

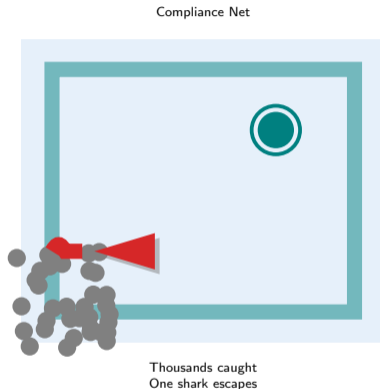
# Why does fighting financial crime cost more than the crime itself – and still miss the worst offenders?

## The compliance paradox:

- Global AML compliance spending exceeds hundreds of billions annually
- Less than one percent of illicit financial flows are intercepted
- Transaction monitoring systems generate thousands of alerts daily
- False positive rates above ninety-five percent waste analyst time
- The few genuine criminals who matter most slip through undetected

## Why the system fails:

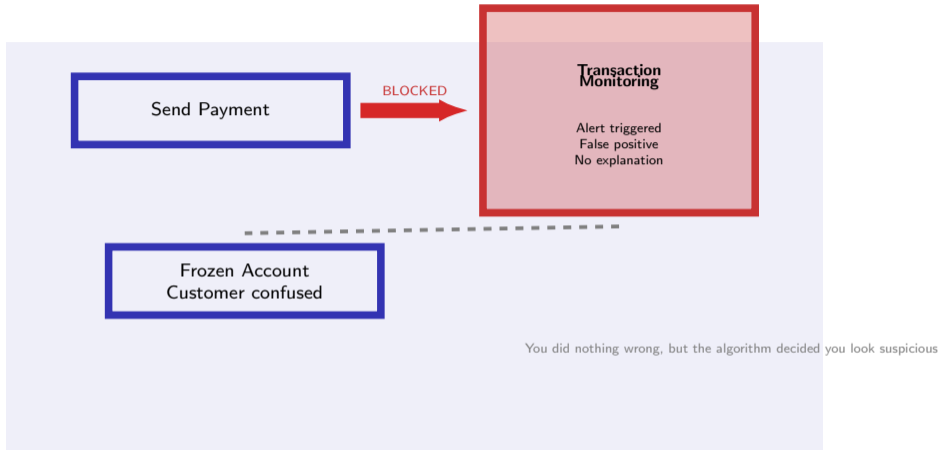
- Sophisticated launderers adapt faster than rule-based systems update
- Volume overwhelms human analysts: alert fatigue causes real cases to be missed
- De-risking pushes transactions underground rather than eliminating them
- Compliance burden falls hardest on small institutions least able to afford it



## Core Insight

Compliance costs billions and catches almost everything, but the few things it misses cause catastrophic damage, and over-compliance harms innocent people.

# Have you ever had a payment blocked or an account frozen for no reason you understood?



## Why This Matters

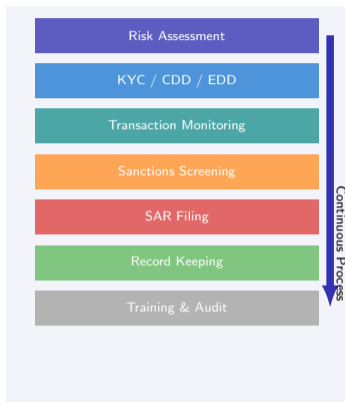
When compliance systems prioritize avoiding false negatives over customer experience, innocent people bear the cost of a system optimized for edge cases.

**Compliance is invisible when it works and incomprehensible when it fails.**

# What are the main components of an anti-money-laundering compliance program?

## Every regulated institution must maintain:

- 1 **Risk assessment:** Identify customer, geographic, product, and channel risks
- 2 **Customer due diligence:** Verify identity, understand business relationship, assess risk level
- 3 **Enhanced due diligence:** Extra scrutiny for high-risk customers including politically exposed persons
- 4 **Transaction monitoring:** Automated systems detect suspicious patterns in payment flows
- 5 **Sanctions screening:** Real-time checks against prohibited persons and entities
- 6 **Suspicious activity reporting:** File reports with authorities when suspicion is raised
- 7 **Record keeping:** Retain all customer data and transaction records for minimum five years
- 8 **Training and testing:** Staff must recognize and report suspicious activity
- 9 **Independent audit:** Regular review of program effectiveness



## Core Insight

An AML program is only as strong as its weakest component. Regulators test all nine elements during examinations.

