QKS Group | Bottomline

QKS AI Maturity Matrix™

Commercial Payment

Most Valuable Pioneer

Anandh Ramaswamy
Practice Director Financial Crime & Compliance

Divya Baranawal
Vice President & Principal Analyst

TABLE OF CONTENTS

Executive Summary

Commercial payment fraud remains a critical risk for financial institutions and corporates around the world. Digital transformation has accelerated the adoption of real-time payment systems, creating new opportunities for revenue growth and customer convenience; however, it has also expanded the attack surface for fraudsters. Traditional rule-based defenses are increasingly insufficient because sophisticated criminals use social engineering, credential theft, malware and coordinated networks to bypass static controls. Regulatory expectations around customer protection, Anti-Money Laundering (AML) and data privacy are rising, adding complexity and cost.

Artificial Intelligence (AI) offers a transformative path for fraud detection and prevention. AI can learn from vast datasets, detect subtle anomalies, adapt to evolving patterns and automate risk management. QKS Group's **AI Maturity Matrix™** provides a structured framework for institutions to evaluate their current capabilities and plan their journey from rule-based syssssstems to self-optimizing AI.

This comprehensive white paper combines quantitative and qualitative analysis, competitor benchmarking, marketing strategy insights, implementation best practices, and regulatory perspectives. The document spans the entire payment fraud lifecycle, articulates current challenges, explores AI-driven solutions, and forecasts emerging trends. The objective is to provide a full-fledged, publishable white paper that not only informs but also persuades potential customers and partners of the value of advanced AI fraud defenses.

Introduction

The global payments ecosystem has undergone dramatic changes in recent years. Innovations such as instant payments, open banking, digital wallets and blockchain are reshaping how money moves. Enterprises are digitizing supply chains and adopting real-time treasury and cash-management solutions. This shift offers tremendous benefits, improved liquidity, faster settlements and better customer experience yet it also creates fertile ground for fraud. Fraudsters exploit new channels, technologies and customer behaviors to execute complex schemes. Financial institutions must adapt quickly or risk reputational damage and financial loss.

AI maturity in commercial payment fraud reflects how far detection and mitigation capabilities have evolved beyond legacy rules-based systems. In the past, financial institutions relied on static heuristics and manual review to identify suspicious transactions. Today's high speed payments ecosystem demands adaptive, intelligent controls that can recognize new fraud patterns in real time.

Several dimensions define the maturity of AI in commercial fraud platforms:

- **Integrated analytics and behavioral profiling.** Mature platforms use statistical and behavioral profiles that continuously learn from current and historical data to detect anomalies. Rather than relying solely on historical fraud labels, the models evaluate login behavior, session activity, changes to beneficiary accounts and patterns across all payment types (Nacha, Fedwire, RTP, FedNow, Swift, etc.). This broad coverage reduces blind spots and improves detection accuracy.
- **Adaptive risk scoring and unsupervised learning.** Advanced AI engines perform dynamic risk scoring that learns from investigator feedback as patterns shift. They use unsupervised machine learning to spot emerging fraud before losses occur, allowing enterprises to respond to novel threats quickly.
- **Natural language and generative AI.** Next generation solutions incorporate conversational assistants and generative AI to simplify onboarding and rule configuration, help investigators interpret alerts, and provide explainable AI

narratives for trust and compliance. These tools guide analysts to the most important cases and articulate model decisions in plain language.

- **Human in the loop tuning and governance.** AI mature platforms maintain human oversight with features such as guided alert prioritization, model performance reports, and tuning recommendations. Fraud AI tuning agents enable automated model optimization while still allowing investigators to review or override changes.
- **Consortium and network intelligence.** High maturity also entails ingesting external intelligence from consortium partners and network peers. By correlating login, session and payment activity across banks, AI can identify account takeover and mule networks earlier.

Assessing AI maturity means evaluating how well a platform incorporates these elements—behavioral analytics, adaptive learning, explainable models, human in the loop governance and cross bank intelligence. As payment channels diversify and fraud tactics evolve, institutions should adopt AI solutions that not only detect known patterns but also adapt to emerging threats and provide transparent, actionable insights to investigators.

The purpose of this paper is to provide a holistic view of the commercial payment-fraud landscape, with particular emphasis on AI maturity. It synthesizes QKS research, industry trends, regulatory requirements, case studies and marketing considerations.

Market Context and Industry Landscape

Global adoption of real-time payments

Real-time payment schemes have proliferated worldwide. Systems such as the UK's Faster Payments Service, India's Unified Payments Interface (UPI), Brazil's PIX, Europe's SEPA Instant Credit Transfer, and the U.S.'s FedNow Service provide immediate clearing and settlement of transactions. Real-time payments support e-commerce, payroll, government disbursements and peer-to-peer transfers. Adoption is driven by consumer expectations for instant gratification, corporate treasury needs for immediate liquidity and regulators' push for modernizing payment infrastructure.

However, the speed of real-time payments reduces the window available for fraud screening. Once funds are transferred, they cannot be reversed easily. Fraudsters take advantage of this by executing fast, high-value transactions across multiple institutions and jurisdictions. AI must operate at the same speed as real-time payments to be effective.

Digital transformation and customer behavior

The COVID-19 pandemic accelerated digital transformation in payments. Businesses and consumers shifted to online channels, increasing digital wallet usage, contactless payments and e-commerce. Corporate treasury functions adopted software-as-a-service (SaaS) solutions for accounts payable and cash management. This digital shift introduced new customer behaviors and transaction patterns. AI models need to learn and adapt to these changes. For example, remote working and flexible hours result in payments initiated outside traditional business hours; AI must differentiate legitimate off-hour activity from suspicious transactions.

