



WHITEPAPER

# Outlook 2026: The Agentic Banking and Financial Services Enterprise



# Introduction:

## The Current State of Play

Banking and financial services providers will see additional pressures in 2026, with economic, regulatory, and technological challenges taking top priority.

From an economic perspective, despite years of automation investment, cost-to-income ratios are deteriorating, while the normalization of credit cycles is increasing risks for lenders.

At the same time, regulatory and compliance issues have ramped up. Not only are regulators intensifying scrutiny of AI fairness, auditability, and transparency, but compliance costs are rising as BFS organizations see increasingly sophisticated fraud and financial crime threats.

Then there's the fact customers expect faster, more personalized interactions, which is made increasingly difficult by a shrinking talent pool with high attrition rates, skill gaps, and higher training costs, as well as complex, fragmented workflows that slow resolution and compound customer frustrations.

Agentic combines autonomous software agents, continuous learning from data, and explicit human governance to be faster and more accountable in how decisions are made.

AI is expected to address these hurdles, with 70% of banking and financial services organizations looking to automation to solve them.<sup>1</sup>

An IDC study reveals that “Frontier Firms” – organizations that embed AI agents across workflows – report returns on AI investments roughly three times higher than slower adopters.<sup>2</sup>

Furthermore, KPMG research shows agentic AI will generate \$3 trillion in corporate productivity gains, with a 5.4% EBITDA improvement for the average company annually, and organizations achieving a 2.3x return on agentic AI investments within just 13 months.<sup>3</sup>

Yet, while institutions are no strangers to AI and automation, there needs to be a shift in how they are used. Not least because, while heavy investments have been made in analytics, machine learning, and rules-based automation, organizations are experiencing uneven outputs largely confined to insight generation.

The problem is that many AI tools don't act, they only update – and often after the window for opportunity has passed. This allows banks to predict fraud risks, for example, but not to act autonomously or work in tandem with humans doing more sensitive tasks.

Agentic AI, however, takes a different approach. It can bridge this gap by moving financial services operations away from linear, step-by-step workflows and toward systems that can sense, decide, and act in real time.

<sup>1</sup> <https://www.sutherlandglobal.com/insights/blog/transforming-bfs-operations>

