NICE Actimize

White Paper

# Beyond the Alert Screen: Research-Based Insights on Rebuilding Fraud Investigations for AI

Anurag Mohapatra, Director, Fraud Strategy & Marketing, NICE Actimize
Kushal Edwankar, Product Manager, NICE Actimize

NICE Actimize

# Table of Contents

# Executive Summary

This paper draws on operational interviews with fraud investigation teams at financial institutions across North America and Europe, covering retail and commercial banking operations. The research examined how investigators allocate their time, which systems and processes they depend on and where friction concentrates in the investigation workflow.

Four operational patterns emerged consistently across all institutions interviewed.

1. Investigators spend most of their SLA gathering data that the investigation platform cannot consolidate.
2. The path investigators take through that data varies, resulting in inconsistent investigative outcomes and quality of documentation.
3. Investigation expertise resides largely with individuals rather than within the workflow, making quality heavily dependent on experience and tenure.
4. The investigation workflow retains no memory across alerts or analysts, so context assembled during one review is unavailable to the next.

These four findings are interconnected. Each one amplifies the operational impact of the others.

The final section of this paper examines how institutions can address these constraints through a combination of data orchestration, process and workflow redesign, and AI-assisted investigation. Together, these interventions address different layers of the investigation challenge — from fragmented data and inconsistent workflows to investigative consistency and context continuity.

Agentic AI has an important role to play in that transformation, but the research also makes clear that technology alone is not enough. Meaningful investigation transformation requires institutions to address the underlying data, workflow and policy decisions that shape how investigations are performed.decisions that institutions themselves must address as part of any meaningful investigation transformation effort.

Finding 1

# Investigators Spend More Time Gathering Data Than Analyzing Risk

When fraud investigation teams were asked to identify the most time-consuming activity during alert review, every institution gave the same answer: accessing and reconciling data from systems outside the alert management platform.

A fraud alert provides a starting point. It shows transaction details, triggered rules, risk scores and any enrichment data the alert management platform is configured to display. It does not provide a complete investigative view of the customer, the counterparty or the transactional context. That picture has to be assembled, manually, from systems that were never designed to work together.

Across the institutions interviewed, investigators routinely consult between six and eight systems during a single alert review. The specific systems vary by institution, but the categories are consistent.

In commercial banking, where SLAs run between 30 and 35 minutes, teams report spending 80 to 90 percent of that window outside the alert management platform. Investigators move between core banking systems for transaction history, external verification services for counterparty data, internal databases for relationship context and open-source lookups for payee validation. In retail operations, where SLAs average 10 to 15 minutes, investigators still spend roughly 30 to 35 percent of their review time outside the platform gathering and reconciling data.

The process remains heavily manual. Investigators copy findings into a notepad or Word document as they work, then transcribe those notes into the alert management platform when the review is complete.

One fraud operations manager described it directly: "They have a notepad open on another screen. They fill in the blanks as they go. At the end, they copy it all out and paste it in."

