

IMPETUS

WHITEPAPER

Redefining Financial Services with **Context-Aware Agentic AI**

With a spotlight on next-gen personalization
and fraud detection



Introduction

The financial services sector is undergoing significant changes due to advancements in technology, with agentic AI playing a crucial role. Agentic AI refers to AI systems that act autonomously across multi-step workflows, with enterprise teams stepping in selectively via 'human-in-the-loop' oversight rather than directing every action. This is particularly useful in the financial services sector, where speed, accuracy, and compliance are crucial. By using agentic AI, financial institutions can improve efficiency, enhance customer experiences, and stay competitive in a fast-changing market.

This paper examines the impact of agentic AI on financial services and insurance, focusing on its ability to deliver solutions that meet regulatory requirements while elevating business outcomes. The report aims to provide a clear understanding of how agentic AI can be effectively used in specific use cases like personalization and fraud mitigation, and how context engineering – designing the information environment and boundaries within which AI agents operate – is the key to making agentic AI effective and safe for banks and insurers.

01 LEVERAGING AGENTIC AI TO NAVIGATE INDUSTRY CHALLENGES AND OPPORTUNITIES

The financial services sector is facing mounting pressure from multiple directions. Incumbent institutions must modernize legacy technology estates while complying with increasingly stringent regulatory and data governance requirements. At the same time, customer expectations are being reset by digital-first competitors that deliver always-on, hyper-personalized experiences. Rising fraud sophistication, shrinking margins, and the need to extract value from fragmented data assets further intensify the challenges. Together, these forces are exposing the limits of traditional, rules-based operating models and driving demand for more adaptive, intelligent, and context-aware solutions.

Around these structural challenges, cybersecurity has become a defining concern for financial institutions. Banks and insurers face a constant barrage of threats ranging from credential stuffing and account takeovers to sophisticated deepfakes and synthetic identities. Legacy controls and static, rule-based fraud systems cannot keep pace with highly automated, adaptive attackers who increasingly exploit AI. The result is a widening gap between traditional cyber/fraud defenses

“ In 2025, global losses to fraud and scams reached an estimated **US\$579 billion** — a **9.2%** increase from 2023, driven in large part by criminals leveraging AI to scale their attacks.ⁱ

and the real-time, intelligence driven protection now required to safeguard customers, preserve trust, and meet tightening regulatory expectations.

This underscores the urgent need for more dynamic and autonomous anti-fraud measures. Agentic AI can help institutions navigate these pressures by accelerating threat detection and response, streamlining compliance, and enabling personalized services. The table below maps three major challenge areas to how agentic AI responds — and the associated business benefits.

CHALLENGES	HOW AGENTIC AI RESPONDS	BUSINESS BENEFITS
Vulnerabilities and fraud risks	Proactively detects and responds to threats, continuously analyzing vast data to surface patterns and anomalies that indicate fraud — so teams can act promptly to protect assets and customers.	Stronger fraud protection Reduce false positives, accelerate threat response, and safeguard customers and assets from emerging fraud threats.
Regulatory and compliance uncertainties	Transparent, auditable agents streamline end-to-end compliance. Embedding AI governance from the start speeds deployment, reduces regulatory risk, and reinforces the institution's reputation as trustworthy.	Lower compliance risk Reduce regulatory risk and penalties while improving reporting accuracy, auditability, and operational efficiency.
Balancing personalization with ethical data use	Transforms customer data into personalized, contextual recommendations while ensuring ethical use and regulatory compliance — deepening relationships and differentiating from competitors.	Improved customer experience Deliver more relevant products and advice, deepen relationships, and increase engagement, conversion, and loyalty.

02 THE RISE OF AGENTIC AI IN BANKING AND INSURANCE

Forward-looking banks and insurers recognize that while fully autonomous finance remains an evolving vision, now is the time to prepare. Industry assessments suggest organizations should begin designing AI workflows, systems, and governance with agentic AI in mind to avoid creating 'organizational straitjackets' that constrain future automationⁱⁱ.

