VIEW POINT

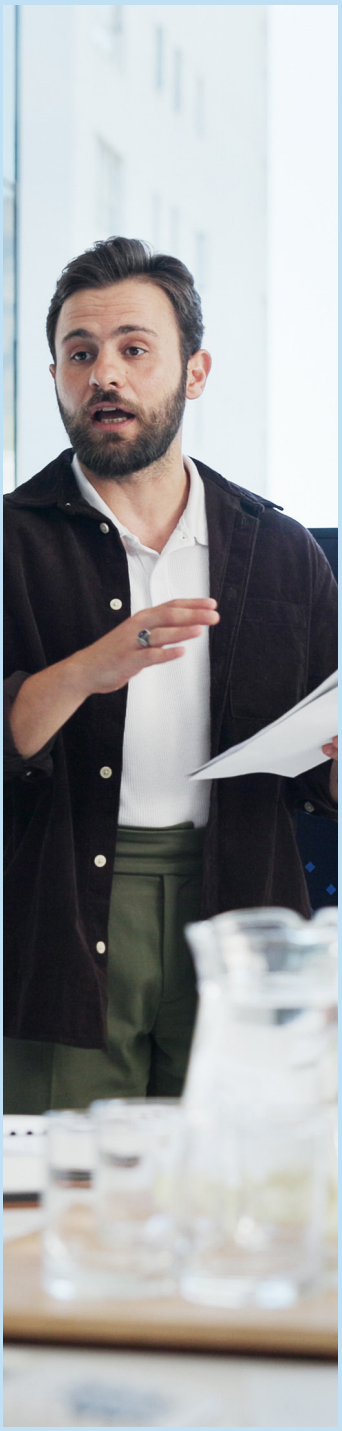

TRANSFORMING FINANCIAL CRIME OPERATIONS WITH PROCESS AUTOMATION AND AGENTIC AI

A STRATEGIC POINT OF VIEW ON HOW BANKS CAN MODERNIZE KYC, AML, INVESTIGATIONS, AND CONTINUOUS RISK MONITORING

Infosys
Navigate your next

Core Message

Financial crime transformation is no longer simply a technology upgrade. It is an operating model shift — from manual, periodic compliance activity to intelligent, always-on risk management.

The winning model uses Agentic AI as the orchestrator, Process Automation as the execution layer, and robust human governance to ensure that speed improves without weakening control.

Executive Summary

Financial institutions are under pressure from every direction. Risk is evolving faster, compliance obligations are expanding, and business leaders are simultaneously expected to deliver faster onboarding, better customer experience, and lower operating costs.

The traditional compliance model cannot keep pace. Too much effort is still spent collecting data, moving cases between teams, reviewing false positives, and repeating point-in-time checks that quickly become outdated.

The right response is not "AI for AI's sake." It is a practical, purposeful reimagining of the financial crime operating model:

Agentic AI should orchestrate end-to-end workflows and support judgment-intensive activities such as document extraction, screening, enrichment, triage, summarization, and monitoring.

Process Automation should handle repeatable operational tasks such as case movement, task execution, approvals, escalations, and structured process steps.

Human oversight should remain central to material decisions — including enhanced due diligence, suspicious activity assessment, and escalations.

When these elements work together, compliance becomes faster, more consistent, more scalable, and more defensible. Rather than operating as a series of disconnected control points, the function begins to perform as a continuous risk intelligence capability.

The Burning Platform: Why Compliance Transformation Can't Wait

$213B
Global AML fines issued once 2008

68%
Of KYC failures traced to manual process gaps

23 days
Average onboarding cycle - global Tier 1banks

4.7X
False positive rate vs. true suspicious

What business leaders should take away:

01

Financial crime transformation is now an operating model issue — not solely a compliance or IT challenge.

02

The greatest opportunity lies in reducing manual effort and false-positive burden while improving the quality of analytical judgment.

03

Perpetual KYC and event-driven monitoring are strategically superior to annual or periodic review cycles.

04

Human-in-the-loop governance is what makes Agentic AI viable in regulated environments.

05

The greatest value comes from phased adoption tied to clear business outcomes — not a big-bang rollout.

Why Change Now?