## Share of Alert-Review SLA Spent Outside the Alert Management Platform

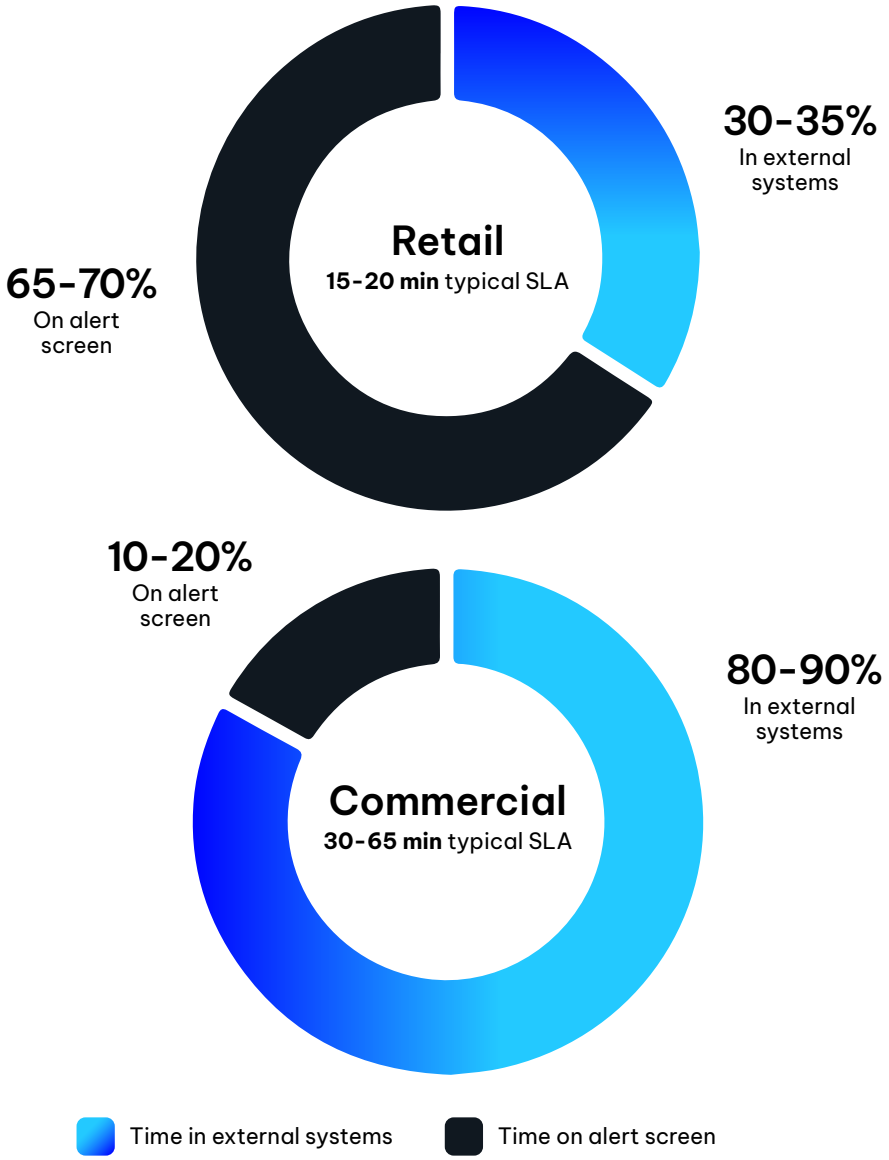


**AI Blueprint:** Agentic AI can help reduce the time investigators spend gathering and reconciling information across disconnected systems By assembling relevant investigative context before the analyst opens the case, AI-assisted workflows can shift investigator time away from manual retrieval and toward analytical decision-making.

Finding 2

# Investigation Quality Depends Too Heavily on Individual Experience

Onboarding a new fraud investigator to full operational proficiency typically takes about six months. The process often begins with four to six weeks of shadowing an experienced analyst, followed by supervised casework before the investigators are expected to manage the full alert queue independently.

Much of that onboarding period is spent learning how to navigate the investigation process – which systems to access, which signals matter, in what sequence and for which fraud typologies. Among the institutions interviewed, only a small number had formal decision matrices mapping alert types to investigation steps and escalation criteria. Most relied on high-level guidance and knowledge gained through shadowing. Once formal oversight ends, investigators tend to adopt investigative habits shaped largely by the analysts who trained them.

**AI Blueprint:** Agentic AI can help reduce variability in how investigators approach alert review by applying consistent investigative logic, workflow sequencing and policy guidance across cases. Analysts still make the final decision, but the information and context presented to them can become more consistent regardless of who handles the investigation.

The operational consequence is inconsistency. Investigation quality becomes partly dependent on which analyst handles the alert. Two investigators with identical training and tenure may approach the same alert differently, surface different signals in a different sequence and sometimes reach different conclusions. The evidence is the same, but the process does not enforce a consistent path to a decision.

This becomes especially visible in scam investigations. Unlike account takeover, which follows relatively structured indicators, scam investigations require analysts to build a narrative that includes the customer's historical banking behavior, the timing and channel of recent contact and the plausibility of the stated payment purpose. Multiple teams described training investigators to ask probing questions rather than follow a script – to take what one operations leader called "a journey" and arrive at a conclusion through reasoning. That capability develops through exposure, at varying rates, and is often difficult to retain when experienced analysts leave. Unlike system access or procedural documentation, investigative judgement does not transfer cleanly.

> " The job of an investigator is to go on a journey through the data and find clues to progress the investigation. Instead, they are going on a journey of collecting and reconciling data."
>
> **Fraud Operations Manager, North American Tier-1 Bank**

Finding 3

# Investigation Context Does Not Carry Forward

At most institutions, the investigation process treats every alert as a new event. It retains no memory of the conversations that preceded it, the context assembled during previous reviews of the same customer or the conclusions reached when the same payee triggered an alert weeks earlier. As a result, investigators often begin reviewing suspicious activity with little visibility into prior investigative work.

