

Suspicious Transaction Detection: A Practical Guide for Banks

Understanding AI-Powered Monitoring, Anomaly Detection, Graph Analytics, and Investigation Workflows

GSDC – GLOBAL SKILL DEVELOPMENT COUNCIL

Introduction to Suspicious Transaction Detection

The financial landscape has undergone a profound transformation over the past decade. The proliferation of digital payment channels, real-time settlement systems, mobile banking applications, and cross-border payment networks has dramatically increased both the volume and complexity of financial transactions that banks must monitor. Whilst this evolution has brought significant convenience to customers, it has also created new opportunities for those seeking to abuse the financial system for illicit purposes.

Financial crime has become increasingly sophisticated, often exploiting multiple accounts, payment channels, and jurisdictions simultaneously to obscure the origin and movement of funds. Traditional rule-based transaction monitoring systems – built on static thresholds and rigid logic – were designed for a simpler era. Today, these systems generate enormous volumes of alerts, many of which relate to entirely legitimate customer activity. The resulting alert fatigue places significant pressure on AML analysts and compliance teams, diverting investigative resources away from genuinely suspicious activity.

Disconnected customer and transaction records compound this challenge. When data held across core banking platforms, card systems, digital channels, and customer information systems cannot be effectively linked, investigators lack the contextual understanding needed to make well-informed, timely decisions. The need for faster, evidence-based investigation workflows – supported by modern technology – has never been more pressing.

Traditional Rule-Based Monitoring

- Static thresholds applied uniformly
- High volumes of false-positive alerts
- Limited contextual awareness
- Disconnected data sources
- Slow, manual investigation workflows
- Reactive rather than predictive

AI-Supported Monitoring

- Dynamic, behaviour-based detection
- Improved alert prioritisation
- Contextual customer and account analysis
- Integrated data across all channels
- Structured, evidence-driven workflows
- Continuous model improvement

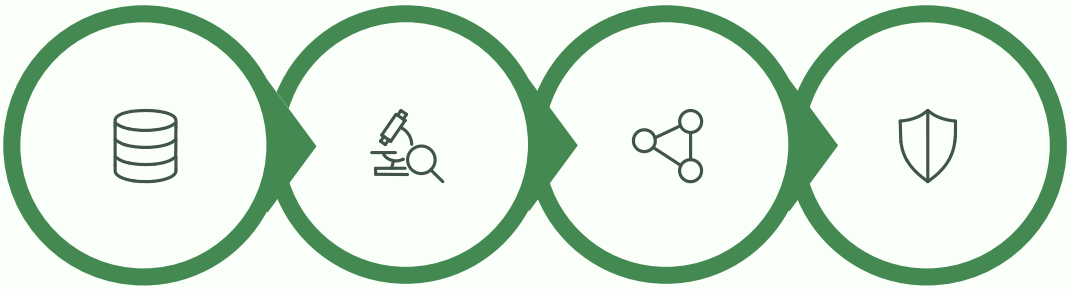
AI-generated alerts indicate potential risks requiring investigator review. They do not independently confirm financial crime or determine regulatory reporting obligations.

Understanding AI-Powered Transaction Monitoring

AI-powered transaction monitoring represents a meaningful advance over traditional rule-based approaches. Rather than applying fixed thresholds uniformly across all customers, AI systems learn from historical transaction data and develop nuanced models of expected behaviour at the individual customer, account, and segment level. This allows monitoring systems to identify deviations that static rules would miss, whilst reducing unnecessary alerts for transactions that are consistent with established patterns.

Effective AI transaction monitoring is built on four interconnected stages, each of which contributes essential information to the overall risk assessment. These stages work together as part of an integrated pipeline, rather than as isolated processes. The output of each stage informs the next, culminating in prioritised, evidence-backed alerts that support structured investigator review.

It is important to emphasise that AI monitoring systems are decision-support tools. Alerts generated by these systems indicate potential risks and flag activity that warrants further investigation. Authorised investigators retain full responsibility for reviewing the available evidence, assessing context, and determining whether further action or regulatory reporting is appropriate.