The strategic risk of not preparing is that competitors or new entrants, unencumbered by legacy processes, will leap ahead with highly automated workflows. The opportunity lies in piloting agentic AI in areas of greatest need – such as functions with low customer satisfaction or high unit costs – and then scaling adoption as trust and efficacy are proven in practice.

AGENTIC AI IN FINANCIAL SERVICES USE CASES

Context-aware
fraud detection

Credit risk memo
generation

KYC/AML onboarding
automation

Always-on
wealth advisory

Regulatory
compliance
monitoring

Personalized customer
engagement

Autonomous
threat detection
and response

Dispute and
chargeback
resolution

Contact center
augmentation

Underwriting and
loan servicing

Portfolio and risk
monitoring

Financial planning
and analysis

FROM AUTOMATION TO AUTONOMY

Financial services and insurance have long used automation such as robotic process automation (RPA) and traditional ML models to improve decisioning and operations. Such approaches rely on predefined rules, static datasets, or narrowly scoped models embedded into fixed workflows. They excel in stable environments but struggle when tasks require multi-step reasoning, dynamic planning, adaptation to novel scenarios, or real-time action across multiple systems. Unlike conventional AI chatbots and assistants, AI agents can independently make decisions and take actions without human intervention, giving financial institutions a competitive edge.

“ Unlike conventional AI chatbots and assistants, AI agents can **independently make decisions and take actions without human intervention**, giving financial institutions a competitive edge.

THE JOURNEY OF EVOLUTION FROM TRADITIONAL AUTOMATION TO AGENTIC AI

TECHNOLOGY	WHAT IT DOES	AUTONOMY LEVEL
Conventional automation (RPA)	Executes predetermined steps with little adaptability. Excellent for structured, repeatable processes; unable to handle exceptions or learn from new data.	Rules-only
Traditional AI / ML point solutions	Provide predictions or classifications (a credit risk score, a fraud likelihood) to inform human decisions — stopping at decision support rather than end-to-end action.	Decision support
Generative AI assistants (copilots)	Use large language models (LLMs) to generate content or retrieve information in natural language, helping people answer questions or draft documents. They speed up work but usually need a human to finalize decisions.	Human-finalized
Context-aware AI agents	Plan and execute multi-step workflows by dynamically integrating with tools and systems while adapting to context in real time. They pursue objectives (e.g., ‘investigate this claim for fraud’) subject to governance guardrails and human-oversight triggers.	Goal-driven & autonomous

In practice, this evolution signifies a major shift from AI assistants (which answer questions) toward AI agents (which achieve outcomes). Core features distinguishing agentic AI include autonomy, goal-orientation, adaptability, real-time decision-making, self-optimization, and the ability to collaborate in multi-agent systems. These capabilities are enabled by a stack of technologies and practices, rather than a single AI model. Key components of this stack include:

ADVANCED DECISIONING MECHANISMS: Agentic systems use traditional ML and reinforcement learning (RL) to support autonomous decision-making. RL helps agents improve through trial-and-error feedback, reducing reliance on large, labelled datasets in volatile financial environments. For example, trading agents can refine strategies using profit-and-loss signals, while fraud detection agents improve through feedback from confirmed cases.

ACCESS TO REAL-TIME, TRUSTED DATA: Agentic AI relies on real-time data signals rather than static historical datasets. Use cases such as point-of-sale credit decisions, instant fraud prevention, and mobile app personalization require millisecond-level decisions based on the freshest data. This requires modernized legacy systems and data pipelines so AI agents can securely access streaming information.

API ECOSYSTEMS AND INTELLIGENT ORCHESTRATION: AI agents need to connect with internal systems such as core banking, CRM, payments, claims, AML/KYC tools, and external services like identity verification and credit bureaus. APIs enable agents to retrieve data and execute transactions across platforms, while orchestration logic determines which

tools to use, manages exceptions, and coordinates multi-agent workflows. Strong orchestration improves efficiency, security, and auditability, but broader API connectivity also increases the attack surface. Therefore, every API call must be secured, monitored, and logged.