Competitive landscape

The commercial payment-fraud prevention market is highly competitive. Vendors range from global enterprises to specialized fintech start-ups. Key players include

NICE Actimize, Feedzai, Bottomline, Featurespace, SAS, Oracle, and Nasdaq Verafin. Each vendor brings different strengths and market focus:

- **NICE Actimize** – Established fraud and AML suite used by large banks; broad functionality but slower pace of AI innovation.
- **Feedzai** – Strong machine-learning and graph analytics capabilities; focuses on retail and card fraud but expanding to commercial payments.
- **Bottomline** – Strong heritage in commercial payments, adaptive AI, deep digital-banking integration, record-and-replay, generative AI, extensive roadmap.
- **Feature space** – Innovator in adaptive behavioural analytics; excels in card fraud; developing commercial payments capabilities.
- **SAS** – Offers advanced analytics and machine-learning tools; often requires significant customization and data-science resources.
- **Oracle** – Provides an enterprise fraud management platform integrated with Oracle banking systems; relies more on rules and conventional analytics.
- **Nasdaq Verafin** – Focuses on small and mid-tier banks; leverages consortium data and network analytics; limited generative AI.

Segment-specific requirements

Different customer segments have distinct needs:

- **Large global banks** require scalable platforms that can handle high transaction volumes, integrate with complex IT infrastructure and comply with multiple regulatory regimes. They value extensible AI, flexible deployment and high configurability.
- **Regional and community banks** often lack internal data-science resources and prefer managed services. They need easy deployment, pre-configured models, and low operational overhead.
- **Corporates and treasury centres** require fraud detection integrated with ERP and treasury systems. They need strong onboarding, adaptive analytics and link analysis to detect insider fraud and vendor impersonation.

- **Fintechs and payment service providers (PSPs)** need API-driven solutions that can be embedded into their own products and scale rapidly.

Market drivers and inhibitors

Drivers:

- Growth of digital commerce and real-time payments.
- Increasing value and complexity of commercial transactions.
- Regulatory mandates for risk management and consumer protection.
- Competitive pressure to offer seamless customer experience.

Inhibitors:

- Legacy IT infrastructure that is difficult to integrate with modern AI platforms.
- Cost and complexity of implementing AI solutions.
- Shortage of skilled data scientists and risk analysts.
- Privacy concerns and data sharing restrictions.

Challenges and Pain Points in Commercial Payment Fraud

Evolving attack vectors

- Commercial payment fraud evolves rapidly. Attackers exploit vulnerabilities at every stage of the payment lifecycle. New attack vectors continue to emerge:
- Synthetic identity fraud – Fraudsters create identities using stolen and fabricated data. They may open accounts, build credit and then exploit credit lines and payment facilities.
- API exploitation – Open banking APIs and payment APIs can be manipulated if authentication is weak or API endpoints are misconfigured. Attackers might use bots to test and exploit API vulnerabilities.
- Deepfakes and voice synthesis – Fraudsters can impersonate executives or suppliers in video or voice calls, convincing employees to approve payments.
- IoT and device spoofing – Compromised devices or IoT sensors can be used to initiate unauthorized payments or provide fraudulent confirmation signals.

Data fragmentation and silos

Data needed for fraud detection reside in multiple systems: core banking, digital channels, payment rails, ERP, CRM, and third-party data providers. Without unified data pipelines, AI models cannot see the full picture. Legacy systems often lack real-time data access, and integration projects can be costly and slow. Modern fraud platforms must support batch and streaming data ingestion, normalize data into a common model (e.g., ISO 20022) and integrate with existing middleware.

Workforce skills and change management

Effective AI deployment requires multidisciplinary teams: data scientists, data engineers, fraud analysts, IT operations and compliance officers. Many organizations lack the internal expertise to build and manage AI models. Training existing staff to interpret AI outputs and tune parameters is essential. Cultural resistance to AI adoption may also hinder implementation. Leaders must communicate the benefits

and empower teams with tools that enhance, not replace, human expertise.

Regulatory complexity

Financial regulations vary by jurisdiction and product. Payment service providers must comply with AML laws (e.g., the Bank Secrecy Act in the U.S.), KYC requirements, sanction screening, data-protection regulations (GDPR, CCPA) and consumer protection rules (PSD2/PSD3). Regulations such as PSD2 enforce strong customer authentication and transaction risk analysis for electronic payments. AI models must incorporate rules that comply with these regulations while avoiding discrimination or bias. Explainability and auditability are key.

Legacy architecture and integration challenges

Banks and corporates often operate multiple core systems, some decades old. These systems were not designed to integrate with modern AI services. Extracting data for real-time analysis requires custom connectors, real-time data pipelines and sometimes a complete overhaul of IT architecture. In addition, many organizations rely on third-party managed service providers for hosting and maintenance; aligning upgrades and integrations with external partners can be complex. These challenges increase the cost and risk of AI adoption.

AI-Driven Automation and Decision Intelligence

Real-time AI architecture

An AI-powered fraud detection architecture consists of several layers:

- **Data ingestion and feature engineering** – Real-time ingestion of payments, login events, devices, geolocation, account changes, third-party threat intelligence, and consortium data. ETL pipelines and stream processors (e.g., Apache Kafka, Flink) feed a feature store where data is enriched and normalized.

- **Machine-learning models** – A combination of supervised, unsupervised and hybrid models scores transactions. Supervised models (e.g., gradient boosting, logistic regression) learn from labelled fraud and non-fraud transactions. Unsupervised models (e.g., isolation forest, autoencoder) detect anomalies in unlabeled data. Hybrid models may combine both to improve accuracy and reduce false positives.

- **Graph analytics and network analysis** – Relationship graphs linking accounts, devices, IP addresses, emails and merchants enable detection of fraud rings and mule networks. Graph-based algorithms such as PageRank and community detection identify suspicious clusters and patterns.

- **Decision orchestration** – Risk scores and business rules are combined to decide whether to allow, hold or block a transaction. Orchestration is governed by configurable policies that weigh customer experience, fraud risk and regulatory requirements. Reinforcement learning can optimize decision policies over time.

