Agentic Fraud Alert Triage

Your fraud tools generate alerts. Strike48 makes them manageable. AI agents process, clear, and prioritize fraud alerts across your existing stack—so analysts focus on the cases that actually matter.

The Challenge: Your Fraud Team Is Buried in Alerts.

Your fraud detection tools work. The problem is what happens after they fire:

- **Analysts buried in low-value triage.** Teams spend their days pulling context from identity systems, authentication logs, device intelligence feeds, and transaction engines—just to determine whether each alert is worth investigating. Most aren't.
- **False positives you can't reduce without increasing risk.** Tightening rules generates more noise. Loosening them lets threats through. There's no good manual answer.
- **Customers paying the price.** Overly restrictive controls mean unnecessary friction, blocked transactions, and degraded experience for legitimate users.

The Solution: An Agentic Layer on Top of Your Existing Tools.

Strike48 sits on top of the tools and rules you already run and applies AI agents to the alert volume problem those tools haven't solved.

- **Shrink the queue.** Agents correlate signals in parallel and clear alerts that don't warrant human attention.
- **Separate real threats from noise.** Agents evaluate alerts across identity, device, authentication, and behavioral dimensions simultaneously—without loosening your detection controls.
- **Assemble the case, not just the score.** When an alert does reach an analyst, context is already pulled and the investigation is halfway done.

How It Works

Agents query your existing log sources in parallel—no manual stitching. Each alert flows through three layers:

Multi-phase triage workflow:

Alert intake from multiple detection systems, parallel correlation across identity elements, geolocation and endpoint verification, risk scoring and routing.

Weighted indicator system:

Prioritized scoring for high-impact patterns, correlation of lower-risk signals, dynamic thresholds by account type, and automatic escalation for critical indicators.

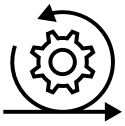

Specialized handling paths:

Different processes by risk level, region-specific regulatory handling, special routing for high-value accounts, and human expert integration for complex cases.

Why Workflow + Agents

The workflow + agent architecture is purpose-built for the alert triage problem:

Orchestrated analysis pipelines allow Strike48 to:

- Execute parallel correlation checks to minimize processing time
- Maintain state across long-running investigations
- Route cases based on dynamic decision points
- Integrate both automated and human-assisted steps

AI-powered intelligence gives agents the ability to:

- Apply NLP to unstructured data analysis
- Contextually understand fraud patterns and risk indicators
- Dynamically formulate queries based on emerging patterns
- Synthesize complex technical findings into human-readable case summaries

Works With Your Existing Stack

- Transaction monitoring platforms
- Identity verification and device intelligence services
- Authentication and access management systems
- Your existing SIEM and log infrastructure

Unified query language across sources. ***No rip-and-replace. No migration.***

Reduce fraud operations cost

Fewer human hours on rote triage. Keep your current tooling but make it operationally sustainable.

Improve customer experience

Cut false-positive friction without loosening detection controls.

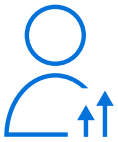

Scale without scaling headcount

Expand fraud coverage with agentic teams while focusing human analysts on high-stakes investigations.

Enterprise-ready

- Deployed in Fortune 500 financial institutions
- 10+ years of platform foundation hardening
- Petabyte-scale storage, billions of events/day ingestion
- Cloud, on-prem, or hybrid deployment
- LLM agnostic—swap models without re-architecting
- Your data never trains our models
- RBAC, tenant isolation, SOC 2-ready
- Full audit trails on every agent action and decision path