CLOUD AND HYBRID ARCHITECTURE: The scale and complexity of agent-driven interactions are pushing financial institutions toward cloud, hybrid-cloud, and multi-cloud strategies. Cloud platforms offer scalable capacity and AI services to support large volumes of agent actions. These capabilities help banks adopt agentic AI at scale while using hybrid setups to keep sensitive data on-premises. Cloud architecture also supports multi-agent communication, continual model updates, and advanced monitoring, while balancing innovation with data residency and security requirements.

DATA QUALITY AND GOVERNANCE: AI agents depend on accurate, well-governed data and clear rules. Poor, siloed, or inconsistent data can lead to errors or hallucinations, so financial firms are improving data quality, lineage, and governance before scaling agentic AIⁱⁱⁱ. This includes modernizing legacy data systems, standardizing business definitions, and enforcing policies on privacy, usage rights, and retention. Many banks are also using data catalogs and knowledge graphs to give AI agents better business context. Strong AI governance frameworks log decisions, enable audits, require human approval for high-risk actions, monitor performance, and provide kill switches if agents behave unexpectedly. These controls help build trust with executives, regulators, and customers.

RIGHT-SIZED MODELS FOR THE RIGHT TASK

Not every task needs a large language model. Financial institutions are combining SLMs and domain-specific models with LLMs, choosing the right model based on accuracy, speed, cost, and complexity.

Wells Fargo's Fargo assistant, for example, combines a Google language model for intent recognition with smaller proprietary models for tasks such as detecting PII and managing workflows^{iv}.

Fintechs like Lightico use lightweight models trained on banking documents to automate ID verification and loan processing more accurately and cost-effectively^v.

The result: greater accuracy, lower costs, and more efficient AI at scale.

MARKET DRIVERS: WHY AGENTIC AI IS RISING NOW

Years of digital transformation and cost reduction efforts are pushing institutions beyond chatbots toward agents that can automate end-to-end workflows, not just triage customer requests. Mounting business, technology, regulatory, and competitive pressures are accelerating the adoption of agentic AI across banking and insurance. Here are some of the key forces driving this shift:

Data and Reporting Demands Intensify

The push toward agentic AI is raising the bar for trusted, well-governed data. At the same time, regulatory requirements such as BCBS 239 demand accurate, timely, and auditable risk data aggregation and reporting. For institutions still operating across fragmented data estates and siloed systems, these pressures are accelerating efforts to build unified, governed data foundations that can support both regulatory compliance and AI at scale.

Regulation Complexities Raise the Stakes

Regulatory scrutiny of AI is intensifying. In Europe, the EU AI Act classifies many financial AI applications, including credit scoring and lending decisions, as high-risk, requiring transparency, human oversight, and documentation^{vi}.

In the United States, regulators including the Federal Reserve, OCC, CFPB, and FTC are applying existing risk management and consumer protection requirements to AI. Banks must validate models, ensure explainable decisions, and maintain robust audit trails. Institutions that build strong AI oversight into their systems will be better positioned to innovate than those that treat governance as an afterthought.

Competitive Pressure Accelerates Adoption

Digital-native challenger banks have demonstrated how automation can enable leaner operations. At the same time, fraud and security threats are escalating as attackers use AI to increase their speed and sophistication. Financial institutions increasingly need agentic AI on both offense—delivering personalized, instant service—and defense, enabling faster risk detection and response.

Early movers are already demonstrating what's possible. Wells Fargo's AI assistant handled 245 million interactions in 2024 without human hand-offs, using a privacy-first design and multiple models for accuracy. This illustrates how AI can operate at massive scale while maintaining control.

CYBERSECURITY EFFORTS IN BANKING & PAYMENTS

EMERGING TECHNOLOGIES AND INITIATIVES


- Proactive banks are using new technologies and ideas to enhance their cybersecurity practices.
- Some of these investments are in their early stages and are not widely adopted.