- **Explainability and human oversight** – Generative AI and explainability tools convert model outputs into understandable narratives. They identify which features influenced the risk score (e.g., new device, unusual amount, unusual hour). Investigators can review explanations and override decisions if

necessary. Dashboards display performance metrics and enable simulation of policy changes.

Behavioral analytics and user profiling

Modern fraud models profile users and entities across multiple dimensions:

- **User behaviour** – Number of logins, frequency, average transaction size, preferred devices, geolocation patterns. Changes such as a high-value transaction initiated from a new device at 2 a.m. may indicate risk.
- **Device and session attributes** – Device ID, operating system, browser fingerprint, IP address, network type. Device velocity (number of devices used in a given time) and session duration are important indicators of ATO or bot attacks.
- **Merchant and beneficiary profiles** – Merchant categories, beneficiary history, transaction history. Unknown beneficiaries or changes to bank account numbers raise risk scores.
- **Temporal patterns** – Seasonality, time of day, day of week. Payments initiated outside normal business hours can signal BEC or insider fraud. Machine learning models capture these temporal anomalies.

Graph intelligence and consortium data

Fraud often involves networks of mule accounts, third-party money mules and collusion between insiders and external actors. Graph analytics reveals connections between seemingly unrelated transactions:

- **Link analysis** – Visualizes and analyses the relationships among customers, accounts, devices and IP addresses. Suspicious clusters or short paths between previously unknown entities can indicate fraud rings.
- **Consortium data** – Banks share anonymized transactional and behavioural data to identify cross-institutional patterns. Consortium data helps detect multi-bank fraud rings, because a single institution may not have enough information to see the full network.

Adaptive authentication and risk orchestration

AI-driven decisioning goes beyond simple allow or deny outcomes. It involves multi-step orchestration:

- **Dynamic friction** – Low-risk transactions proceed without interruption; medium-risk transactions might trigger additional verification, such as one-time passwords, push notifications or dynamic knowledge-based questions. High-risk transactions are blocked or held for manual review.
- **Reinforcement learning** – Models learn from outcomes (e.g., confirmed fraud, false positives) to adjust thresholds and actions. They optimize the trade-off between fraud prevention and customer convenience. As data volumes grow, reinforcement learning algorithms can fine-tune risk policies better than static rules.

Commercial Payment Fraud Lifecycle and Risk Distribution

Lifecycle stages and risk concentration

The payment lifecycle comprises several stages, each with different fraud risks:

- **Vendor onboarding and invoice submission** – Fraudulent vendors or invoices can be introduced. Risk mitigation includes vendor verification, duplication checks and behavioural profiling of invoices.
- **Accounts payable (AP) and enterprise resource planning (ERP)** – Changes to beneficiary details or payment amounts can occur. Segregation of duties, dual approvals and real-time validation controls reduce risk.
- **Digital banking and treasury initiation** – Users log into digital portals to initiate payments. Attackers may attempt ATO using stolen credentials. Multi-factor authentication, behavioural biometrics and device fingerprinting mitigate this.
- **Core banking system** – Final payment instructions are processed. Real-time monitoring and AI scoring ensure last-minute changes or suspicious instructions are flagged.
- **Payment rail execution** – Transactions are cleared through ACH, wire or instant payment networks. Payment rail-level controls (e.g., real-time risk scoring with credit limits) prevent outflow of fraudulent funds.
- **Reconciliation and settlement** – After payments are executed, reconciliation occurs. Post-transaction monitoring can identify anomalies missed earlier and support refund or recall efforts.

Risk distribution

The shows that digital banking and treasury initiation carry the highest risk, followed by accounts payable/ERP, Vendor onboarding and invoice submission, carry moderate risk, as do core banking systems and payment rail execution. Institutions should prioritize detection and controls at high-risk stages while not neglecting lower-risk areas, as fraud can shift.

Enabling upstream detection

Upstream detection is critical because fraud identified at invoice submission or AP approval prevents funds from leaving the account. AI models should ingest invoice data, supplier profiles, approval histories and user activity to detect anomalies early. For example, if a supplier's bank account changes, the model can cross-check with known patterns and flag the change before payment is initiated. This reduces reliance on last-minute wire checks and core banking controls.

AI Model Capability Framework

AI models for payment fraud vary in complexity, training data requirements and interpretability. QKS proposes a capability framework to guide selection and deployment:

Model type	Description	When to use
Rule-based heuristics	Predefined rules and thresholds based on known fraud patterns (e.g., payments above a certain amount).	Useful as a baseline, for regulatory compliance or when data is scarce. Should be combined with more advanced models.
Supervised classification models	Models trained on labeled fraud and non-fraud data. Examples include logistic regression, random forests, gradient boosting machines and neural networks.	When historical data is available. Effective at detecting known patterns but may miss novel fraud.
Unsupervised anomaly detection model	Models that detect outliers without labeled data. Algorithms include isolation forest, autoencoders, k-means clustering.	Useful for detecting unknown fraud and emerging patterns. Must be calibrated carefully to avoid excessive false positives.
Semi-supervised and weakly supervised models	Models trained on small labeled datasets and large unlabeled datasets (e.g., teacher-student models, self-training).	Appropriate when labels are expensive. Provide a balance between detection and false positives.
Graph models	Models that learn from graph structure (e.g., Graph Neural Networks,	Essential for detecting networked fraud, mule rings and collusion.

	PageRank) to detect suspicious link patterns.	
Reinforcement learning	Models that optimize decision policies by trial and error. They learn the best actions to minimize fraud and maximize customer experience.	Suitable for dynamic risk orchestration and adaptive authentication. Requires careful design to avoid undesired outcomes.
Generative AI and natural-language models	Large language models and generative models that create human-readable explanations, summarize alerts and assist in rule creation.	Useful for explainability, analyst productivity, customer communication and model tuning.

Training data and model management

Building robust AI models requires high-quality data. Labelled fraud data is scarce because fraud rates are low; synthetic data and augmentation techniques can expand training datasets. Data augmentation might include generating synthetic fraudulent transactions to simulate emerging patterns. Model management includes versioning, performance monitoring, retraining schedules and rollback processes.