Most financial crime teams are still working around structural friction. Data is spread across onboarding systems, core banking platforms, transaction monitoring tools, case management applications, and external screening sources. Analysts spend considerable time assembling the full picture before they can begin applying judgment.

This perpetuates a familiar set of business problems: high false-positive queues, slow investigation turnaround, inconsistent case handling, stale customer risk profiles between review cycles, and manual reporting that compounds operational risk.

The business context has also changed materially. Digital channels have increased transaction volumes and velocity. Cross-border activity has expanded. Customer expectations have risen. Regulators now demand stronger evidence, better auditability, and more timely action.

In short, the legacy operating model is being asked to do far more than it was designed for.

That is why this conversation should not be framed as "how do we automate a few tasks?" It should be framed as "how do we redesign financial crime operations so the business can scale safely?"

The Strategic Shift: From Point-in-Time Compliance to Continuous Risk Intelligence

The most significant strategic change is moving away from compliance as a periodic review exercise. Traditional models depend on onboarding reviews, annual refreshes, and manual follow-up when something goes wrong — leaving long gaps during which a customer's risk profile may no longer reflect reality.

A stronger model is event-driven and dynamic. Customer risk should be refreshed when ownership structures change, when adverse media emerges, when sanctions lists are updated, when behavior shifts, or when transaction patterns become unusual. This is the foundation of perpetual KYC and continuous financial crime monitoring.

For business leaders, the benefit extends beyond better control — it enables better allocation of analytical attention. Instead of reviewing everything with uniform intensity, teams can focus on the highest-risk cases while the platform handles monitoring, enrichment, and routing in the background.

Dimension	Traditional Model	Target Model
Risk view	Periodic and backward-looking	Continuous and event-driven
Analyst effort	Heavy focus on data gathering and false positives	Heavy focus on exceptions and judgment

What Changes for the Business

4.1 Onboarding and KYC Become Orchestrated, Not Document-Heavy

Customer onboarding should transition from a manual queue to a structured digital workflow. Identity documents and corporate records can be extracted automatically, key data validated, entities matched across registries, and beneficial ownership assembled — all before an analyst steps in. This shortens the path from intake to decision and improves the quality of the initial review.

4.2 Screening Becomes Continuous, Not a One-Time Check

Sanctions screening, politically exposed person checks, relatives and close associates screening, and adverse media monitoring should not stop at onboarding. They should function as living controls — reacting to new information and triggering review only when material signals arise.

4.3 Level 1 Investigations Become Far More Efficient

Some of the most costly manual effort sits in alert triage and Level 1 investigations. Data gathering, enrichment, risk scoring, red-flag detection, duplicate alert identification, first-draft case summaries, and recommended next actions can all be automated. This is where institutions can eliminate a significant share of low-value manual effort without reducing control quality.

4.4 Reporting and Case Management Become More Standardized

Case handling should follow a guided workflow with clear evidence capture, standard process steps, task routing, and complete audit history. Draft narratives for suspicious activity reports can be prepared automatically — but the final decision must remain with accountable human reviewers.

4.5 Risk Scoring Becomes Dynamic

Customer risk should no longer remain static until the next periodic review. It should be recalculated as new data arrives, so that enhanced due diligence is triggered sooner and investigative effort is concentrated where it matters most.

How the Model Works in Practice

The technology architecture matters, but business leaders do not need to navigate technical detail to understand the logic. The model works because three capabilities reinforce one another.

Agentic AI — the orchestration layer. It directs the flow of work, prioritizes actions, coordinates specialized agents, and determines when cases should be routed, escalated, or reviewed by humans.

Process Automation — the execution layer. It carries out repeatable operational tasks such as case creation, workflow routing, approvals, escalations, service-level tracking, and structured process steps across systems and teams.

Human oversight should remain central to material decisions — including enhanced due diligence, suspicious activity assessment, and escalations.

A simple way to view the architecture is shown below.

Layer	Business Role
Data and input layer	Collects customer, transaction, screening, and external intelligence data.
Integration and events layer	Enables the platform to react when material events occur and synchronizes systems in real time.
Intelligence layer	Uses Agentic AI to orchestrate decision flow, coordinate specialized agents, and prioritize the next best action.
Workflow and automation layer	Uses Process Automation to execute workflow steps, move cases, trigger approvals, and enforce process discipline.