A commercial customer making repeated payments to a new payee may trigger multiple alerts across different sessions and analysts. Each analyst opens the alert with no visibility into what colleagues have already established. As a result, customers may be contacted repeatedly, asked the same questions and taken through the same verification steps multiple times. On transactions running into the millions, that experience signals to the customer that the bank does not know them, regardless of how long the relationship has been in place.

Between 30 and 40 percent of retail alerts require an outbound customer call – to verify suspicious activity, gather context the data does not provide, or in scam cases, assess whether the customer is acting under a fraudster's influence. In retail operations, the investigating analyst typically makes that call directly, with the full investigation context live in front of them. In commercial operations, the call goes to a separate customer contact team . What transfers is the notes field: a summary written under time pressure, reflecting what the analyst chose to document. The next reviewer works from that condensed summary rather than the full investigative context or underlying evidence.

In many institutions, the investigator effectively becomes the system's only source of retained context. When that investigator is unavailable, the context is unavailable. When they leave the institution, much of that institutional knowledge leaves with them.

**AI Blueprint:** Agentic AI can help institutions retain and connect context across alerts, customers and payees. By surfacing prior investigative activity, customer interactions and related alert history during case review, AI-assisted workflows can reduce repeated investigative work and improve continuity across analysts and teams.

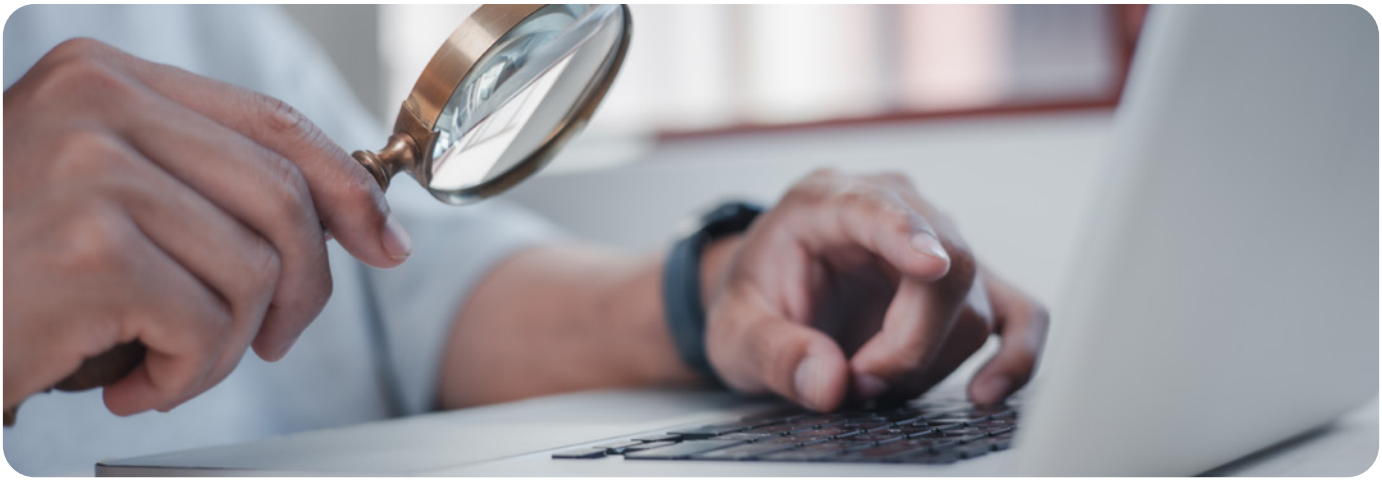

NICE Actimize

Finding 4

# Investigation Workflows Do Not Adapt to Risk Signals

Between 10 and 20 percent of daily alert volume are what operations teams call quick-resolution cases — alerts where the relevant signals are clear within the first minute. The device is recognized, the payee is established and the transaction is consistent with the customer's history. These risk signals indicate that the alert is unlikely to be fraudulent.

Yet investigators are still required to complete the full investigative flow.

Compliance requirements and internal policy mandate the same structured review regardless of how quickly the risk signals resolve. Every external system on the checklist must be consulted. Every section of the notes template must still be completed before the alert can be dispositioned. The SLA clock continues running whether the alert requires meaningful investigation or not.

The second-order effect is less visible but more consequential. Time spent on procedural work for low-risk alerts reduces the time available for investigations that require deeper analytical judgment. An investigator who might otherwise spend 40 minutes on a complex scam case may spend only 20 because 15 minutes were consumed by a quick-resolution alert that required little investigative analysis.