Explainability and trust

Explainability is necessary for model adoption by regulators, auditors and risk committees. Techniques include:

- **Feature importance** – Identifying which features drive a model's prediction (e.g., new device, unusual amount).
- **Local explanations** – Methods like LIME or SHAP that explain individual predictions.
- **Generative narratives** – Natural-language explanations generated by large language models (LLMs) to describe the rationale behind decisions.
- **Counterfactuals** – Illustrating how small changes to inputs would change the prediction (e.g., if the amount were lower, the transaction would be approved).

Explainable models foster user trust and facilitate regulatory compliance.

AI ethics and fairness

AI systems must avoid bias and ensure fairness. Models trained on historical data may perpetuate discrimination. Best practices include:

- **Diverse training data** – Use data representing different customer demographics and behaviours.
- **Bias detection and mitigation** – Test models for disparate impact and adjust features, training data and thresholds. Use fairness metrics such as equalized odds or demographic parity.
- **Human oversight** – Retain human review for high-stakes decisions. Provide mechanisms for customers to challenge automated decisions.
- **Transparency and disclosure** – Inform customers when AI is used and how decisions are made.

Certification and standards

Obtaining independent certifications enhances credibility. Relevant standards include:

- **ISO 20022** – International standard for financial messaging. Certification demonstrates the enablement of richer data and interoperability.
- **ISO 27001** – Information security management. Certification demonstrates robust security controls.
- **SOC 2** – Service Organization Controls, ensuring controls over security, availability, processing integrity, confidentiality and privacy.
- **PCI DSS** – Payment Card Industry Data Security Standard, important for card-related payments.
- **GTF Certification** – Relevant for cross-border payments and compliance with SWIFT guidelines.

QKS AI Maturity Matrix™ and Market Landscape

Framework overview

QKS's AI Maturity Matrix™ assesses vendors across four levels AI Explorers, Building Momentum, Scaling for Impact, and Industry Pioneers—along three dimensions: AI vision and roadmap, AI maturity and AI-first productization. The matrix measures a vendor's ability to deliver predictive fraud detection, automate decision making and mitigate risk in real time.

- AI Vision & Strategy: How forward-thinking, intentional, and embedded is the vendor's AI roadmap across core insurance functions?
- AI Productization: How deeply is AI integrated into actual platform capabilities, and how scalable and accessible are those capabilities to insurers?
- AI Maturity: How sophisticated and adaptive a platform's AI is in handling complex fraud scenarios and emerging threats. How system uses advanced, explainable analytics that continually learn, integrate outside intelligence and operate across the full transaction lifecycle to deliver consistent value.

Vendors are placed into four categories:

- AI Explorers: Early in their AI journey, with limited production capabilities
- Building Momentum: Demonstrating growing investment but lacking end-to-end maturity
- Scaling for Impact: Possessing strong AI foundations with selective deployment across product lines
- Industry Pioneers: Vendors with embedded, scalable AI capabilities powering enterprise transformation

Mapping vendors on the AI Maturity Matrix

The diagram below places major vendors into the four maturity stages.