Data Integration

Anomaly Detection

Graph Analytics

Risk Scoring

Key AI Technologies Used in Suspicious Transaction Detection

Modern transaction monitoring platforms draw on a range of complementary AI and data science technologies. Each technology addresses a specific aspect of the detection challenge, and together they provide a more comprehensive and adaptive capability than any single approach could offer alone. Understanding the purpose and application of each technology is essential for banking professionals evaluating or implementing AI monitoring solutions.

Technology	Purpose	Application in Banking	Example
Machine Learning	Identifies patterns in historical and current transaction data	Builds predictive models of expected customer behaviour	Classifying transactions by risk category based on learned patterns
Anomaly Detection	Flags transactions deviating from established behavioural baselines	Identifies unusual amounts, timing, or frequencies for a given customer	Flagging a sudden high-value transfer from a low-activity account
Graph Analytics	Maps relationships between accounts, entities, and transactions	Reveals fund movement patterns across connected accounts	Identifying a cluster of accounts sharing common beneficiaries
NLP	Analyses unstructured text and investigation documents	Extracts relevant information from payment references and case notes	Identifying unusual payment references that may indicate structuring
Predictive Analytics	Supports risk assessment and alert prioritisation	Ranks alerts by estimated likelihood of requiring escalation	Prioritising high-risk accounts for senior investigator review

These technologies are most effective when integrated within a coherent data architecture that provides consistent, high-quality transaction and customer information. The value of AI monitoring is directly proportional to the quality and completeness of the underlying data available to the system.

Understanding Suspicious Transaction Patterns

Effective transaction monitoring requires investigators and analysts to understand the types of patterns that may warrant further examination. It is essential to note that the presence of any of these indicators does not constitute evidence of illegal activity. Each pattern must be assessed in the context of the customer's known profile, business activity, and the broader transactional environment. The purpose of monitoring is to surface activity that requires contextual review, not to make determinations of wrongdoing.

Unusual Transaction Frequency

A significant increase in the number of transactions over a short period, particularly where the volume is inconsistent with the customer's established pattern of activity.

Sudden Changes in Transaction Amounts

Large transactions that are disproportionate to the customer's typical activity, or a pattern of transactions structured just below reporting or monitoring thresholds.

Multiple Transfers Involving Connected Accounts

Funds moving through a series of interconnected accounts in a short timeframe, particularly where the relationships between account holders are not clearly established.

Rapid Movement of Funds

Funds received into an account and immediately transferred to other accounts, sometimes described as "pass-through" behaviour, which may indicate layering activity.

Unusual Timing or Geographical Patterns

Transactions occurring at unusual times or involving geographies inconsistent with the customer's known activity, travel history, or business operations.

Activity Inconsistent with Customer Profile

Transaction behaviour that does not align with the customer's stated occupation, business type, or the customer risk profile held by the bank.

These indicators require contextual review by authorised investigators. Pattern identification is the starting point for investigation, not a conclusion about the nature of the activity.

Anomaly Detection in Banking Transactions

Anomaly detection models are designed to identify transactions or patterns of behaviour that deviate materially from what would be expected for a given customer, account, or customer segment. Unlike rule-based systems that apply uniform thresholds, anomaly detection models learn from historical data to establish individual behavioural baselines, making them sensitive to changes that might not cross a fixed threshold but are nonetheless unusual for a specific customer.

The development and maintenance of anomaly detection models involves several interconnected components. Historical transaction analysis forms the foundation, establishing what normal activity looks like for each customer over time. Feature engineering – the process of selecting and transforming relevant data attributes – is critical to model quality. Relevant features may include transaction amounts, frequency, timing distributions, beneficiary diversity, geographical spread, and channel usage patterns.

Normal Transaction Behaviour

- 3 - 5 domestic transfers per month
- Consistent transaction amounts (£200 - £1,500)
- Established recurring beneficiaries
- Transactions during standard business hours
- Single geographic region
- Consistent channel usage (mobile banking)

Flagged Unusual Behaviour

- 12 transfers in a 48-hour period
- Multiple transactions of £9,800 - £9,950
- 8 new beneficiaries added simultaneously
- Transactions at 02:00 - 04:00 local time
- Transfers to 4 different countries
- Simultaneous use of web and mobile channels

Model evaluation is an ongoing process. Precision, recall, and false-positive rates must be monitored continuously to ensure models remain effective as customer behaviour and financial crime methods evolve. Feature drift – where the statistical properties of input data change over time – can degrade model performance if not actively managed. Banks should establish regular model review cycles and maintain clear governance frameworks for model updates and deployment.

