The Industry Solved the Wrong Problem

Efficiency became the industry's dominant objective. Adaptability must become the next one.

Opening illustration for Chapter 1: The Industry Solved the Wrong Problem.

EXECUTIVE INSIGHT

For more than two decades, financial crime compliance has pursued a remarkably consistent objective: reduce false positives, process alerts faster and lower the operational cost of control. Those goals were understandable, and they delivered meaningful improvements. Yet they also encouraged the industry to confuse efficiency with effectiveness. A quieter monitoring programme is not necessarily a safer one. The central question is no longer whether controls operate efficiently, but whether they continue to reflect the risks that exist today.

Figure 1.1: The Industry Solved the Wrong Problem - key framework.

The Wrong North Star

For more than two decades, the financial crime industry has organised much of its transformation agenda around operational efficiency. The language has become familiar: lower false positives, higher investigator productivity, faster case resolution and reduced cost per alert. These measures appear in vendor demonstrations, board reporting packs and annual strategic plans because they are concrete, measurable and easy to explain. They also respond to a real pressure. As digital payments and cross-border commerce expanded, compliance functions had to process more activity without allowing operational cost to grow without limit.

The difficulty is that a metric can become so dominant that it begins to distort the objective it was originally intended to support. False-positive reduction matters only if the programme continues to identify the right risk. Faster investigations matter only if investigators are looking at the activity that best represents the institution's exposure. Lower cost matters only if the institution has not quietly reduced its sensitivity to the forms of financial crime that are emerging around it.

A highly efficient monitoring programme can still become less effective. It may produce fewer alerts, meet service-level targets and show steady productivity while the assumptions behind its detection logic become progressively outdated. The dashboard can remain green at precisely the moment governance should be asking whether the programme still understands the risk environment it is supposed to protect.

Technology Advanced. Governance Often Did Not.

The industry's technological progress has been significant. Rules engines have become more configurable. Machine-learning models have improved segmentation and prioritisation. Cloud

infrastructure allows institutions to process vast datasets. Screening tools incorporate richer lists and better matching. Workflow platforms have improved case management and auditability. These advances have delivered real benefits, and no serious argument for adaptive compliance should understate them.

Yet regulatory findings and industry experience continue to point to the same underlying problem: the technology may improve while the operating model around it remains largely periodic. Scenarios are reviewed after scheduled intervals. Thresholds are adjusted after evidence becomes compelling. Risk assessments catch up to business change rather than anticipating it. Product launches alter customer behaviour before monitoring assumptions are reassessed. Data quality problems remain tolerated because ownership is fragmented. Governance identifies deterioration only after internal audit, a backlog or an external examiner forces the issue.

The issue is therefore not that rules, models or platforms are unimportant. They remain essential. The issue is that technology is a tool for expressing an institution's understanding of risk, and that understanding has a shorter half-life than it once did. If governance does not continuously test whether assumptions remain valid, even sophisticated controls can become outdated with impressive efficiency.

Efficiency Is Not Effectiveness

Efficiency asks whether the organisation is doing the work with less effort. Effectiveness asks whether the work is producing the intended outcome. In financial crime compliance, the intended outcome is not simply fewer alerts or faster closure. It is the ability to detect, prevent, escalate and explain financial crime risk as that risk changes.

The distinction sounds obvious, but it has significant implications. A scenario that generates fewer alerts may have been successfully tuned, or it may have become less sensitive. A model that performs well against historical validation data may still be poorly aligned to a new product or customer segment. A risk rating methodology may produce consistent outputs while failing to reflect how a market has changed. In each case, the institution can point to operational evidence that the process is working, while the more important question remains unresolved: is the control still describing reality?

This is why the industry's next stage cannot be defined by efficiency alone. The next stage requires governance capable of asking whether controls remain true to the environment. It requires a willingness to treat every model, rule and methodology as a hypothesis that must be continuously tested against changing facts.

The Real Problem Is Adaptation

Financial crime is adaptive. Criminal networks experiment, learn and move rapidly across products, jurisdictions and payment rails. Fraud typologies can spread internationally in days. State-sponsored actors exploit legitimate digital platforms. Sanctions exposure can change overnight. Customers adopt new products, establish new counterparties and behave in ways that historical baselines could not anticipate.

