Deloitte.

Together makes progress

Artificial intelligence adoption in Financial Crime Compliance

Planning and priortisation in a fast-moving market

Regulators increasingly expect financial institutions (FIs) to demonstrate meaningful progress in adopting advanced technologies, including artificial intelligence (AI). However, adoption is not without risk, and many institutions are not yet structured to scale such capabilities effectively. The challenge, therefore, is clear: How can compliance and technology teams introduce AI in a way that delivers value, satisfies leadership expectations, and meets regulatory scrutiny while maintaining robust risk management?

The momentum behind AI adoption in financial services continues to accelerate. Regulators across APAC have placed increasing emphasis on AI within their 2026 supervisory agendas, particularly in areas such as fraud detection, risk analytics, and compliance. The global AI in banking market, valued at approximately USD 34.58 billion in 2025[1], is projected to grow significantly over the next decade, reflecting widespread adoption across core banking functions.

In Singapore, this momentum is particularly pronounced. As a leading global financial hub, Singapore faces heightened exposure to increasingly sophisticated financial crime threats. Traditional rule-based systems are no longer sufficient to address the scale and complexity of these risks. As a result, AI has shifted from a theoretical advantage to a strategic necessity for building resilient, future-ready compliance frameworks.

Why AI?

AI is fundamentally reshaping how institutions detect and respond to financial crime. By leveraging advanced analytics and machine learning, firms can identify non-obvious patterns across large, fragmented datasets, enabling earlier and more accurate detection of suspicious activity. Unlike manual processes, AI operates at scale and in near real time, providing continuous monitoring across transactions, channels, and customer behaviours.

From a regulatory perspective, the adoption of AI aligns closely with the expectations of the Monetary Authority of Singapore (MAS). Singapore's regulatory framework emphasises a risk-based approach to anti-money laundering (AML) and countering the financing of terrorism (CFT), requiring institutions to dynamically assess and mitigate risks.

For example, AI enables this by continuously refining customer risk profiles and enhancing transaction monitoring precision. In doing so, it strengthens compliance with MAS guidelines while improving the quality and timeliness of suspicious transaction reporting.

Moreover, AI enhances transparency as well as auditability, and standardises quality and assurance reviews; critical factors in regulatory supervision. Advanced systems can generate detailed audit trails and data-driven explanations, supporting more effective regulatory reviews and inspections. This is particularly relevant as MAS continues to promote responsible AI adoption through initiatives such as the Veritas Initiative, which underscores the importance of fairness, ethics, accountability, and transparency (FEAT) in financial services.

The regulatory benefits are clear: institutions that adopt AI-driven financial crime controls are better positioned to meet compliance obligations, reduce the risk of enforcement actions, and safeguard their reputations. In an environment where regulatory scrutiny is intensifying, proactive adoption of AI can serve as a differentiator, demonstrating both robustness in risk management and alignment with supervisory expectations.

With the Singapore Government having announced its new National AI Impact Program (NAIIP), which will train over 100,000 people in AI literacy, the push for AI adoption has never been greater[2]. At the same time the March 2026 MAS AI Risk Toolkit[3] shows that careful application is of paramount importance in terms or regulatory compliance.

[1] *Artificial Intelligence (AI) in Banking Market, Precedence Research, January 2026*
[2] *Singapore to train 100,000 AI-savvy workers by 2029, The Straits Times, March 2026*
[3] *MindForge AI Risk Management Operationalisation Handbook, Monetary Authority of Singapore, January 2026*

As a matter of priority, FIs can focus on the following areas:

Incremental process automation and data reconciliation

What might the ideal target operating model look like in a transaction monitoring investigation context? One could envisage a fully automated framework: agentic AI driving data collection and transformation, machine learning models analysing transaction behaviour, and generative AI producing case narratives and even recommending case outcomes.

While such an end-state is increasingly feasible, attempting to implement it holistically from the outset is both impractical and high risk.

In reality, the primary constraint in the investigation process is not decision-making or narrative writing; it's data fragmentation. Investigators are often required to access multiple systems (e.g. core banking platforms, KYC utilities, screening tools, and third-party data providers), manually extract relevant information, and consolidate it into a case management system. This process is time-consuming, inconsistent, and prone to error.

As a result, the most effective starting point for AI adoption is not narrative automation, but data extraction and structuring.

In the transaction monitoring case investigation scenario, an analyst may be required to access five or more different systems (e.g. business platform, KYC tool, screening engine, third party vendors), manually retrieve information and consolidate it into a case management tool.