**AI Blueprint:** AI-assisted investigation workflows can help institutions better understand how investigator time is allocated across alert types, workflows and procedural requirements. By identifying which investigative steps contribute meaningful analytical value and which create operational overhead, institutions can make more informed decisions about workflow redesign and investigative prioritization.

# The Compounding Effect

The findings discussed in this paper do not operate independently. Each challenge reinforces the next, creating a compounding cycle of operational inefficiency.

Fragmented data access creates the need for institutional knowledge – the ability to navigate multiple systems, workflows and investigative signals efficiently. That reliance on institutional knowledge introduces inconsistency because the capability is developed through individual experience rather than embedded within the investigation process itself. As variability increases, institutions become more dependent on detailed documentation to preserve investigative context across analysts and teams. Yet even extensive documentation cannot fully compensate for the absence of shared investigative continuity across alerts and investigations.

Every incremental fix addresses one part of the problem without resolving the underlying interconnected workflow constraints. For example, adding analysts increases operational capacity but does not improve per-alert efficiency. Similarly, improving guidance documentation may accelerate onboarding, but it does not reduce the number of systems investigators must navigate during alert review.

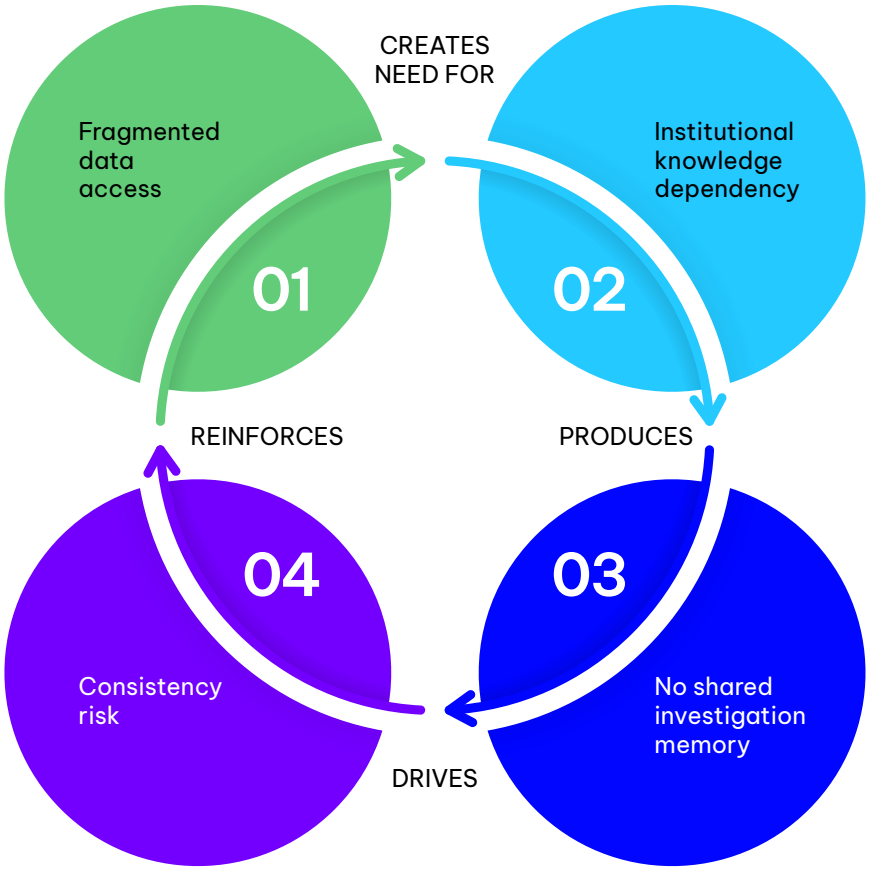

# The Path Forward

The findings in this paper point to three broad areas of intervention: data orchestration, process and workflow redesign, and agentic AI. Each addresses a different layer of the investigation challenge. Some operational constraints stem from fragmented data and workflow design, while others are rooted in investigative consistency, context continuity and decision support. What follows is a framework for aligning the appropriate intervention to each challenge.

