Catch what rules miss.

Explainable ML fraud scoring for banks and credit unions.

Rule-based fraud systems were built for a world where fraud patterns were predictable. That world no longer exists.

THE CHALLENGE

Rules-based fraud tools are static. Once fraudsters learn the rules, they route around them. Fraud patterns evolve daily. Compliance teams need not just alerts, but explanations they can defend in an exam. The result: a system that generates too many false positives to action, and misses the sophisticated attacks that cause the most damage.

OUR SOLUTION

BankGuard is a five-engine intelligence platform: Fraud Detection, Risk Scoring, Compliance Monitoring, Customer Retention, and Predictive Analytics. The Fraud Detection engine runs a four-component ensemble model — explainable rules, temporal behavior analysis, supervised ML scoring, and heterogeneous graph analysis — returning a score, a risk level, and a plain-English reason list your team can act on and defend.

Explainable by Design

Every score includes specific reasons: new device, amount anomaly, high-risk country, mule account linkage. Defensible for examiners and BSA officers.

Runs on Your Infrastructure

BankGuard connects to your SQL Server or Oracle transaction database. No data lake migration. No rip-and-replace. Your DBA approves the read-only connection.

Five Engines, One Platform

Fraud detection, risk scoring, AML compliance monitoring, customer retention analytics, and predictive cash flow — in a single deployment.

BUILT FOR

Banks · Credit Unions · Community Banks · BSA / AML Officers · Risk & Compliance Teams · CFOs & Finance Leadership · Core Banking Teams

See BankGuard score your own transactions.

30-day proof of value · your data · no long-term commitment required

Request a POV

Why Rule-Based Fraud Detection Is No Longer Sufficient

Fraud has become an adaptive adversary. Financial institutions that rely exclusively on static rule sets are increasingly exposed to sophisticated attack patterns that rules, by definition, cannot anticipate.

< 1%
FRAUD CAUGHT BY RULE-BASED SYSTEMS IN BANKGUARD VALIDATION TESTING

5.7×
LARGER THAN AVERAGE: MEDIAN FRAUD TRANSACTION AMOUNT VS. LEGITIMATE

98%
OF FRAUD EVENTS DRAIN THE ORIGIN ACCOUNT TO ZERO — INVISIBLE TO AMOUNT-ONLY RULES

Industry Context

The economics of fraud have shifted. Modern fraud operations are sophisticated, automated, and adaptive. Mule account networks — coordinated groups of accounts used to move and launder stolen funds — operate faster than rule update cycles. Rules-based systems that flag on transaction amount or MCC code are easily circumvented by breaking transactions into smaller amounts, using new merchant categories, or routing through intermediate accounts.

The result is a system that generates thousands of low-value alerts — keeping fraud operations teams on a treadmill — while the highest-value attacks slip through undetected. Account takeover, synthetic identity fraud, and coordinated mule networks all exploit the gaps that rules leave open.

The BankGuard Approach

BankGuard is a five-engine platform: Fraud Detection, Risk Scoring, Compliance Monitoring, Customer Retention, and Predictive Analytics. The Fraud Detection engine scores every debit transaction across four independent components: an explainable rules engine (25% weight), a temporal behavior scorer that detects deviations from each account's own historical pattern (20%), a supervised ML model (25%), and a heterogeneous graph model that detects device, payee, and merchant linkages to known or suspected mule accounts (30%).

Each component contributes a score and a reason list. The final composite score — 0 to 100 — maps to a risk level and a recommended action: Allow, Step-Up Authentication, Manual Review, or Hold for Review. No transaction is denied automatically; BankGuard is a decision-support layer, not an enforcement oracle.

Explainability as Compliance Infrastructure

Regulatory expectations around model explainability in financial services are increasing. Examiners expect financial institutions to demonstrate that algorithmic systems can produce human-readable explanations for adverse actions. BankGuard's design — named components, reason strings, feature vectors — means the basis for every score is readable, loggable, and auditable.

When a BSA officer asks why a transaction was flagged, the answer is already in the record: "New device for this account," "Amount 8.3× account median," "Payee linked to flagged accounts." That is the examiner answer, built into the score output by design.

solutions@answerpoint.com · answerpoint.com ·