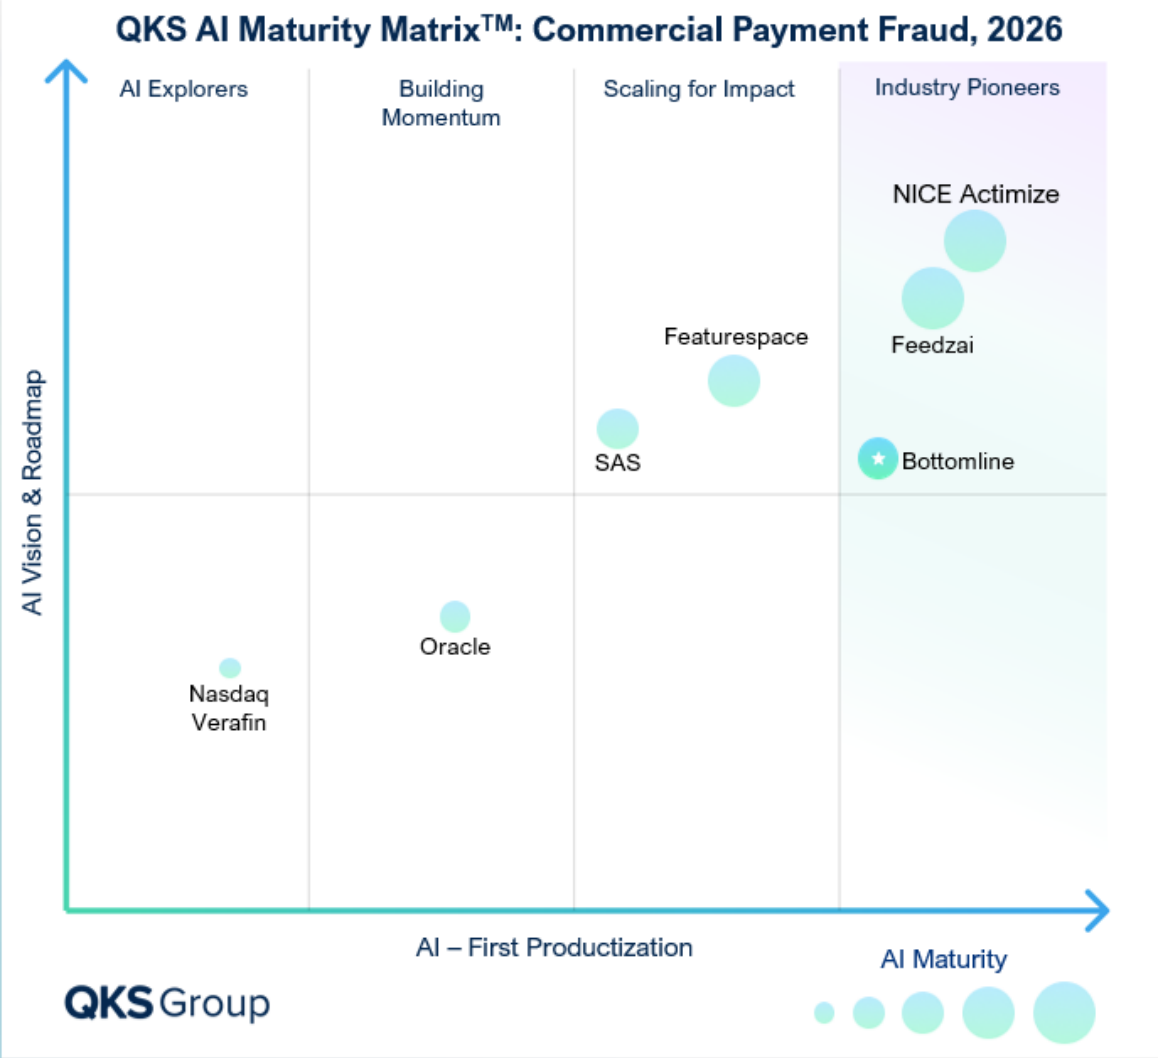

Vendors in the Industry Pioneers quadrant NICE Actimize, Feedzai & Bottomline exhibit the most advanced AI maturity, including self-optimizing models, generative explainability, consortium data integration and user-centric design. Vendors in Scaling for Impact Featurespace& SAS deliver strong AI models but lack some aspects of automated tuning or generative AI. Building Momentum vendors Oracle offer basic AI features but rely more on rules and conventional analytics. AI Explorers

include smaller players and in-house bank teams experimenting with AI but not yet deploying at scale.

Analysis of leading commercial payment fraud platforms

This section expands on the competitor analysis introduced earlier. For each vendor, we provide more granular insight into strengths, weaknesses and strategic direction:

Bottomline Market Position and Strategic Strengths

Bottomline delivers a **payments-centric financial crime solution** with strong capabilities in **commercial payments fraud detection and payment rail integration.** The solution is designed to operate across multiple payment

infrastructures, enabling institutions to apply consistent fraud controls across real-time and batch-based payment environments.

A key strength lies in its **adaptive AI and model tuning capabilities,** which allow institutions to continuously refine detection strategies based on evolving fraud patterns. The platform supports dynamic risk scoring, feedback-driven model optimization, and reinforcement-based learning approaches, enabling institutions to align detection strategies with changing threat landscapes.

Bottomline is also advancing its capabilities in **graph analytics, consortium data integration,** and generative AI, enabling improved contextual detection and enhanced explainability through natural-language outputs. Its **flexible deployment model (SaaS, managed services, and on-premises),** combined with strong user interface capabilities such as record-and-replay and link analysis, supports efficient investigation workflows and operational scalability.

Overall, Bottomline provides a **robust, payment-focused fraud platform,** well-suited for institutions seeking to strengthen controls across complex commercial payment ecosystems while progressively adopting advanced AI-driven capabilities.

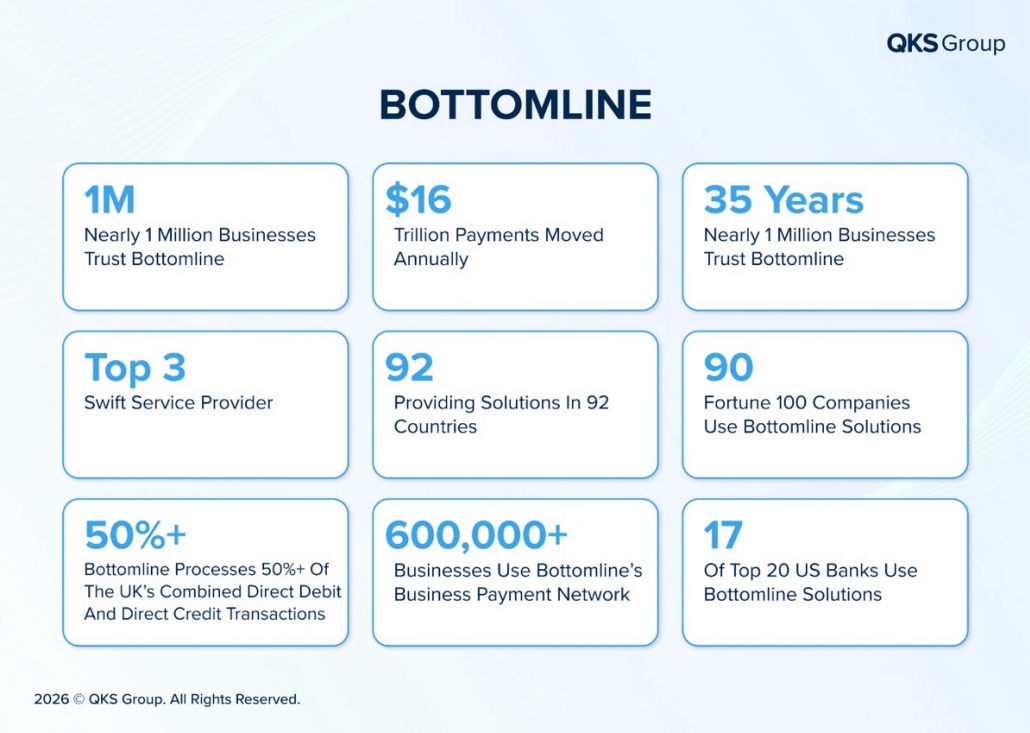

The company moves **US$16 trillion** in payments annually and serves **1 million businesses.** Its solutions underpin corporate banking, payments automation, cash management and fraud prevention for banks and corporates in 92 countries. This scale places Bottomline in a unique position to observe evolving fraud patterns and develop innovative defenses

"Bottomline sets the benchmark in commercial payments fraud by embedding adaptive, self-optimizing AI directly into the payment lifecycle – unifying behavioral intelligence, real-time interdiction, automated model tuning, and governed explainability into a purpose-built defense architecture for high-value commercial risk."