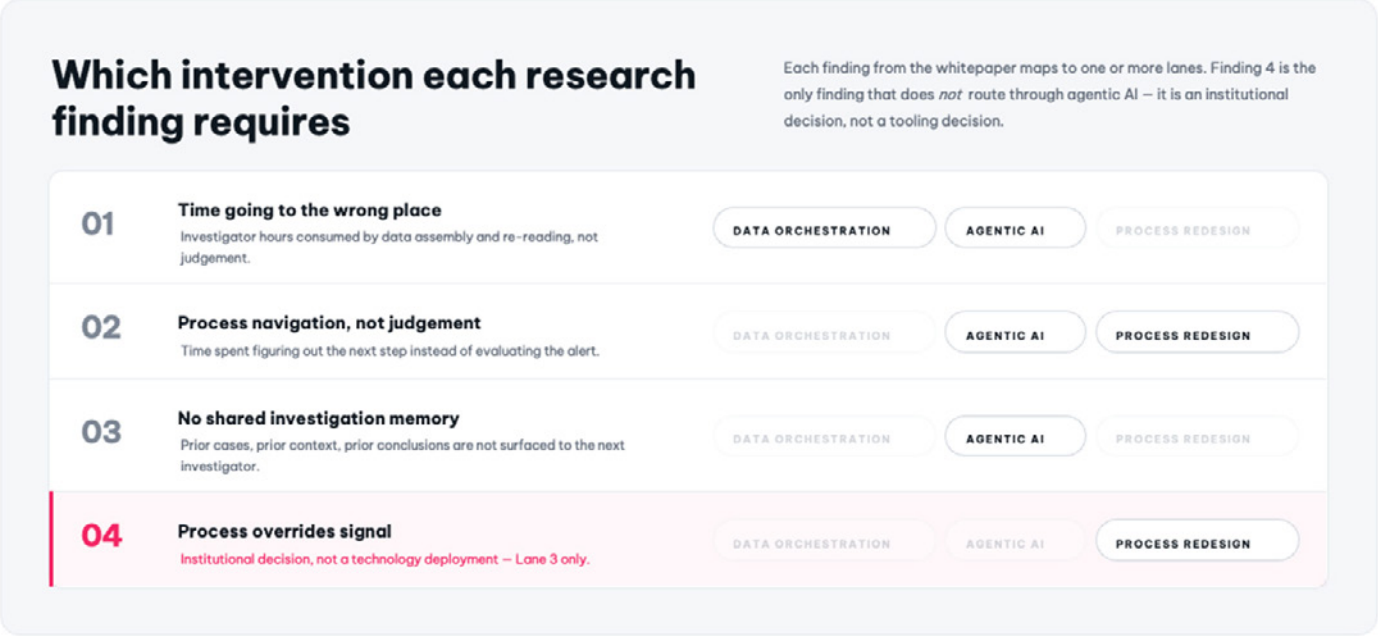


## Data Orchestration and Consolidation

The most immediate constraint investigators face is that the data they need exists but is scattered. Core banking history, counterparty verification, device intelligence and IP context are available, but they are rarely accessible within a single investigative workflow. Consolidating those sources into a unified investigation view is primarily an integration challenge rather than an AI challenge. Before deploying any AI-assisted investigation capabilities, institutions need a reliable and consolidated data foundation. Without it, AI systems operate with incomplete investigative context.

Achieving that consolidation is not straightforward. Connecting core banking platforms, third-party enrichment services, device intelligence feeds and consortium databases into a coherent investigation view requires significant integration effort. Modern investigation platforms are increasingly applying AI to data normalization and source mapping to accelerate the process, but the underlying integration layer remains one of the most resource-intensive components of any investigation transformation program. Institutions that underestimate this step risk deploying AI capabilities against fragmented or incomplete data environments.

## Process and Workflow Redesign

Investigator efficiency is affected by process requirements that mandate full due diligence on every alert regardless of early signal clarity. The research also identifies a consistency problem: most institutions lack a formal decision matrix, leaving investigators to rely on training and individual judgment to determine how to approach each alert type.

Institutions that want to address these gaps need to examine the investigation workflow from the ground up. The question is not simply which tool to deploy. It is whether investigative workflows are designed to distinguish between alerts that require deep analytical review and those that do not. Institutions must determine which procedural steps contribute genuine analytical value and which persist primarily because workflows evolved incrementally over time That distinction matters: a process that treats a quick-resolution alert identically to a complex commercial wire investigation is not operationally neutral. It is a source of wasted capacity.

A formal decision matrix, mapping alert types to investigation steps and escalation criteria, is one output of that examination. It is not a substitute for it. Closing the consistency gap requires deliberate policy decisions about how investigations are structured, not just better tooling layered on top of an unchanged process. Ultimately, this is an operational and governance decision that institutions must make themselves.