The industry's defining challenge is therefore not simply building better controls. It is building governance capable of recognising when controls no longer describe the world. That is an organisational capability, not merely a technical one. It depends on information moving between

teams, assumptions being challenged, performance being continuously interpreted and leadership treating change as a normal condition rather than an exceptional event.

The next era of compliance begins with a different question. Instead of asking only how to make controls more efficient, organisations must ask how they know those controls still reflect today's risks. That question changes the role of governance. It turns monitoring from a static system into a continuously tested hypothesis about how risk appears in the institution.

From Principle to Practice

The strategic argument becomes operational only when an institution changes the unit of management. The unit cannot be the alert alone. It must be the control hypothesis: the risk the control is intended to detect, the population to which it applies, the evidence that it works, and the conditions that would make it obsolete.

A practical effectiveness program does not begin by tuning thresholds. It begins by building a reliable inventory and asking whether every control still has a defensible purpose. The institution then separates operational burden from risk coverage so that it can reduce noise without treating alert reduction as evidence of success.

Figure 1.2 · The Effectiveness Review Cycle

Step 1: Build a control inventory that can be challenged

At minimum, the inventory should identify the control owner, risk typology, customer and product scope, data inputs, logic or model version, operational output, known limitations, dependencies and the most recent evidence of effectiveness. This is more demanding than a technical catalogue. It connects the control to a specific risk statement and makes assumptions visible.

The inventory should be maintained as a governed record rather than a one-time transformation deliverable. Product changes, data changes, emerging typologies and enforcement themes should trigger review. A control that cannot be linked to a current risk should not continue indefinitely simply because it has always existed.

Step 2: Measure coverage before optimizing workload

Coverage assessment asks whether material typologies, customer segments, products and payment corridors are addressed by an appropriate combination of preventive and detective controls. It should distinguish between nominal coverage - a scenario is said to address a risk - and evidenced coverage, where test results, investigations, intelligence and outcomes demonstrate that the control can identify the behavior in the relevant population.

This assessment should be performed before major alert-reduction targets are set. Otherwise, the organization may remove the very signals needed to understand where coverage is weak.

Step 3: Decide whether to tune, retire, replace or redesign

Tuning is appropriate when the control objective remains valid and noise can be reduced without materially weakening sensitivity. Retirement is appropriate when the risk is covered more effectively elsewhere or the scenario no longer addresses a relevant exposure. Replacement is appropriate when the current technology or data cannot support the control objective. Redesign is required when the risk itself has changed or when the control was built around an assumption that no longer holds.

Each decision should be recorded with the evidence considered, the expected effect on risk coverage and customer impact, the approver, and the monitoring period after implementation. This converts tuning from a technical exercise into a governed risk decision.

Working Template

Field	Purpose	Minimum evidence
Risk typology and objective	Defines what the control is intended to identify or prevent.	Current risk assessment, typology library and regulatory rationale.
Population and product scope	Shows where the control applies and where it does not.	Customer segments, entities, corridors, products and exclusions.
Data and dependencies	Makes upstream fragility visible.	Fields, lineage, latency, quality rules and known gaps.
Logic / model version	Creates traceability to the deployed control.	Rule, features, threshold, code version and approval record.
Outcome evidence	Demonstrates whether the control produces useful risk information.	Investigations, filings, confirmed events, quality testing and back-testing.
Known limitations	Prevents false confidence.	Documented blind spots, assumptions and compensating controls.