– ***Anandh Ramaswamy,*** *Practice Director, QKS Group*

This statement encapsulates QKS's view that Bottomline's Payments Fraud Defense platform leads the market in AI maturity. Bottomline unifies multiple layers of intelligence—behavioural analytics, graph analysis, real-time interdiction and automated tuning—directly within the payment lifecycle. The platform's extensible architecture allows rapid deployment of new AI models and dynamic tuning, enabling it to self-optimize as threats evolve. Governed explainability makes model decisions transparent and auditable, which is vital for compliance and trust. Bottomline thus delivers not only advanced technology but also practical, deployable solutions that reduce fraud losses and enhance operational efficiency, justifying its Most Valuable Pioneer status in our AI Maturity Matrix.

Payments heritage and scale

Bottomline's foundation in payments processing is its core strength. Serving high-value wire and ACH transactions, as well as real-time payment rails, allows Bottomline to embed fraud detection at the earliest possible stage. This heritage provides access to rich transaction and behavioral data. Large-scale adoption by **17 of the top 20 US banks and 50 %+ of UK direct debit/credit transactions** underscores market trust.

AI and innovation

Bottomline introduced predictive analytics and machine learning for fraud detection in **2014** and has continuously evolved its platform. The latest generation of Payments Fraud Defense features generative AI for onboarding and natural-language explanations, unsupervised anomaly detection, graph analytics and reinforcement learning. A key innovation is the **Fraud AI Tuning Agent,** which automates model optimization and produces AI-narrated performance reports. The roadmap includes consortium data integration and performance dashboards.

Integration and deployment

Bottomline's Payments Fraud Defense is pre-integrated with its digital banking and financial messaging solutionS, which accelerates time to value for customers. The solution is available as SaaS, managed service or on-premise, enabling customers to match deployment preferences. Integration connectors support core banking systems, treasury management systems, ERP, and payment rails. The extensible architecture ingests external intelligence and is API-driven, allowing for custom integration with third-party analytics and threat feeds.

Differentiators

When considering solutions for commercial payment fraud defense, Bottomline stands out for several reasons. The following differentiating factors highlight how Bottomline's capabilities compare favorably to other providers in this space:

- **Deep payment workflow integration** — Bottomline enables intervention at every stage of the payment lifecycle, from invoice creation and accounts payable approval to digital banking initiation and core banking settlement. This upstream detection helps prevent fraud before funds are moved, offering a proactive approach that distinguishes Bottomline.
- **Adaptive AI and automated tuning** — The platform's self-learning models continually refine fraud detection and minimize false positives. This adaptability ensures ongoing improvements and responsiveness to evolving threats, setting Bottomline apart in its ability to deliver reliable results.

- **Generative AI and explainability** — Natural-language explanations of AI decisions foster transparency and build user trust, while meeting regulatory requirements. This commitment to clarity and compliance is a notable advantage compared to other solutions.
- **Record & replay and link analysis** — Investigations are streamlined with visual tools that allow users to review and analyze transactions, providing clear evidence and supporting faster resolution of fraud cases.
- **Comprehensive roadmap** — Bottomline is dedicated to ongoing innovation, with plans to expand features such as consortium data integration, advanced AI tuners, performance dashboards, and additional AI-driven enhancements. This forward-thinking approach reflects a strong commitment to continual improvement.
- **SaaS and managed service offerings** — Flexible deployment options, including SaaS and managed services, enable rapid implementation and reduce operational demands, making Bottomline an appealing choice for institutions seeking efficiency and scalability.

Industry recognition and case studies — Bottomline's solutions have earned multiple industry accolades and have demonstrated success in preventing high-value fraudulent wire transfers. These proven outcomes reinforce its reputation as a trusted partner for fraud prevention.

Case Studies and Testimonials from Bottomline

01 Case Study
Large global bank: Cross-channel fraud orchestration

- **Background:** A large global bank handling $2 billion in daily wire transactions experienced multiple cases of BEC and APP fraud. Attackers compromised email accounts and inserted malicious payment instructions into invoices. The bank used separate monitoring tools for ACH and wire, resulting in inconsistent risk scoring and delayed detection.
- **Implementation:** Bottomline's commercial payments fraud solution was integrated across digital banking, ACH and wire platforms. Real-time data ingestion allowed models to detect anomalies across channels. Risk orchestration triggered step-up authentication for medium-risk transactions and blocked high-risk payments pending review. Investigators used record & replay to analyse suspicious sessions and link analysis to identify colluding accounts.
- **Results:** Within nine months, the bank reduced wire fraud losses by 60 %, decreased false positives by 45 %, and improved customer satisfaction scores. The bank's fraud investigation team handled 25 % more alerts without additional staff. Compliance teams used generative AI explanations to respond to regulatory inquiries swiftly and accurately.

02 Case Study
Regional credit union: Managed service transformation

- **Background:** A mid-sized credit union suffered from account takeover attempts targeting its business members. It lacked the resources to deploy and manage complex AI models. False positives caused delays in legitimate payroll and invoice payments, damaging member satisfaction.
- **Implementation:** The credit union adopted Bottomline's commercial payments fraud solution. Models scored events in the provider's secure cloud, with results delivered back to the credit union's digital banking system in real time. The credit union enabled generative AI narratives to explain alerts and tuned parameters through QKS consultation.

- **Results:** Within six months, the credit union reduced false positives by 38 % and account takeover attempts by 70 %. Investigators received clear AI-generated explanations that improved confidence in decision making. The credit union gained the ability to focus on relationship management rather than fraud detection.