Practical Example: A retail customer who typically makes a small number of domestic transfers each month suddenly initiates multiple high-value transfers to previously unknown beneficiaries over a 72-hour period. An anomaly detection model identifies this deviation from the customer's baseline and raises an alert for investigator review.

Graph Analytics for Identifying Suspicious Networks

Graph analytics provides a powerful capability that complements individual transaction analysis. Whilst anomaly detection examines the behaviour of individual accounts or customers in isolation, graph analytics maps the relationships between entities – customers, accounts, beneficiaries, devices, and merchants – to reveal patterns that would be invisible at the individual transaction level. This network perspective is particularly valuable in the context of layering and structuring activity, where funds move through multiple connected accounts before reaching their destination.

In a graph model, entities are represented as nodes and relationships between them are represented as edges. The properties of both nodes and edges can be enriched with transactional and behavioural attributes, allowing analysts to query the network for specific patterns. For example, investigators can identify accounts that share common beneficiaries, devices used across multiple customer accounts, or chains of transfers that move funds rapidly through a series of intermediary accounts.



Customer-to-Account Relationships

Mapping individuals to the accounts they control, including joint accounts and accounts where they have signatory authority, reveals the full scope of a customer's banking relationships.



Account-to-Account Transfers

Visualising fund flows between accounts highlights transfer chains, pass-through patterns, and the velocity at which funds move through the network.



Shared Beneficiaries and Devices

Accounts sharing common beneficiaries or devices may indicate coordination between nominally unrelated customers, which warrants further contextual investigation.



Highlighted Investigation Clusters

Graph algorithms can automatically identify densely connected sub-networks – clusters of entities with unusually high transaction volumes or shared relationships – for prioritised investigator review.

Graph analytics does not replace investigator judgement. The relationships identified by graph models provide important context that investigators can use to understand the wider environment surrounding a suspicious activity alert. The presence of network connections requires contextual assessment, not automatic escalation.

How Banks Can Reduce False-Positive Alerts

One of the most significant operational challenges in transaction monitoring is the management of false-positive alerts – cases where legitimate customer transactions trigger monitoring rules or AI-based indicators. High false-positive rates consume investigative capacity, contribute to analyst fatigue, and can reduce the overall effectiveness of the monitoring programme by diverting resources from genuinely suspicious activity. Addressing this challenge requires a disciplined, evidence-based approach to model and threshold management.

Metric	Definition	Why It Matters
Precision	Of all alerts raised, the proportion that relate to genuinely suspicious activity	High precision reduces unnecessary investigative work
Recall	Of all suspicious activity that occurs, the proportion successfully identified by the system	High recall ensures suspicious activity is not missed
False-Positive Rate	The proportion of alerts raised that relate to legitimate activity	High rates reduce investigative efficiency and increase costs
Alert Prioritisation	The ranking of alerts by estimated risk level to guide investigator workload management	Ensures highest-risk cases receive prompt attention