INDUSTRY BEST PRACTICES

- Developing bespoke internal standards for cyber risk management and incident response can help bolster cybersecurity compliance and reassure stakeholders

REGULATION

- Many countries mandate that banks must comply with a laundry list of cybersecurity regulations, with consequences such as fines or imprisonment for non-compliant companies.
- While these help banks ensure data and transactions are secure, they usually overlap with each other, creating confusion and inefficiencies.



03 THE POTENTIAL OF AGENTIC AI-DRIVEN HYPER-PERSONALIZATION IN FINANCIAL SERVICES

Traditional customer service models are limited in meeting evolving customer expectations as they can be slow and reactive, with issues being moved between siloed departments. Agentic AI tackles these limits with a 'Customer 360' approach that connects dynamic customer journeys, multiple data sources, and conversational banking to improve engagement, satisfaction, and product uptake. The opportunities include:

- **RICHER PERSONALIZATION:** Using data from customer relationship management, pricing, servicing, and digital channels to respond to the full customer context.
- **SEAMLESS, CONSISTENT EXPERIENCES:** Leveraging APIs and shared data to attain a single view of the customer across channels, including chat, voice, and mobile.
- **PROACTIVE, END-TO-END ASSISTANCE:** Helping customers spot issues such as likely overdrafts or expiring rewards with agents that act across systems and proactively suggest fixes.
- **SCALABILITY WITH CONTROL:** Cloud platforms, small language models, and strong data governance enable this at scale with clear controls and auditability.

STRATEGIC USE CASES

01

PROACTIVE CREDIT LIMIT MANAGEMENT


A retail bank uses an agent that monitors card utilization and repayment patterns, draws on income and tenure data, and then proactively offers tailored credit limit increases before customers hit their limit. This reduces decline rates and call center traffic, while improving satisfaction.

02

CONTEXT-AWARE REWARDS AND TRAVEL

A card issuer links travel bookings, loyalty balances, and expiry dates so an agent can suggest seat upgrades or lounge access funded by soon-to-expire points. This drives an increase in engagement and redemption rates.

IMPETUS' SCALABLE REFERENCE ARCHITECTURE



04 THE POTENTIAL OF AGENTIC AI IN FRAUD DETECTION AND COMPLIANCE

FROM REACTIVE FRAUD DETECTION TO REAL-TIME DEFENSE

As fraud techniques grow more sophisticated, legacy systems that rely on static rules or after-the-fact checks are increasingly insufficient. Agentic AI offers a more proactive and adaptive approach to risk mitigation. AI agents can monitor multiple data streams simultaneously—transactions, device metadata, geo-location, merchant information, network traffic, and more—and cross-correlate signals to flag anomalies in real time.

By continuously learning from new fraud patterns and incorporating external intelligence, these agents become more effective over time. They can also reduce false positives by using contextual information to distinguish unusual but legitimate behavior from truly suspicious activity.

ACCELERATING AML INVESTIGATIONS AND REGULATORY COMPLIANCE

Similar approaches can strengthen risk management and compliance. In anti-money laundering, for example, an AI agent might not only flag a suspicious transaction but autonomously gather supporting information such as past account activity, related entities, and relevant news, then draft a report for investigators.

In regulatory compliance, agents can enforce rules by design—preventing trades that violate limits or ensuring customer communications stay within approved guidelines, escalating to humans only when needed. Embedding compliance into automated workflows helps institutions reduce errors and respond faster to regulatory changes.

MOVING HUMANS TO HIGHER-VALUE DECISIONS

Real-world use cases are already demonstrating these benefits. Agentic AI doesn't replace humans in fraud and risk; it shifts them from frontline detection to complex, strategic cases. By automatically catching more issues earlier, it reduces losses, better protects customers, and supports regulatory compliance.