Control Prioritization Matrix

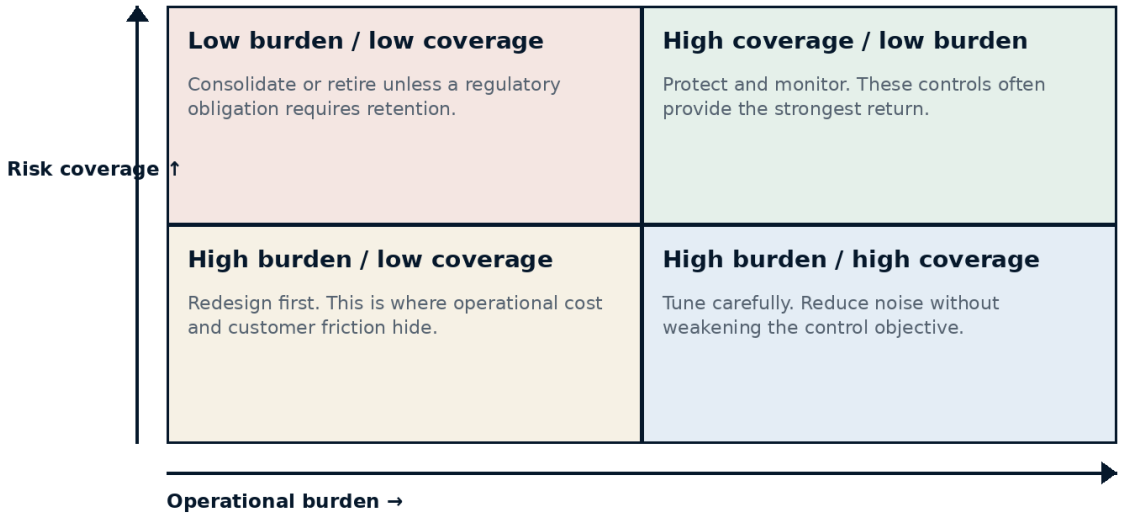

Figure 1.3 · Control Prioritization Matrix

Decision Framework

Condition	Primary action	Governance question
High coverage / high noise	Tune	Can burden be reduced while preserving demonstrated detection capability?
Low coverage / high noise	Redesign or replace	Why is the institution paying for activity that does not address material risk?
Low coverage / low noise	Retire or consolidate	Is the control retained only because it is familiar?
High coverage / manageable burden	Protect and monitor	What changes in the environment would threaten performance?

Executive Case Study

Transaction-monitoring modernization

A global payments business had invested heavily in alert reduction. Productivity improved, but executives still could not explain whether the program covered the most important typologies across customer segments and corridors. High-volume scenarios received the most attention because their operational cost was visible; low-volume coverage gaps received less attention because they did not create backlogs.

Implementation

The organization built a control inventory and mapped every scenario and model to a current typology library. Outcomes were segmented by product, geography and customer type. A cross-functional review then placed each control in a coverage-versus-burden matrix and assigned one of four decisions: tune, retire, replace or redesign.

Governance

Changes were approved through a monthly effectiveness forum with Compliance accountable, Data and Technology responsible for implementation, Operations consulted on workflow impact, and independent testing responsible for post-change verification. Each material change had a defined monitoring window and rollback criteria.

Outcome and lesson

The most important result was not a single percentage reduction. It was a better decision system. The organization could show where burden was reduced, where coverage was strengthened and why each change was defensible. The lesson is that modernization succeeds when the objective is explicit: improve detection effectiveness and customer outcomes together.

Common Failure Modes

- Starting with a false-positive target before establishing coverage.
- Allowing technology teams to tune controls without a documented risk decision.
- Treating true-positive counts as a complete measure of effectiveness.
- Keeping legacy scenarios because ownership for retirement is unclear.
- Failing to monitor the effect of a change after deployment.

A Practical 30 / 60 / 90-Day Plan

Period	Actions and evidence
First 30 days	Create the scenario and model inventory; map material typologies; identify the five highest-burden controls; agree the effectiveness forum and decision record.
Days 31-60	Complete segmented outcome analysis; assess coverage gaps; select two tuning candidates and one redesign candidate; define post-change measures and rollback criteria.
Days 61-90	Implement the first changes; perform independent outcome testing; update the board dashboard to show coverage, burden and change decisions; set the next quarterly review cycle.

Implementation checkpoint

At the end of the first ninety days, leadership should be able to point to at least one changed decision, one implemented control or governance improvement, and evidence that the result was tested. Activity alone is not proof of adaptation.

Executive Reflection

The industry did not fail by pursuing efficiency. It failed only when efficiency became the destination. The future belongs to organisations that can improve efficiency without losing sight of effectiveness - and that can adapt before yesterday's assumptions become tomorrow's failures.

Questions for Leaders

- What assumptions underpin the metrics your organization celebrates most?
- If false positives declined significantly, what evidence would prove that risk coverage did not decline with them?
- Which controls are most dependent on assumptions made more than twelve months ago?
- Where does your governance model test whether the environment has changed, not merely whether the control performed?