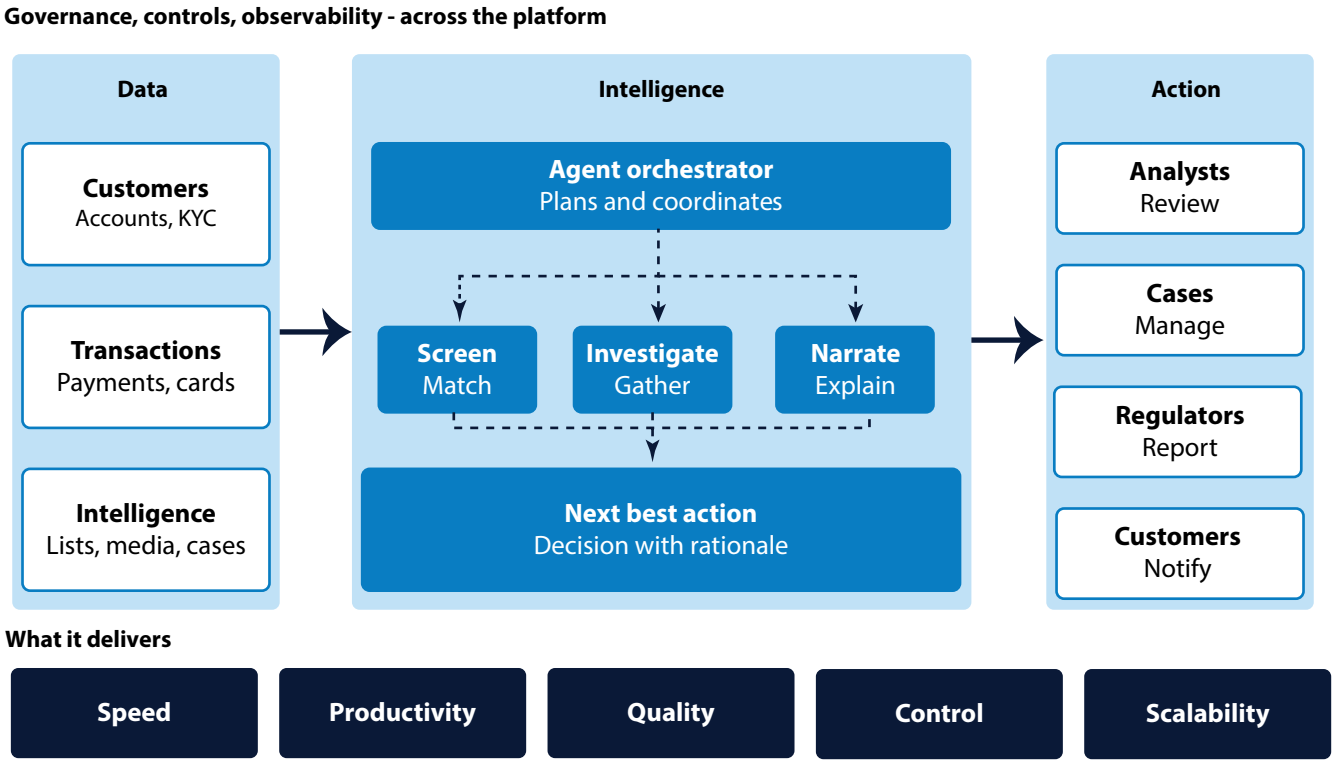

Why Governance Is the Deciding Factor

The strongest financial crime proposition is not full autonomy — it is controlled autonomy. That distinction matters because business leaders need to improve productivity without creating new regulatory exposure.

Any credible Agentic AI implementation in financial crime must meet four tests:

- Humans must retain accountability for material decisions.
- Every action must be traceable through a clear, immutable audit trail.
- AI outputs must be explainable enough for reviewers to understand why a recommendation was made.
- Access and approvals must follow strict role-based controls.

When these foundations are in place, automation becomes easier to trust. Without them, efficiency gains will not be sustainable — because the business will hesitate to rely on the output.