Regulators generally support advanced analytics with appropriate controls, such as MAS guidance on AI in fraud and AML. But to fully unlock agentic AI's potential, banks and insurers must bridge the gap between what AI knows and what their enterprise knows—which is discussed in the next section.

STRATEGIC USE CASES

01

MULTI-SIGNAL FRAUD DETECTION

A large bank adds an AI 'brain' on top of its existing fraud tools that fuses many signals in real time: transaction patterns, device fingerprints, merchant risk, and identity/behavioral data. The agent builds a risk score per transaction and follows playbooks:

- **Low risk:** let it through.
- **Medium risk:** send extra verification (e.g., text or biometric).
- **High risk:** place an immediate hold and alert a human analyst, with a rich AI-generated summary.

This multi-signal AI layer catches more fraud faster, reduces false positives, and cuts manual reviews by up to 50%^{vii}, while keeping legitimate customer transactions smooth.

02

REAL-TIME INTERVENTION IN PAYMENTS

AI agents in a bank's payments system act as a real-time fraud 'forcefield' during high-risk transactions. While a customer is on a call trying to wire a large amount (potentially to a scammer), different agents simultaneously:

- Analyze the live call transcript for signs of coaching, pressure, or distress.
- Check the recipient's account against fraud databases for red flags.
- Assess the transaction context (time, place, and amount vs. customer history).

If risk is high, the system instantly alerts bank staff, temporarily blocks the payment, and prompts additional questions. In the example, AI runs continuously in the background, intervening only when needed and always paired with human oversight.

05 CLOSING THE CONTEXT GAP — THE KEY TO EFFECTIVE AGENTIC AI

Powerful model capabilities alone cannot ensure reliable AI outcomes. AI agents must understand the proprietary data, unstructured information, operational logic, regulatory constraints, and domain expertise of an institution. This context-awareness is crucial for minimizing hallucinations and delivering accurate, trusted AI outcomes. Most AI pilots fail not because the model is wrong, but because the context surrounding it is fragmented, stale, or structurally invisible to AI.

“ In a highly regulated financial services environment, **context is what takes agentic AI beyond routine automation and turns autonomous decisions into trusted action.** ”

UNDERSTANDING THE CONTEXT GAP

The Context Gap is the distance between what today's AI models know and what makes your enterprise unique. It has four distinct layers:

THE DATA GAP

Legacy, siloed, and inconsistently governed data that AI systems cannot reliably access or interpret. Closing this gap requires modernizing data estates (including schemas, pipelines, lineage, quality & entitlements) so agents have a trusted, AI-ready data foundation to reason with.

THE SEMANTIC GAP

Generic AI models do not natively understand a firm's products, policies, decision logic, domain specifics, or regulatory constraints. Techniques such as retrieval-augmented generation, semantic modelling, and knowledge graphs are needed so responses reflect the institution's own logic rather than generic answers.

THE EXECUTION GAP

Agents lose context across steps, tools, and hand-offs, leading to drift, errors, or hallucinations in longer workflows. Addressing this calls for well-designed multi-agent workflows, a unified context layer, and intelligent orchestration before agents are exposed to real customers or transactions.

THE TRUST GAP

Without the right guardrails, monitoring, and auditability in place, agentic AI cannot be scaled safely. Organizations need strong governance frameworks and clear policies on autonomy and human-in-the-loop oversight, so that all AI decisions are explainable.

Bridging these four gaps can help financial institutions move from pilots that impress in demos to agentic AI initiatives that translate into real-world business impact and ROI. They can start by selecting strategic use cases, assessing their current AI-readiness, and designing a shared, robust context layer. From there, they can prototype intelligent agentic workflows, measure outcomes, refine controls, and then scale successful patterns into wider use across the enterprise.

06 HOW IMPETUS BRIDGES THE CONTEXT GAP

Impetus Technologies – a leader in software-driven data and AI services – helps enterprises close the context gap and translate agentic AI from concept to real-world outcomes. Through its Leap™ AI Solutions & Services Family, Impetus delivers a repeatable context engineering methodology that addresses each layer of the context gap. This approach ensures that autonomous AI agents are grounded in the enterprise's data, knowledge, and policies from day one.