Deloitte Diligence Insights can significantly reduce case processing time, allowing FIs to operate with a leaner workforce, and investigators to focus on higher-value risk assessment. From this starting point, FIs can demonstrate tangible cost savings, secure management buy-in, and establish a strong base for phase-two initiatives, such as generating AI-driven case narratives from the automated data.

It should also be said that while the Agentic AI-driven approach may allow FIs to run models that select only upstream source system data that is most pertinent and/or relevant to the specific review, FIs can achieve material time and cost savings even without an immediate AI component, but rather via Application Programming Interface (API) enabled rules-based Extract, Transform, Load (ETL). This approach can be applied to Customer Due Diligence (CDD) processes as well as name and transaction screening.

This reflects a broader dependency within financial crime investigations:

Data → Analysis → Narrative

Narrative generation is the final step in the process. Without structured and complete data, analysis is limited; and without robust analysis, narratives lack defensibility. Automating narrative generation in isolation therefore risks producing well-written but poorly substantiated outputs, which may not withstand regulatory scrutiny.

From rules to models

FIs are increasingly exploring the deployment of AI and machine learning (ML) in transaction monitoring detection architecture, moving from the typical rules-governed set up to a dynamic model-based surveillance framework. There are various key benefits in this regard. This shift is particularly relevant in Singapore. Regulatory expectations are evolving beyond isolated alert reviews toward the identification of network-level risk, especially in relation to scam activity and money mule networks. Individual accounts may appear low risk in isolation, but when assessed collectively, can form part of coordinated financial crime structures.

While previously FIs could review rules performance and recalibrate thresholds once every one to three years, regulators now expect far more frequent, even dynamic ongoing recalibration. Such retuning initiatives have historically been manual processes involving reviewing transaction alert data samples and recommending adjustments that are then required to be assessed via below the line testing prior to official re-deployment.

Applying a model-based approach to threshold retuning allows FIs to speed up the process significantly and allowing a larger data set to be analysed rather than sample-based methodologies. Cost savings are realised not only in the reduced headcount requirement for retuning exercises, but also in downstream processes such as case review wherein model-based recalibration reduces false-positive alerts and overall alert volumes. Moreover, as criminals adapt to rules a dynamic approach to retuning allows FIs to stay on top of developing trends in financial crime.

Taking this use case one step further, FIs may look to replace traditional rules-based detection scenarios entirely, in favour of agentic AI driven solutions. Such an approach may replace traditional detection scenarios and instead favour alert generation via risk-scoring.

Network analysis

Not long ago, AI technology advancements in areas such as facial recognition promised a new era of seamless, minimal touchpoints in the client onboarding process. Fast forward a few years and the irony is stark: the same technologies are now being exploited by scammers to compromise victims and infiltrate financial systems.

Between 2022 to 2023, scammers reportedly stole an estimated USD 1.02 trillion, with victims in Singapore losing the most on average. Why Singapore bears this unfortunate distinction is open to debate. Those of us fortunate to call the Little Red Dot home live in a high trust, relatively high-income society, surrounded in the Southeast Asian countries at varying stages of development. This unique geographical and sociocultural makes Singapore fertile ground for would-be scammers.

Most scam activity in Singapore is cross-border; facilitated by complex criminal networks located in countries such as Cambodia and Myanmar. This author has visited the Singapore Police Force's Anti-Scam Command Centre and observed first-hand the incredible work being done; however, the SPF's jurisdiction is naturally limited. FIs are increasingly looked to for collaboration in scam and fraud prevention, no longer only in the form of answering subpoenas and requests for information, but now in the form of proactive, preventative measures, especially beyond the borders of Singapore.

One emerging response is model-based surveillance to detect money mule activity. In addition to the Monetary Authority of Singapore's Collaborative Sharing of Money Laundering/Terrorism Financing Information & Cases (COSMIC) initiative, some banks have deployed machine learning models to better identify potential mules.

From an AI standpoint, FIs should shift toward assessing customer risk holistically for potential mules and scam risk, rather than waiting for a transaction to occur. FIs may for example build AI agents that look at a new customer's email address, phone number, and IP address and seek to identify other accounts and touchpoints at the FI sharing the information to expose potential networks; especially those spanning country borders. By identifying potential networks, FIs can move from a transaction-focused reactive transaction monitoring strategy to a predictive model-driven solution that addresses customer risk holistically across the organisation.

This has important implications for technology strategy. Capabilities such as entity resolution, link analysis, and cluster detection are no longer confined to complex, high-risk investigations. Instead, they are increasingly required across the broader alert population to identify weak signals and emerging patterns. This positions network analytics platforms as a core intelligence layer, rather than a niche investigative tool.