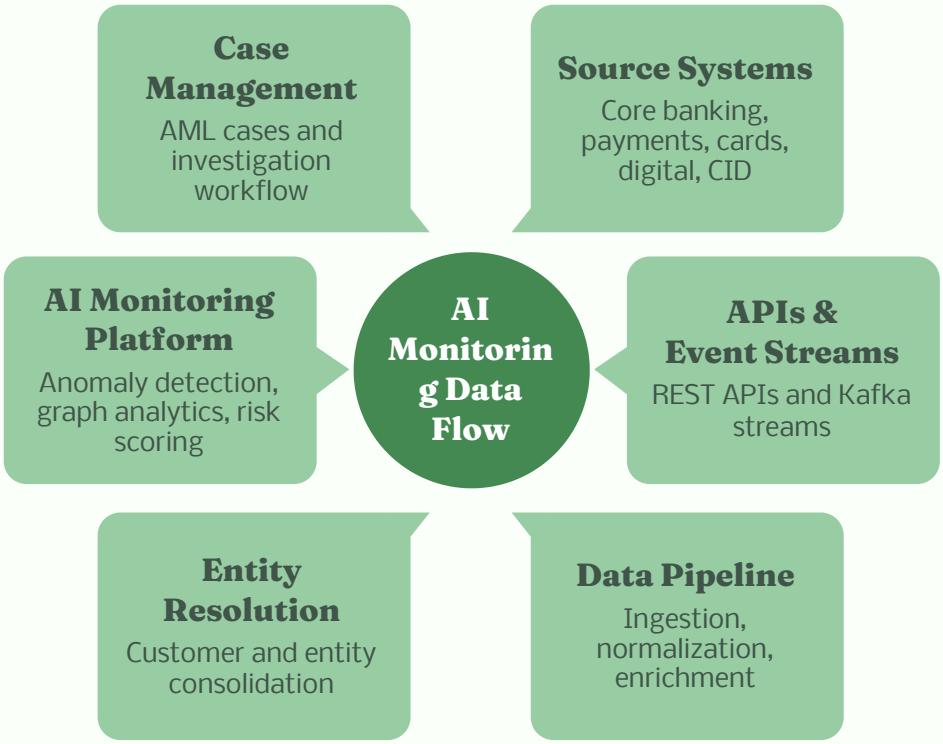
Reducing false positives effectively requires enriching monitoring models with better customer context. When a monitoring system has access to detailed customer profile information – including their occupation, business activity, expected transaction volumes, and known counterparties – it can distinguish more reliably between unusual and genuinely suspicious behaviour.

Practical Example: A business customer operating a wholesale import company regularly makes high-value international transfers to supplier accounts. Without adequate customer context, these transactions may trigger monitoring alerts repeatedly. Enriching the monitoring system with the customer's known business profile allows the model to recognise this as expected activity, reducing unnecessary alerts whilst maintaining detection of genuinely unusual transfers.

It is critical that efforts to reduce false-positive rates do not compromise the monitoring system's ability to identify genuinely suspicious activity. Threshold adjustments and model refinements should always be evaluated against their impact on recall, and governance frameworks should require formal sign-off before changes are deployed to production environments.

Integrating Banking Systems for AI Monitoring

The effectiveness of any AI transaction monitoring solution depends fundamentally on its ability to access consistent, complete, and timely data from the full range of banking systems in which customer and transaction information resides. For most banks, this represents a significant engineering challenge. Transaction data is typically distributed across multiple source systems – each with its own data formats, identifiers, and update cycles – that were not originally designed to interoperate.



REST APIs

Enable structured, authenticated data exchange between banking source systems and the AI monitoring platform.

Apache Kafka

Provides event-driven, real-time data streaming to support near-real-time transaction analysis and alert generation.

Entity Resolution

Links records relating to the same customer or account across different source systems, ensuring a consistent view of each entity.

Data Pipelines

Orchestrate the ingestion, transformation, validation, and delivery of data to the monitoring platform from multiple sources.

Critical engineering considerations include secure authentication and encrypted data transmission between all system components, robust handling of duplicate events and failed messages within the event streaming infrastructure, and consistent use of entity identifiers across source systems to enable reliable data linkage. Data validation and reconciliation processes should verify data completeness and integrity at each stage of the pipeline. Delayed processing – where transaction data arrives late from specific source systems – must be handled gracefully to avoid gaps in the monitoring coverage.

Suspicious Transaction Investigation Workflow

A structured, AI-supported investigation workflow enables banks to manage alert volumes systematically, ensure that evidence is presented consistently to investigators, and maintain clear audit trails of investigation decisions. The following six-stage workflow illustrates how AI capabilities can support each phase of the investigation process, whilst preserving the central role of authorised investigators in reviewing evidence and making determinations.

1

1. Detect Unusual Activity

Monitoring rules and AI anomaly indicators identify transactions or patterns warranting further review.

2

2. Connect Related Records

Link relevant accounts, beneficiaries, customers, and transaction histories into a unified investigation view.

3

3. Analyse Patterns

Examine transaction sequences, fund movement, frequency, timing, and network relationships in detail.

4

4. Prioritise Alerts

Risk indicators and scoring models rank alerts to guide investigator workload and focus attention appropriately.

5

5. Review Evidence

Investigators access transaction timelines, customer information, previous alerts, and relevant supporting records.