“Through its Leap™ AI Solutions & Services Family, Impetus delivers a **repeatable context engineering methodology** that addresses each layer of the context gap.

IMPETUS LEAP AI SOLUTIONS, PURPOSE-BUILT TO POWER THE CONTEXT-AWARE ENTERPRISE



MODERNIZE

Impetus LeapLogic™ Suite

Transforms legacy data estates (data warehouses, ETL, mainframes, BI) into AI-ready foundations with an ontology-led approach.



ENGINEER

Impetus Context Fabric™

The semantic foundation connecting ontology, knowledge graphs, business rules, policies, lineage, and memory into enterprise context.



ORCHESTRATE

Impetus Agent Solutions™

Domain-specific Skills, Kits, and Agents that automate decision-heavy workflows across business and technology operations.



GOVERN

Impetus Prism™

A unified control plane for AI observability, tokenomics, governance, and intelligent DataOps.

HOW THESE SOLUTIONS COME TOGETHER TO DELIVER CONTEXT-AWARE AGENTIC AI

MODERNIZING LEGACY DATA INTO AN AI-READY FOUNDATION:

Using the Impetus LeapLogic™ Suite, siloed legacy systems (like data warehouses, ETL pipelines, mainframes) are automatically migrated and transformed into modern cloud or hybrid data platforms. This process bridges the **data gap** and creates a unified, clean, accessible data layer for AI, while preserving business logic. During migration, LeapLogic enriches metadata, connects lineage, and validates all transformed workloads in the new environment, ensuring accuracy and performance. One global bank leveraged this to consolidate customer data from 27 disparate sources into a single cloud platform, providing the needed 360° view for its personalization AI agent.

BUILDING A SEMANTIC INTELLIGENCE LAYER:

Impetus Context Fabric™ creates an enterprise ontology and knowledge graph layer that encodes business meaning. It connects data with business definitions, policies, and rules, transforming raw data into structured, AI-consumable context. For example, Context Fabric can define how a 'high-value customer' is determined (combining balance, tenure, products owned, etc.) so that all AI agents consistently recognize which customers fit that category. This semantic layer, combined with retrieval techniques, supplies agents with on-demand contextual knowledge – bridging the **semantic gap**. Agents can now answer customer questions or make decisions using the company's specific terminology and logic, not just generic training data.

AUTONOMOUS WORKFLOW ORCHESTRATION WITH GUARDRAILS:

With clean data and context in place, Impetus helps design intelligent workflows using domain-specific Skills, Kits, and Agents with Agent Solutions™. These are pre-built components and templates for common financial scenarios (fraud investigation, claims processing, customer onboarding, etc.) that can be customized. Multiple agents can be orchestrated to work together, and a human-in-the-loop control layer is integrated for regulated industries. The orchestration ensures agents don't lose context between steps and that there are checkpoints for any sensitive decisions – effectively bridging the **execution gap**. Impetus' methodology emphasizes testing these workflows thoroughly in sandbox environments to iron out hand-offs and communication between agents before going live.

UNIFIED OBSERVABILITY AND GOVERNANCE:

Impetus Prism™ provides a control plane that monitors data pipelines, model performance, and agent decisions in one unified view. It offers real-time observability – for instance, tracking drift in an agent's behavior or flagging if an agent starts operating outside its bounds. Prism enforces compliance rules (like preventing PII data from being used by an AI agent if not permitted), proactively detect anomalies, and automates remediation. This helps close the **trust gap** by preventing hallucinations or unauthorized actions and by ensuring cost-efficient operations (with deep insights into tokenomics and other costs). In one case, a large insurer used Prism to catch and correct an AI claims agent that began missing a step in the adjudication process after a software update – the issue was detected and resolved within hours, avoiding regulatory non-compliance.