Where to start?

A common starting point for many financial institutions is the adoption of generative AI tools to automate case narrative writing. While this delivers visible productivity gains, it does not address the underlying constraint within the transaction monitoring process. In practice, investigators spend the bulk of their time finding, extracting, and preparing data, not writing narratives. As such, focusing on narrative automation addresses a symptom, not the root cause. Effective AI adoption must therefore follow the clear dependency chain

(Data → Analysis → Narrative)

Accordingly, AI tools should be positioned as productivity enhancers, supporting narrative generation and investigator workflow. However, they do not address core analytical challenges.

What	Why	How	Benefits
Data extraction & structuring	Regulatory increasingly expects complete, accurate and timely data usage in AML. Current bottleneck: Investigations spend excessive time gathering fragmented data.	Automate ingestion from PDFs, transactions and upstream systems. Standardise into a unified case dataset.	Regulatory: Stronger data integrity, completeness and audit trails. Efficiency: Reduces manual effort significantly; faster case preparation.
Network analytics	Regulatory focus on scam and mule networks requires detection beyond single alerts. Low-risk cases may form part of wider suspicious networks.	Implement entity resolution, link analysis, and clustering across all alerts. Enrich cases with networks context.	Regulatory: Aligns with regulatory expectations on network-level detection and holistic risk assessment. Efficiency: Earlier detection of organised activity; reduces repeat investigations.
Case management	Regulatory expects robust, consistent workflows and auditability. Existing systems already support this.	Retain current systems; integrate enriched data and analytic outputs into workflows.	Regulatory: Maintains auditability and control consistency. Efficiency: Avoid costly system replacement; leverages existing investments.
AI narrative	Regulatory expects clear, well-supported STR/SAR narratives. Manual writing is time-consuming but not the core bottleneck.	Use GenAI enable solutions to draft narratives based on structured data and analytics outputs, with human validation.	Regulatory: Improves consistency and clarity of reporting Efficiency: Reduces documentation time; accelerates case closure.
Dependency model	Regulatory's FEAT principles (Fairness, Ethics, Accountability, Transparency) require controlled and explainable AI usage.	Ensure AI consumes structured, validated inputs. Build upstream capabilities before scalling GenAI.	Regulatory: Ensures outputs are explainable and defensible. Efficiency: Avoids rework and poor-quality outputs.
Risk & governance (FEAT-aligned)	Regulatory emphasis on defensibility and explainability requires strong underlying data and analysis. Narrative alone is insufficient.	Implement human-in-the-loop, audit logs, model validation, and secure deployment environments.	Regulatory: Meets FEAT expectations and reduces model risk. Efficiency: Minimises regulatory remediation costs and operational risk.
Strategic outcome	Regulatory is driving a shift toward proactive, intelligence-led AML rather than reactive alert handling.	Build a layered architecture combining structured data, network analytics, and AI augmentation.	Regulatory: Future-proofs compliance with evolving regulatory expectations. Efficiency: Delivers both cost savings and improved detection effectiveness.

Lastly, AI adoption in financial crime compliance must be source-system and agnostic. FIs must establish an AI framework that sits above its system architecture and can be redeployed quickly and accurately should changes occur (for example a change case management software). Deloitte is well-placed to ensure FIs remain dynamic and resilient to change from a business continuity perspective in their AI adoption and change initiatives.

Deloitte has assisted numerous clients, from traditional banks to social media and payments companies, on AI solution implementation and risk management and is therefore well positioned to help FIs adopt and scale further AI capabilities.

A measured path forward

Meaningful AI adoption in financial crime compliance is achievable in the near term, provided institutions take a measured, phased approach. Rather than attempting wholesale transformation, organisations should prioritise foundational capabilities, particularly data structuring and network intelligence, before scaling advanced AI use cases.

Ultimately, successful adoption is not defined by the deployment of the most visible tools, but by addressing the most fundamental constraints in the investigative process. Institutions that do so will not only realise efficiency gains, but also build more effective, defensible, and regulator-aligned AML frameworks in an increasingly complex threat environment.

Contact us

Graham Dawes
Forensic & Financial Crime Leader
Deloitte Southeast Asia
gdawes@deloitte.com

Kalyani Vasan
Forensic & Financial Crime Partner
Deloitte Southeast Asia
mvasan@deloitte.com

Jamie King
Forensic & Financial Crime Director
Deloitte Singapore
tking@deloitte.com

Deloitte.