Key governance requirements include:

- Human review for enhanced due diligence, suspicious activity assessment, and escalations
- Immutable audit history for all material decisions and overrides
- Clear evidence, confidence levels, and rationale behind AI-supported recommendations
- Separation of duties across analyst, manager, auditor, and administrator roles
- Feedback loops to continuously improve rules, prompts, thresholds, and models

Where Value Becomes Visible First

High-risk onboarding: A high-risk customer or entity can be screened, enriched, and risk-scored rapidly, with the analyst stepping in only where enhanced due diligence or escalation is required.

Sanctions and watchlist false positives: Context-aware screening can reduce unnecessary reviews caused by name-match variations, spelling differences, and incomplete data.

Perpetual KYC: The platform refreshes customer risk when material changes occur, rather than waiting for annual review cycles.

Level 1 investigation automation: Data gathering, alert prioritization, case creation, and first-draft summaries can be automated, freeing investigators to focus on substantive judgment.

Regulatory reporting support: Structured evidence collection and automated draft preparation can improve the timeliness and consistency of regulatory submissions.

Business Value and Executive Outcomes

The business case is straightforward. This model helps institutions move faster, reduce avoidable manual effort, improve consistency, strengthen audit readiness, and expand risk coverage — without scaling headcount proportionally.

The precise outcomes will vary by organization, but the direction of value is consistent: lower false-positive burden, faster onboarding and investigations, more timely reporting, better use of analyst capacity, and a more current view of customer risk.

Value Area	What It Means for the Business
Speed	Shorter turnaround for onboarding, triage, investigations, and reporting.
Productivity	Less analyst time spent on low-value review and data assembly.
Quality	More consistent evidence gathering, narrative preparation, and case handling.
Control	Better auditability, explainability, and role-based accountability.
Scalability	A stronger ability to absorb growth in volume and complexity without proportional cost growth.

A Pragmatic Adoption Path

This should not be approached as a single large transformation. The better path is phased adoption, starting with use cases that address visible operational pain and deliver measurable value.

A practical sequence begins with onboarding and document intelligence, then expands into screening and alert triage, followed by dynamic risk scoring and perpetual monitoring, and finally scales to enterprise-wide governance, reporting, and controls.

This approach gives business leaders an early value story while building the operating discipline needed for broader, confident adoption.

Phase	Focus
Phase 1	Document intake, data extraction, and workflow discipline
Phase 2	Screening automation, case routing, and alert triage
Phase 3	Dynamic risk scoring and perpetual KYC monitoring
Phase 4	Enterprise controls, regulatory reporting support, and scale

Closing Point of View

Financial crime operations are at an inflection point. The choice is no longer between manual work and unchecked automation. The real choice is whether institutions build a controlled, intelligent operating model — one that lets people focus on judgment while the platform handles the repeatable, high-volume work.

The most credible path forward is governance-first transformation: use Agentic AI to orchestrate the work, use Process Automation to execute repeatable operational tasks, and keep humans accountable for the decisions that matter most.

That is how financial crime teams can become both more efficient and more trusted.

FROM PERIODIC COMPLIANCE TO CONTINUOUS RISK INTELLIGENCE
How financial crime operations evolve from reactive control activity to intelligent, always-on risk management.

The strongest financial crime model is not more manual effort - it is smarter orchestration, faster execution, and better human judgment.

About the Authors

Ankush Jain
Industry Principal, EAIS Consulting – Infosys

Ankush Jain is an Industry Principal at Infosys, leading EAIS Consulting with a sharp focus on Process Excellence, Agentic Automation, and platform-led enterprise transformation. He serves as a trusted advisor to CXO and senior leadership across industries, helping organizations reimagine operating models through AI-driven processes, business architecture, and domain-anchored transformation strategies. Ankush brings a rare blend of strategic foresight and hands-on execution, guiding enterprises from vision to value realization by embedding intelligent agents, data, and process orchestration at scale.

Bhavik Chugh
Digital Solution Specialist, EAIS Consulting – Infosys

Bhavik Chugh is a Digital Solution Specialist at Infosys, driving digital solutioning, agent-enabled system design, and IT transformation initiatives for complex enterprises. Working closely with business and technology leaders, Bhavik plays a critical role in translating strategic intent into executable architectures. With strong grounding in process automation, domain workflows, and platform integration, he bridges the business–IT divide by designing scalable solutions that anchor agentic capabilities into day-to-day operations.

For more information, contact askus@infosys.com

Infosys®
Navigate your next