**The framework is layered: each component provides defense in depth against money laundering.**

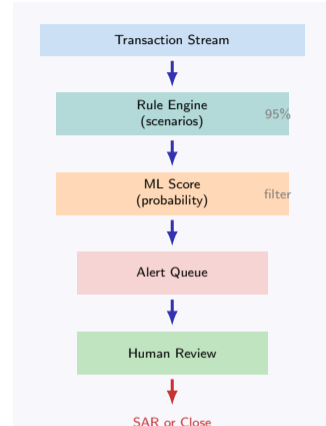
# How does a transaction monitoring system decide which payments to flag?

## Detection logic operates in stages:

- 1 **Enrichment:** Add customer profile, historical behavior, geographic data to each transaction
- 2 **Rule matching:** Check against scenarios like structuring, rapid movement, threshold breaches
- 3 **Threshold comparison:** Flag if amount or frequency exceeds predefined limits
- 4 **Risk scoring:** Apply machine learning model to calculate suspicion probability
- 5 **Alert generation:** Transactions above combined threshold enter investigation queue
- 6 **Analyst review:** Human examines context, customer history, business rationale
- 7 **Disposition:** Analyst closes as false positive or escalates to suspicious activity report

## Key challenge:

- Rules alone generate ninety-five to ninety-nine percent false positives
- Machine learning reduces noise but requires explainability for regulators
- Balance: catch all true crime without overwhelming human analysts



## Core Insight

Transaction monitoring is a funnel: millions of payments compress to thousands of alerts, which human analysts triage to hundreds of suspicious activity reports.

Each stage filters legitimate activity while preserving signals of genuine crime.

# How are rule-based and ML-based financial crime detection systems architected?

## Rule-based architecture:

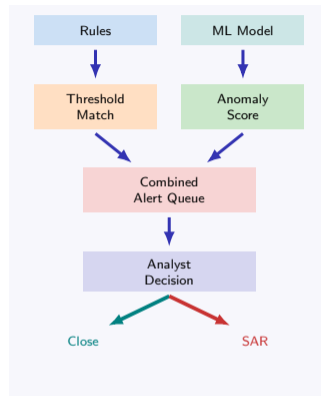
- Hard-coded scenarios: transactions just below reporting threshold, rapid movement, geographic risk
- Static thresholds that do not adapt to customer behavior
- Transparent and auditable but easy for criminals to circumvent
- Cannot detect novel patterns not anticipated by rule designers

## Machine learning architecture:

- Supervised models trained on labeled suspicious activity reports and cleared alerts
- Unsupervised models detect anomalies without prior examples
- Features: transaction velocity, counterparty networks, time-of-day patterns, geographic anomalies
- Continuously retrained as new crime patterns emerge

## Hybrid architecture (most common):

- Rules enforce regulatory mandates and known typologies
- Machine learning scores everything else to reduce false positives
- Human analyst makes final disposition with both signals available



## Core Insight

Hybrid systems leverage rule transparency for regulatory compliance and machine learning precision for operational efficiency.

**The best systems combine deterministic rules with probabilistic models and human judgment.**

# What happens when a compliance system generates so many false positives that analysts stop paying attention?

## The alert fatigue cycle:


- Rules tuned too sensitively generate thousands of daily alerts
- Analysts have thirty to sixty minutes per alert but hundreds in queue
- Pressure to clear backlog incentivizes rubber-stamping
- True positives are indistinguishable from noise at volume
- Real suspicious activity is closed without investigation

## Consequences of alert fatigue:

- Criminal money flows through undetected despite triggering alerts
- Regulatory examinations find missed cases in backlog
- Enforcement actions cite inadequate investigation quality
- Institution faces fines despite having monitoring system in place

## Why this is hard to fix:

- Lowering alert volume requires regulatory approval
- Regulators fear false negatives more than false positives
- Technology vendors over-promise alert reduction without delivering



## Core Insight

Volume is the enemy of quality. High false positive rates do not just waste resources, they actively reduce the detection of real crime.


# Where is financial crime detection most effective and where does it fail?

## Detection works best for:

- Structuring and threshold evasion: automated rules catch this reliably
- Sanctions violations: name-matching technology is mature and accurate
- Large single-institution flows: transaction monitoring sees complete picture
- Known typologies: rules encode past enforcement case patterns

## Detection fails for:

- Layering across multiple institutions: no single bank sees the full pattern
- Small amounts below monitoring thresholds: terrorist financing operates here
- Novel techniques: machine learning helps but always lags criminal innovation
- Trade-based laundering: invoices and goods flows obscure financial crime
- Professional money laundering networks: sophisticated actors avoid known red flags



Detection rates drop

sharply for cross-institutional schemes and novel techniques that evade pattern matching.