6

6. Document Outcomes

Findings are recorded in the case management system; investigators determine whether further action or regulatory reporting is required.

Throughout this workflow, authorised investigators remain responsible for reviewing the available evidence, assessing contextual factors, and making all substantive investigation decisions. AI systems support this process by organising information, surfacing relevant evidence, and prioritising workloads – they do not make investigation or regulatory reporting decisions independently. Banks should ensure that their investigation workflows are documented, auditable, and subject to appropriate quality assurance oversight.

Practical Case Study – Investigating a Suspicious Transaction Network

This scenario is entirely illustrative and hypothetical. It does not represent a real bank investigation, real individuals, or real accounts. It is provided for educational purposes only.

A bank's transaction monitoring platform identifies an unusual pattern across several retail accounts that individually do not cross existing monitoring thresholds. However, when examined collectively, these accounts demonstrate a series of characteristics that the AI monitoring system flags for further review.

1 AI Monitoring Identifies Unusual Patterns

The anomaly detection model identifies unusual transaction frequency and a rapid increase in the number of new beneficiaries across a cluster of accounts within a 96-hour window. Individually, no single account crosses the monitoring threshold, but the collective pattern is identified as a material deviation from established baselines.

2 Graph Analytics Connects the Accounts

Graph analytics reveals that the flagged accounts share two common beneficiaries and that funds flow between the accounts in a pattern consistent with rapid layering. Several accounts share a common device identifier used during the account opening process.

3 Transaction Sequence Analysis

Detailed analysis of the transaction sequences shows that funds received into each account are transferred to the shared beneficiaries within hours of receipt. The timing pattern is consistent across all accounts in the identified cluster.

4 Risk Scoring and Prioritisation

The risk-scoring model combines the network connectivity findings, transaction velocity indicators, and new beneficiary patterns to assign a high-risk score to the cluster, prioritising it for senior investigator review within the case management system.

5 Investigator Review and Documentation

The case management system presents investigators with transaction timelines, network relationship visualisations, customer profile information, and the full history of previous alerts. Investigators review the assembled evidence, document their findings, and determine the appropriate next steps, including whether regulatory reporting obligations are engaged.

Technical Challenges and Implementation Checklist

Deploying AI transaction monitoring in a banking environment involves significant technical and operational complexity. Banks that approach implementation with a clear understanding of the key challenges are better positioned to make informed decisions about technology selection, data architecture, and governance frameworks. The following table identifies the most important challenges and the primary considerations associated with each.

Technical Challenge	Why It Matters	Implementation Consideration
Data Fragmentation	Incomplete data degrades model quality and creates gaps in monitoring coverage	Map all relevant data sources; establish data integration pipelines before model deployment
Real-Time Processing	Delayed detection reduces the ability to intervene before funds are moved further	Implement event-driven architecture with low-latency data streaming capabilities
Entity Resolution	Inconsistent identifiers across systems prevent reliable linking of related records	Deploy entity resolution tooling; standardise customer and account identifiers
Model Explainability	Investigators and regulators require understandable, auditable explanations for alerts	Select models with interpretable outputs; document reasoning for all alerts generated
False-Positive Alerts	High alert volumes reduce investigative efficiency and increase operational costs	Enrich models with customer context; monitor precision and recall continuously
Data Privacy and Access Control	Transaction monitoring involves sensitive personal and financial data	Implement role-based access controls; ensure data handling complies with applicable regulations
Model Drift	Customer behaviour and financial crime methods evolve, degrading model performance over time	Establish model monitoring and regular retraining cycles with formal governance sign-off
Audit Trails and Governance	Regulators expect comprehensive records of monitoring decisions and investigation outcomes	Ensure all alerts, investigations, and model decisions are logged and retained appropriately

AI Transaction Monitoring Implementation Checklist

- ◯ All relevant data sources identified and mapped
- ◯ API and system integration points documented
- ◯ Data quality and validation processes established
- ◯ Entity resolution tooling deployed and tested
- ◯ Model training, testing, and evaluation completed
- ◯ Precision, recall, and false-positive rates reviewed
- ◯ Secure access controls and encryption implemented
- ◯ Model explainability outputs reviewed by investigators
- ◯ Investigator workflow integration tested end-to-end
- ◯ Audit trail and logging capabilities verified
- ◯ Model drift monitoring and retraining schedule agreed
- ◯ Governance and sign-off frameworks documented

The Role of Forward Deployed Engineers in Banking AI

The gap between AI capability and practical, production-ready implementation in a banking environment is often wider than it first appears. Data science teams can develop sophisticated models, and technology vendors can supply powerful platforms, but the work of connecting these capabilities with existing banking infrastructure, operational workflows, and the practical realities of a regulated financial institution requires a specific and valuable skill set. Forward Deployed Engineers (FDEs) are professionals who specialise in bridging this gap.