PRE-BUILT, CONTEXT-AWARE SOLUTIONS FOR FINANCIAL SERVICES:

Recognizing that time-to-market is critical, Impetus offers purpose-built agentic solutions for high-value use cases in banking and insurance. These include an agent for real-time payment fraud detection, an agent for continuous customer engagement and cross-sell, and others for compliance monitoring. Each solution comes 'context engineered' with financial domain knowledge and can be further tuned to a client's specific data and policies. This jump-starts adoption by providing proven blueprints, which are then tailored to fit the organization's environment.

In essence, the Impetus Leap™ AI family makes it easier and safer for financial institutions to deploy context-aware agentic AI. By systematically engineering the surrounding context – data, semantics, workflows, and governance – Impetus empowers banks, payment companies, and insurers to harness the full potential of autonomous AI agents while ensuring they are well-grounded within the organization's unique context landscape.

“ Differentiated agentic AI in FSI isn't just automating a workflow — it's bringing the right context to bear at each decision point, **which is what actually earns trust in a regulated environment.** ”

07 HOW IMPETUS DELIVER AI PROGRAMS

Models don't win with agentic AI — method does.


The Impetus Context Engineering Delivery Lifecycle (CEDL) brings that method to life: a four-phase loop that engineers context, certifies every agent before launch, and keeps improving it in production. Repeatable by design, provable at every step.

C → E → D → L, and **L** feeds back into **C**. The continuous loop never closes: what an agent learns in production makes the next version smarter.

WHY IT MATTERS

CEDL is what makes Impetus Leap AI™ production-ready rather than demo-ready. It provides the accountability structure most pilots lack — the reason so many stall at hand-off. With CEDL, no agent reaches production without a context bundle, an evaluation record, and an explainability trace.

C	E	D	L
Contextualize	Engineer	Deliver	Learn
Discover + Model	Build + Validate	Activate + Certify + Handoff	Benchmark + Observe
Map the enterprise knowledge landscape and identify what context the agent needs to reason correctly.	Build the context layer: pipelines, embeddings, tools, and evaluation harnesses—before any agent is deployed.	Production deployment with a signed-off audit trail. No agent ships without passing CEDL certification.	Live performance feeds back into context quality. What the agent learns in production makes the next version smarter.



SUCCESS STORIES

PAYMENT FRAUD DETECTION AT SCALE

CHALLENGE

A Fortune 100 bank faced increasingly sophisticated payment fraud, including account takeovers and mule networks that evaded its old rule-based systems.

SOLUTION

Impetus implemented a unified, real-time fraud intelligence platform powered by agentic AI. It integrated thousands of data points (transactions, device telemetry, geolocation, merchant info, customer behavior) into an entity graph to recognize linked activities across accounts. On top of this graph, multiple AI models – from anomaly detection to graph analytics – were deployed, with an agent orchestrating them and handling alerts. Critically, every alert came with an explanation generated by the system (e.g., highlighting the chain of events linking a device to multiple suspicious accounts).

The platform intercepts fraudulent payments before they execute and automates case creation – investigators intervene only on the most complex incidents. All decisions are fully auditable, helping the bank meet stringent regulatory requirements and internal model risk management standards.

<100ms
response time

↓ LOSSES
substantial drop in
fraud losses

DOUBLE-DIGIT
false-positive improvement

REAL-TIME PERSONALIZED CUSTOMER ENGAGEMENT

CHALLENGE

A large credit card company was struggling with low uptake on its marketing campaigns and offers. Analysis showed that their batch, segment-based approach was too slow and generic – by the time an offer reached a customer, the moment of relevance had often passed.

SOLUTION

Impetus helped implement a real-time customer intelligence layer combined with agentic AI for decisioning. Streaming data from card transactions, mobile app interactions, location pings, and web activity now feeds into a unified profile for each customer. An AI scoring engine running in milliseconds evaluates a customer's context in the moment – for example, detecting if a customer is traveling, if they just got paid, or if their spending pattern has changed – and an AI agent uses those insights to select and deliver the optimal message or offer across channels (push notification, email, etc.).

