

Know the field.

How BankGuard positions against every alternative.

Prospects are rarely evaluating BankGuard in isolation. Understanding what they are already running — and what they have considered — determines which strengths to lead with.

RULES-BASED TOOLS (FIS, FISERV, JACK HENRY)

Already in place at most community institutions. Do not position BankGuard as a replacement — position as additive. BankGuard's graph engine surfaces mule account networks that rules engines are not designed to detect. The wedge: the mule network question. Rules don't model account-to-account linkages. BankGuard does.

ENTERPRISE PLATFORMS (ACTIMIZE, SAS, FEATURESPACE)

Powerful but expensive and slow to implement. Position BankGuard on speed, simplicity, and fit for mid-market. Key wedge: time to first score. Enterprise implementations take 12–18 months. BankGuard POV is 30 days. If they have been "evaluating" Actimize for two years, BankGuard offers results without the commitment.

vs. Rule Engines

Additive, not competitive. Lead with the graph engine. Mule network detection is the capability rules engines explicitly don't have.

vs. Enterprise Platforms

Lead with time-to-value and fit. Enterprise platforms are built for \$50B+ institutions. BankGuard is built for community and regional banks.

vs. In-House ML

Lead with build cost and timeline. 18–36 months to production. BankGuard is 30 days to a scored dataset.

BATTLE CARDS

NICE Actimize: Complex data model, 12–18 month implementation, \$200K–\$500K+/year at community scale. Our wedge: time, fit, cost.

SAS Fraud Management: Requires SAS expertise to configure and maintain, pricing opaque. Our wedge: simplicity, explainability, no talent dependency.

Featurespace / ARIC: Strong ML engine, growing US presence. Our wedge: US community banking focus, SQL Server/Oracle native connectors, direct support.

Hawk AI: Growing in AML/compliance monitoring, less strong on transaction fraud scoring. Our wedge: BankGuard covers both fraud and compliance in a single platform.

"We'll build it ourselves": Almost always a 2–3 year project without a dedicated data science team. Our wedge: 30-day POV proves the concept at a fraction of the build cost.

Lead with the POV. It turns a competitive evaluation into BankGuard vs. the prospect's own data.

That is a competition we win.

Internal Reference

Competitive Landscape: Positioning BankGuard in the Mid-Market Financial Institution Segment

The fraud intelligence market is bifurcated: enterprise platforms for large institutions and rules-based tools for community banks. BankGuard occupies the middle — ML-level capability at community institution scale.

Community

THE UNDER-SERVED SEGMENT:
INSTITUTIONS UNDER \$10B IN ASSETS,
LIMITED ACCESS TO ENTERPRISE-
GRADE FRAUD ML

30 days

POV TIMELINE: COMPRESSES
EVALUATION TO A WINDOW
ENTERPRISE COMPETITORS CANNOT
MATCH

Additive

BANKGUARD'S POSITIONING VS.
EXISTING RULES TOOLS — NOT A
REPLACEMENT, A CAPABILITY
EXTENSION

The Market Gap

The fraud intelligence software market is effectively segmented by institution size. Institutions above \$50B in assets have access to enterprise platforms — NICE Actimize, SAS, Featurespace — with the implementation budgets, internal technical teams, and data infrastructure to deploy them. Institutions below \$1B in assets typically rely on core banking rule sets embedded in FIS, Fiserv, or Jack Henry.

The gap is the \$1B–\$10B segment: large enough to have meaningful fraud exposure, too small to justify enterprise platform implementation, but outgrowing what core banking rule sets can address. BankGuard is designed for this segment.

The Explainability Differentiator

One capability that consistently differentiates BankGuard in competitive discussions is explainability. Enterprise platforms produce scores; BankGuard produces scores with reasons. For institutions where BSA officers and examiners need to understand and document why a transaction was flagged, the reason list is not a nice-to-have — it is a compliance requirement.

Competitors who produce a score without a readable explanation require the institution to reverse-engineer the decision for the examiner. BankGuard builds the explanation into the score output by design.

Why the 30-Day POV Wins Competitive Evaluations

In direct competitive evaluations, the BankGuard POV structure has a structural advantage: it compresses the evaluation timeline to a point where no enterprise competitor can match it. A prospect evaluating NICE Actimize and BankGuard simultaneously is looking at a 12–18 month Actimize implementation on one side and a 30-day BankGuard scored result on the other.

Even if Actimize is ultimately the better long-term fit for a large institution, BankGuard consistently wins the "show me results first" competition. For community institutions where fraud problems are live and growing, waiting 18 months is not an

acceptable answer.