03 Case Study
Supply chain company: Invoice fraud prevention

- **Background:** A multinational corporation with thousands of suppliers across Asia and Europe wanted to prevent invoice fraud and vendor impersonation. Fraudsters were sending fake invoices with altered bank details after intercepting email communications between buyers and suppliers.
- **Implementation:** The company integrated Bottomline's commercial payments fraud solution at the invoice ingestion stage within its ERP. AI models analysed invoice metadata, amounts, item descriptions, vendor history and email patterns. Any deviation from expected patterns triggered an alert. The models also correlated supplier bank changes with known fraud risk indicators.
- **Results:** The company identified and blocked 12 fraudulent invoices totaling $8 million in the first quarter. Vendor onboarding validation improved, ensuring only legitimate suppliers were paid. Link analysis revealed that certain suppliers were targeted more frequently, leading to additional training and security measures.

04 Case Study
Government treasury: Real-time payroll protection

- **Background:** A government treasury responsible for disbursing wages to public-sector employees across multiple jurisdictions faced frequent attempts by criminals to redirect payroll payments. Attackers used phishing emails to capture login credentials and changed payment details to their own accounts.
- **Implementation:** Bottomline's commercial payments fraud solution was deployed in the treasury's payment initiation layer. User behaviours were profiled across hundreds of thousands of employees. AI models flagged unusual IP addresses,

device changes and beneficiary updates. Step-up authentication via government-issued identification was triggered for high-risk transactions.

- **Results:** Fraudulent payroll transactions were reduced by 90 %. The treasury saw a 50 % reduction in manual investigations because false positives dropped dramatically. Citizens received timely payments without disruption.

These case studies demonstrate the versatility of Bottomline's commercial payments fraud solution across sectors and highlight the value delivered when QKS and Bottomline collaborate closely with clients.

Other Vendors Evaluated On Commercial Payment Fraud Platform

NICE Actimize

NICE Actimize delivers a **comprehensive, enterprise-grade financial crime platform** with strong capabilities across fraud, AML, and compliance. The platform is widely adopted by global financial institutions and supports complex transaction environments, particularly in commercial banking.

A key strength lies in its **end-to-end coverage across the financial crime lifecycle,** including onboarding risk, transaction monitoring, and investigation workflows. The platform supports a wide range of fraud typologies and operates across multiple payment rails.

NICE Actimize emphasizes an **explainability-first approach to AI,** providing transparent and auditable model outputs. Its capabilities in behavioral analytics, risk scoring, and case management support both detection effectiveness and regulatory compliance.

The platform is designed for **high-performance, real-time processing at scale,** enabling institutions to manage large transaction volumes without impacting customer experience. Its structured approach to model governance and operational stability supports consistent performance across evolving risk environments.

Overall, NICE Actimize provides a **robust and scalable financial crime platform,** combining regulatory alignment, operational maturity, and advanced analytics to support complex banking operations.

Feedzai

Feedzai delivers an **AI-native financial crime prevention platform** with strong capabilities in **real-time risk decisioning and behavioral intelligence.** The platform processes large volumes of transactional and behavioral data, enabling institutions to detect fraud across multiple channels with high accuracy.

A key strength lies in its **advanced machine learning and behavioral biometrics capabilities,** which combine transaction data, device intelligence, and user behavior to generate highly contextual risk scores. Its streaming architecture supports real-time analysis of historical and current data, enabling dynamic risk assessment at the point of transaction.

Feedzai also provides a **unified, omnichannel view of risk,** integrating signals across payments, digital banking, and external data sources. Its approach enables detection of complex fraud scenarios, including social engineering, account takeover, and cross-channel attacks.

The platform supports strong **model governance and explainability,** with transparent decision logs, audit trails, and performance monitoring capabilities. Its focus on operational efficiency through intelligent alert prioritization enhances investigator productivity and reduces unnecessary friction.

Overall, Feedzai offers a **highly adaptive, AI-driven platform,** enabling institutions to respond effectively to evolving fraud threats in real-time, digital-first environments.

Featurespace

Featurespace offers a **behavioral analytics-driven fraud detection platform,** with strong capabilities in **real-time anomaly detection and adaptive behavioral modeling.** The platform leverages unsupervised machine learning to identify deviations from normal activity patterns, enabling detection of both known and emerging fraud typologies.

A key differentiator is its **adaptive behavioral analytics engine,** which continuously learns from transactional patterns to generate highly contextual risk scores. This approach allows institutions to detect subtle behavioral anomalies while maintaining low false-positive rates, supporting both fraud prevention and customer experience.

Featurespace is recognized for its ability to deliver **real-time decisioning with low latency,** making it particularly effective in high-volume transactional environments. The platform is expanding its capabilities into broader payment ecosystems, including commercial payments, while continuing to build on its strong foundation in transaction-level behavioral analytics.

Overall, Featurespace provides a **specialized, analytics-focused platform,** enabling institutions to adopt a behavioral and intelligence-led approach to fraud detection across dynamic transaction environments.

SAS Fraud Management

SAS Fraud Management delivers a **comprehensive analytics platform** built on strong foundations in **data management, statistical modeling, and advanced analytics.** The platform supports a wide range of fraud and financial crime use cases, with particular strength in enabling institutions to design and deploy **customized detection strategies.**

A key strength lies in its ability to provide **flexible model development and advanced feature engineering capabilities,** allowing institutions to tailor detection frameworks to specific business requirements. Its robust data integration capabilities enable the processing of large volumes of structured and unstructured data across enterprise environments.

The platform supports a combination of **rule-based systems, predictive modeling, and machine learning techniques,** enabling institutions to apply multi-layered detection strategies. SAS is particularly well suited for organizations with strong in-house data science capabilities seeking greater control over analytics and model optimization.

Overall, SAS Fraud Management offers a **highly flexible and scalable analytics platform,** supporting deep customization and enterprise-wide deployment of financial crime detection strategies.

Oracle Financial Services Analytical Applications

Oracle Financial Services Analytical Applications (OFSAA) provides a **broad, integrated platform** for financial crime, risk, and compliance management. The platform combines AML, fraud detection, and regulatory reporting capabilities within a unified data and analytics framework.