All this happens within 250 milliseconds of an event (like a purchase or login). The system is also used to power precise fraud warnings and credit risk decisions, since the same real-time profile can indicate potential fraud or credit stress. Thanks to an explainability module, every offer or decision is accompanied by a rationale (stored internally) like 'Offer A was shown because customer is within 5 miles of Partner Store and has high likelihood to redeem based on past behavior.'

+300%
increase in offer acceptance

3.5M
personalized decisions
per day

250ms
from event to decision

These case studies demonstrate tangible outcomes of agentic AI when paired with strong context engineering. Importantly, metrics like faster fraud response, higher offer uptake, and improved accuracy translate to direct business value and ROI. The journey from pilot to production is not trivial, but with a methodical approach to building context and the right expertise, financial institutions can rapidly advance on this frontier.

CONCLUSION

Agentic AI offers a transformational opportunity for financial services and insurance firms, enabling them to tackle long-standing challenges and seize new opportunities in personalization and fraud mitigation. By empowering AI systems to autonomously execute multi-step processes – grounded in the institution's data, rules, and oversight – organizations can achieve breakthroughs in efficiency and customer experience while maintaining control and compliance.

Hyper-personalization agents can drive growth by delighting customers with tailored advice and proactive service, and fraud/risk agents can act as ever-vigilant guardians, protecting both the institution and its clients in real time. Crucially, our analysis shows that success with agentic AI hinges on context-awareness. Banks and insurers must invest in closing the **context gap** through robust data foundations, domain-specific knowledge integration, intelligent orchestration, and strong governance. Those that do so will be able to move from AI that works in demos to AI that delivers in production – improving customer satisfaction, reducing losses, and ensuring regulatory confidence.

“ As the financial industry continues to evolve, the ability to deploy **trusted, context-aware AI agents** will become a key differentiator.

As the financial industry continues to evolve, the ability to deploy trusted, context-aware AI agents will become a key differentiator. Early adopters are already seeing competitive advantages, from cost savings and faster processes to higher customer loyalty. We stand at a juncture where technology makes it possible to redefine personalization and fraud detection simultaneously – achieving both greater customer-centricity and stronger security. By embracing agentic AI with the right context engineering, financial institutions can chart a path of sustainable innovation, meeting the demands of a digital-first market and building deeper trust with those they serve.

REFERENCES

- i. <https://verafin.com/nasdaq-verafin-global-financial-crime-report>
- ii. <https://www.mckinsey.com/industries/financial-services/how-we-help-clients/banking-on-innovation-how-ing-uses-generative-ai-to-put-people-first#the-opportunity>
- iii. <https://www.mckinsey.com/industries/financial-services/our-insights/the-ai-powered-bank-rewiring-for-excellence-in-customer-care>
- iv. <https://venturebeat.com/ai/wells-fargos-ai-assistant-just-crossed-245-million-interactions-with-zero-humans-in-the-loop-and-zero-pii-to-the-llm>
- v. <https://www.lightico.com/>
- vi. <https://digital-strategy.ec.europa.eu/en/library/draft-commission-guidelines-classification-high-risk-ai-systems>
- vii. <https://isrgpublishers.com/wp-content/uploads/2025/09/ISRGJEBM4262024I.pdf>

IMPETUS

About Impetus Technologies

Impetus Technologies is a global software-led Data and AI services company that helps enterprises make AI actually work. Built on a rich engineering heritage, deep industry expertise, and a robust technology partner ecosystem, Impetus enables organizations to modernize data foundations, engineer enterprise context, orchestrate trusted agentic AI systems, and govern them at scale to deliver measurable business outcomes. Impetus partners with leading organizations across financial services, consumer packaged goods, healthcare, life sciences, manufacturing, travel and hospitality, and other industries, worldwide.

To learn more about how context-aware agentic AI can transform your organization and the Impetus Leap AI Family of Solutions and Services visit www.impetus.ai