## Agentic AI

With a consolidated data layer and a redesigned process in place, agentic AI can address operational gaps that neither integration nor workflow redesign alone can fully resolve: the synthesis, consistency and memory problems at the core of Findings 2, 3 and 4.

An AI agent embedded in the investigation workflow can gather relevant context across all connected systems, summarize key findings, identify related prior alerts involving the same customer or payee and recommend next steps aligned with institutional policy — all before the analyst opens the case. The analyst reviews the agent's output, refines it where needed and makes the final decision. The agent does not make final disposition decisions. It prepares investigators to do so with better information and less time spent retrieving and reconciling context.

Three specific improvements follow from this model.

- **Consistency:** The agent applies the same sequence and the same policy logic to every alert, removing individual habits from the decision path.
- **Shared investigative context:** The agent identifies that a customer has triggered multiple alerts, that a payee has appeared in prior investigations and that a colleague already spoke to the customer last week — connections that might otherwise remain fragmented across systems, alerts and teams.
- **Structured investigation summaries:** The agent generates a consistent documenting the assembled context, the identified signals and recommended next steps. This creates more standardized investigative documentation regardless of which analyst handled the alert or how much time was available during review.

The sequencing of these interventions matters. Institutions that deploy agentic AI before consolidating their data layer will find the agent working with incomplete investigative context. Institutions that deploy AI before re-designing investigation workflows risk making inefficient processes faster rather than fundamentally improving them. The value of agentic AI is highest when the underlying data is reliable and the process is sound. In that environment, AI-assisted investigation not only improves efficiency but also helps raise the baseline consistency and quality across analysts, alerts and shifts.

# Conclusion

The fraud operations teams behind this research are experienced, capable and committed. The tools available to them are more sophisticated than at any point in the history of the discipline. The challenge is neither the people nor the technology in isolation. It is the workflow that connects them.

That workflow was assembled incrementally – system by system, guidance document by guidance document – over years of operational evolution. Each addition addressed a legitimate operational need at the time. Collectively, however, they produced a process where investigator judgment is consistently applied last, after most of the available time has been consumed by work that does not require it.

Progress looks different depending on which operational constraint an institution is trying to address. It may be measured by how much investigator time is spent on analysis rather than data retrieval, the consistency of outcomes across analysts reviewing the same alert type or whether a second analyst begins with established context instead of rebuilding the investigation from scratch. It may also be reflected in how quickly new investigators reach proficiency because the workflow guides decision-making rather than relying on habits developed through experience.

Agentic AI addresses most of these challenges directly. It gathers, synthesizes and surfaces context that currently consumes the majority of investigators' time to assemble manually. It creates the continuity and investigative context the workflow has historically lacked. It helps make investigative performance more consistent, scalable and resilient across teams and workflows. The institutions best positioned to benefit are those that approach transformation deliberately: consolidating their data foundation first, examining their workflows honestly and deploying AI where it addresses a meaningful operational constraint rather than where it appears most attractive on a technology roadmap. In that environment, agentic AI elevates what fraud investigation teams are capable of.


**Anurag Mohapatra,** Director, Fraud Strategy & Marketing, NICE Actimize

Anurag Mohapatra is a financial crime prevention expert with twenty years of experience across fraud, compliance, and trade surveillance. His background spans implementation, pre-sales, product management, and go-to-market strategy, giving him deep expertise in how financial crime solutions are built, bought, deployed, and scaled. He currently leads go-to-market strategy for the fraud portfolio at NICE Actimize.

**Kushal Edwankar,** Product Manager, NICE Actimize

Kushal Edwankar is a Product Manager with 19+ years of experience in Banking and Financial Services, including over 9 years building Enterprise Fraud Prevention products across digital, check, card, and payments channels. He specializes in New Account Fraud, Scams and Mule typology prevention, and has partnered closely with Tier 1 banks across the U.S., Europe, and Asia to drive customer outcomes and product innovation.

NICE Actimize

## Know more. Risk less.

info@niceactimize.com
niceactimize.com/blog
@NICE_actimize
/company/actimize
NICEactimize

### About NICE Actimize

As a global leader in artificial intelligence, platform services, and cloud solutions, NICE Actimize excels in preventing fraud, detecting financial crime, and supporting regulatory compliance. Over 1,000 organizations across more than 70 countries trust NICE Actimize to protect their institutions and safeguard assets throughout the entire customer lifecycle. With NICE Actimize, customers gain deeper insights and mitigate risks. Learn more at www.niceactimize.com