A key strength lies in its **deep integration with core banking systems and enterprise data environments,** enabling institutions to leverage existing infrastructure for financial crime detection and compliance processes. This integration supports consistent data usage across multiple risk and business functions.

Oracle offers a wide range of **analytical and reporting capabilities,** allowing institutions to implement risk models, monitor transactions, and meet regulatory requirements within a single platform. Its architecture supports large-scale deployments across global banking operations, making it suitable for complex, multi-entity institutions.

Overall, Oracle OFSAA delivers a **comprehensive and integrated financial crime solution,** with strong capabilities in data consolidation, enterprise analytics, and regulatory alignment.

Nasdaq Verafin

Nasdaq Verafin provides a **cloud-native financial crime platform** focused on delivering **integrated fraud and AML capabilities** through a SaaS-based model. The platform is designed to simplify financial crime operations while enhancing detection through collaborative intelligence.

A key differentiator is its **consortium data network,** which enables participating institutions to identify fraud patterns across multiple organizations. This shared intelligence approach enhances visibility into broader fraud schemes and supports detection of cross-institutional risk.

The platform offers strong capabilities in **case management, alerting, and regulatory reporting,** allowing institutions to streamline investigation workflows and improve operational efficiency. Its SaaS delivery model supports scalability and ease of deployment, making it accessible across a wide range of institutions.

Overall, Nasdaq Verafin delivers an **intelligence-driven, cloud-based financial crime platform,** enabling institutions to leverage shared insights and integrated workflows to enhance fraud and AML effectiveness.

Future Outlook and Innovation

Emerging technologies

- **Quantum computing** – Quantum computers will eventually challenge current cryptography. Payment security protocols must be quantum-resistant. Fraud detection models may need to adapt to new encryption techniques and analytic challenges.
- **Blockchain and distributed ledger** – Payment networks built on blockchain can reduce settlement risk but introduce new fraud patterns (e.g., double spending, smart-contract exploits). AI models must monitor blockchain transactions and smart contract interactions.
- **Central bank digital currencies (CBDCs)** – CBDCs will create new payment rails and require fraud detection mechanisms. AI models must handle CBDC-specific transaction metadata and anti-money laundering (AML) requirements.
- **Edge AI** – Deploying AI at the edge (e.g., on mobile devices or POS terminals) can enable faster, local fraud detection while protecting privacy. Federated learning techniques may allow models to learn from dispersed data without centralizing it.
- **Synthetic biology and deepfakes** – Biometric spoofing, such as fake fingerprints or deepfake facial recognition, may challenge authentication systems. AI must detect synthetic biometrics and differentiate them from genuine ones.

Evolution of regulation

Regulators will continue to update payment and AI guidelines. PSD3 will likely impose stricter requirements on third-party providers and expand consumer protection. AML and KYC standards will require more granular risk analysis and faster reporting. Data-protection laws will evolve, emphasizing privacy-preserving AI and consent management. Institutions must stay agile and adapt models and processes accordingly.

AI democratization and commoditization

AI components will become commoditized through open-source libraries and cloud APIs. Differentiation will shift from model performance to data quality, domain expertise, integration depth, and user experience.

Focus on ESG and Sustainability

ESG considerations are gaining attention in financial services. Fraud prevention ties into governance (G) by ensuring ethical and compliant operations. Institutions will need to align fraud systems with broader ESG commitments, such as reducing fraud related to environmental crime (e.g., illegal wildlife trafficking) and ensuring fair access to financial services.

Research Methodology

QKS employed a multi-faceted research approach:

- **Primary research** – Interviews with fraud experts, payment specialists, regulators, and technology leaders. Site visits to banks and corporates provided context on operational challenges and technology adoption.
- **Secondary research** – Review of industry reports, regulatory publications, vendor materials, and academic papers. Analysis of global payment statistics and fraud trends.
- **Vendor surveys** – Structured questionnaires sent to leading vendors to gather information on product capabilities, AI maturity, deployment models, pricing and roadmaps.
- **Case studies** – Detailed investigation of multiple real-world deployments across banks, credit unions, corporates and government treasuries. QKS worked closely with clients to quantify results and lessons learned.
- **AI maturity assessment** – QKS applied the AI Maturity Matrix to evaluate vendor capabilities and customer adoption levels. Vendors were scored based on vision, maturity and productization.
- **Market modelling** – Quantitative models to estimate market size, growth rates and ROI. Scenario analysis explored different adoption and attack patterns.
- **Peer review** – Drafts were reviewed by industry experts, academic advisors and regulatory advisors to ensure accuracy and objectivity.

Limitations and assumptions

- Data on fraud losses and detection rates are often confidential; assumptions were made based on anonymized case studies and industry benchmarks.
- ROI models are illustrative and based on average industry metrics. Actual ROI depends on specific institution factors such as transaction volume, internal processes, and fraud patterns.
- Competitive analysis is based on public information and QKS surveys. Vendors continuously innovate; capabilities may change after publication

Conclusion

The commercial payments landscape is undergoing profound transformation. Real-time rails, open banking and digital commerce are accelerating growth but also introducing new fraud risks. Traditional rule-based systems cannot keep pace with agile, persistent adversaries. AI provides the means to detect, prevent and respond to fraud in real time, but its adoption requires a clear maturity roadmap, organizational readiness, and ethical governance.

Bottomline stands out as a leader in AI maturity for commercial payment fraud, leveraging its payment heritage, adaptive models, generative AI, record & replay tools, flexible deployment and extensive roadmap. By continuing to invest in consortium data, reinforcement learning, partner ecosystem integration and user-centric design, Bottomline can maintain its leadership and serve as a trusted partner for financial institutions worldwide. This white paper aims to inform and inspire stakeholders—banks, corporates, regulators and fintechs—to pursue advanced fraud defenses and to collaborate in building a safer, more efficient payment ecosystem.

www.qksgroup.com

www.bottomline.com