Understand Banking Requirements

FDEs work directly with banking stakeholders, AML teams, compliance professionals, and investigation teams to develop a detailed understanding of the operational challenges, regulatory requirements, and existing workflows that the AI implementation must support.



Connect Banking Systems

FDEs design and implement the data pipelines, API integrations, and event-driven architecture components needed to connect core banking platforms, payment systems, and investigation applications with AI monitoring capabilities.



Support AI Implementation

FDEs help integrate anomaly detection models, graph analytics engines, and risk-scoring components into the bank's existing technology environment, ensuring that AI-generated insights are accessible to authorised investigators in a structured and auditable format.

Improve Investigation Workflows

Ensure AI insights and supporting evidence are presented effectively within investigator-facing applications and case management platforms.

Support Production Deployment

Work with engineering teams on testing, performance monitoring, system reliability, and the ongoing operational stability of deployed AI monitoring solutions.

FDEs work collaboratively with banking experts, data scientists, compliance professionals, and engineering teams. They do not replace authorised investigators or compliance decision-makers, and they do not make AML or regulatory reporting determinations.

Build Your Forward Deployed Engineering Knowledge With GSDC

The GSDC Certified Forward Deployed Engineer Certification is designed for professionals who want to develop a structured understanding of the FDE role and build relevant knowledge in customer-focused engineering and practical technical implementation. Whether you work in banking technology, AI implementation, compliance systems, or financial data engineering, this certification provides a structured framework for developing and demonstrating FDE competencies.

1

Connect Business and Technology

Understand how FDEs translate business requirements into practical technical implementations that deliver operational value.

2

Explore AI Applications

Explore practical AI applications and implementation approaches relevant to regulated industries including banking and financial services.

3

APIs and System Integration

Develop knowledge of APIs, engineering tools, event-driven architecture, and the technical patterns used in enterprise system integration.

4

Support Real Business Workflows

Understand how technical solutions are designed and deployed to support real operational workflows in complex, regulated environments.

5

FDE Responsibilities and Skills

Explore the responsibilities, competencies, and professional skills associated with the Forward Deployed Engineering role.

> **Ready to develop your Forward Deployed Engineering knowledge?** Explore the GSDC Certified Forward Deployed Engineer Certification and take the next step in your professional development.

🌐 https://www.gsdcouncil.org/certified-forward-deployed-engineer

Explore the Certification

GSDC – GLOBAL SKILL DEVELOPMENT COUNCIL

CERTIFIED FORWARD DEPLOYED ENGINEER

GSDC
Global Skill Development Council

Globally Recognised Certification Body
190+ countries · 2,50,000+ Certified Professionals

IAS ACCREDITED Personnel Certification Body

Certified Forward Deployed Engineer



ABOUT GSDC CERTIFICATION

EBOOK

Extensive and exclusive Ebook created by world's experts to help you with understanding core concepts.

LEARNING MATERIALS

Get access to learning materials such as videos, ebooks, templates, and practice exams, which will help you clear the certification exam.

CREATED BY EXPERTS

GSDC certifications are created and authored by world's leading experts in the field.

LEARNING OBJECTIVE

- **Gain a strong understanding of Forward Deployed Engineering concepts, roles, tools, technologies, and development practices.**
- **Identify business and customer challenges and design AI-driven solutions aligned with real-world requirements.**
- **Apply modern AI frameworks, APIs, data integration, and deployment practices to build and implement solutions within existing business systems.**
- **Apply testing, evaluation, security, monitoring, and optimization practices to deliver reliable, scalable, and production-ready AI solutions.**

Enroll now with the code **LEARN20** To avail **20%** discount

Enroll Now

www.gsdcouncil.org