic embedded multidisciplinary FDE pods within the bank's security, risk, compliance, platform, and business-domain teams. Each pod combined Databricks and Unity Catalog architects, security engineers, cloud and network specialists, DevSecOps practitioners, data engineers, and banking-domain experts.

The same team reverse-engineered entitlements, translated policy language, implemented infrastructure, built security functions, integrated enterprise systems, tested real production scenarios, supported cutover, and transferred capabilities to the bank. This eliminated the gap between advisory architecture and executable controls. Computomic was accountable for measurable outcomes: fewer direct grants, broader sensitive-data coverage, faster provisioning, reduced exceptions, shorter investigation cycles, and automated regulatory evidence.

Conclusion

Computomic transformed a fragmented collection of security products and manual governance processes into a unified enterprise control framework on Databricks. Unity Catalog provided the foundation, while Computomic connected enterprise identities, classifications, approvals, policies, lineage, monitoring, and evidence into an operational platform. Through its FDE model and reusable accelerators, Computomic delivered a scalable architecture capable of protecting data, analytics, machine learning, and AI across a large global bank.

Computomic

www.computomic.ai