## Core Insight

Detection systems excel at catching clumsy criminals and known patterns, but sophisticated launderers exploit the gaps between institutions and regulations.

**The system is optimized for visible threats, leaving invisible ones undetected.**

# Who bears the cost of compliance and who is inadvertently excluded by it?

## Who pays for compliance:

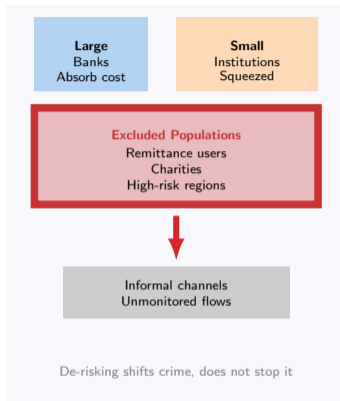
- Large banks: hundreds of millions annually for staff, systems, penalties
- Small banks and credit unions: compliance cost per customer is higher, squeezing margins
- Fintech startups: regulatory burden creates barrier to entry
- Customers: costs passed through as fees and reduced access

## Who gets excluded:

- **De-risking targets:** money transfer operators, charities in conflict zones, cannabis businesses, entire high-risk countries
- **Reason:** institutions exit relationships to avoid compliance cost and regulatory risk
- **Consequence:** excluded populations pushed to informal channels that are harder to monitor
- **Irony:** de-risking intended to reduce crime risk actually increases it by driving flows underground

## The equity problem:

- Compliance is regressive: costs fall hardest on those least able to bear them
- Immigrants sending remittances pay highest fees after correspondent banking withdrawal



## Core Insight

Compliance is meant to protect the financial system, but over-compliance harms vulnerable populations and pushes transactions into channels where crime is harder to detect.

# Three questions to evaluate a financial crime program's real-world effectiveness

## The AML Effectiveness Test:

### 1 What is the false positive rate and how does it affect analyst behavior?

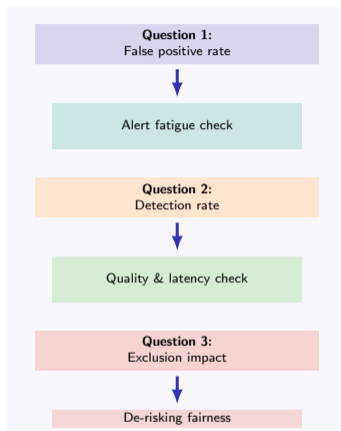
- Calculate: alerts investigated divided by alerts closed without SAR filing
- Red flag: rate above ninety-five percent indicates alert fatigue risk
- Check: are analysts spending adequate time per alert or rushing to clear backlog

### 2 What is the detection rate for genuinely suspicious activity?

- Reverse test: take known enforcement cases and ask if your rules would have caught them
- Check quality: what percentage of filed SARs lead to regulatory feedback or action
- Measure latency: time from first suspicious transaction to SAR filing

### 3 Does the system inadvertently exclude legitimate users from financial services?

- Track: account closures and relationship exits by customer risk segment
- Investigate: are de-risking decisions proportionate to actual crime risk
- Test fairness: do similar customers receive similar treatment regardless of demographics



## Core Insight

Effectiveness is not measured by alert volume or SAR count, but by catching real crime without overwhelming analysts or excluding innocent customers.

**A good AML program balances detection, efficiency, and fairness.**

**Design a transaction monitoring rule for a specific crime pattern. Estimate the false positive rate. Propose one ML-based improvement.**

- 1 **Choose a crime pattern:** structuring, rapid movement, geographic risk, or trade-based laundering
- 2 **Write a rule:** define the trigger logic using thresholds, time windows, and customer segments
- 3 **Estimate false positives:** consider how many legitimate transactions match your rule
- 4 **Propose ML improvement:** what additional features could reduce false positives while maintaining detection

### Example starter:

- Pattern: structuring (deposits just below reporting threshold)
- Rule: multiple cash deposits between eight thousand and nine thousand nine hundred in rolling seven days
- False positive estimate: small business owners making daily deposits
- ML improvement: exclude customers with consistent deposit patterns over six months, flag only sudden changes

### Deliverable

A one-page rule specification with estimated precision and recall, plus one concrete machine learning feature that improves it.

**The goal is not perfection, but understanding the tradeoffs between coverage and noise.**