<sup>2</sup> <https://marketingassets.microsoft.com/gdc/gdcvSFjHm/original>

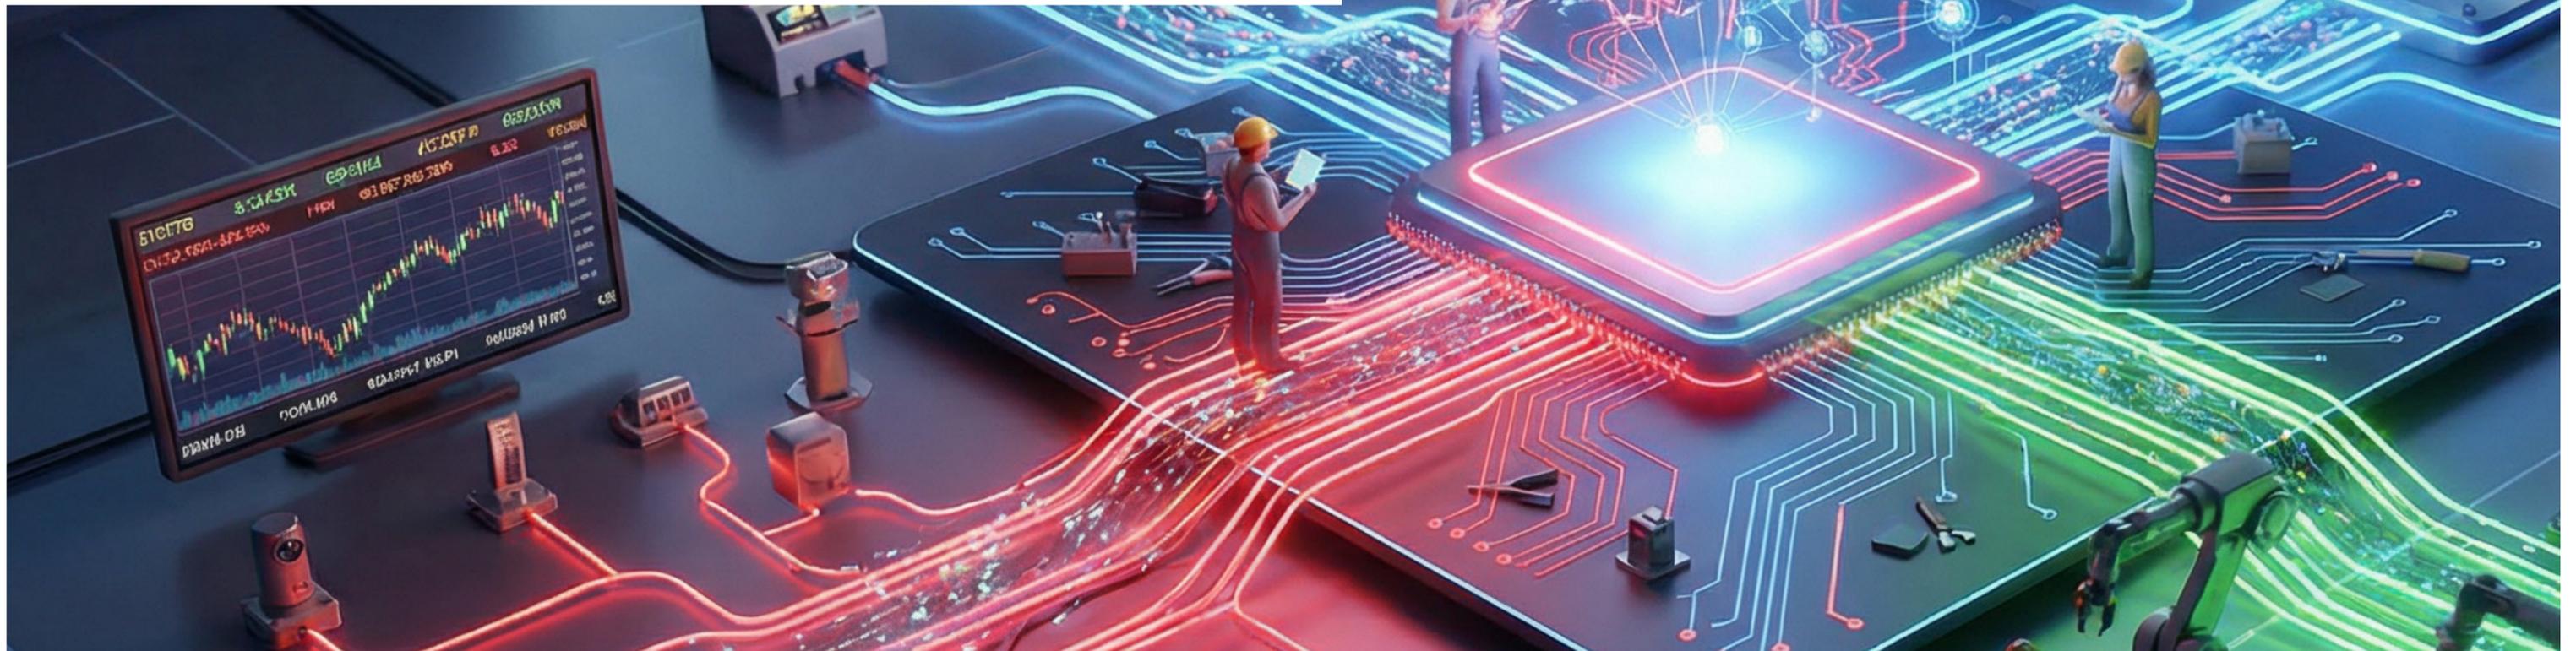
<sup>3</sup> <https://kpmg.com/kpmg-us/content/dam/kpmg/pdf/2025/kpmg-agentic-ai-advantage.pdf>



Instead of AI tools that only generate alerts or reports, agentic architectures embed decision-making into the flow of work. These agents coordinate across risk, operations, and customer teams to trigger next-best action. They also document rationale and escalate edge cases to human experts when needed.

Done well through proven strategies with reliable technology, agentic AI can grow organizations while preserving the legacy values that built them.

The following sections explore three areas where agentic AI can create measurable impact for banking and financial services organizations, and practical steps leaders can take to prepare.



# 3 Opportunities for Banking and Financial Services Enterprises

## Opportunity 1: Intelligent Customer Onboarding and KYC Orchestration

### The Challenge

Even with enhanced tool integrations and real-time data checks, customer onboarding and Know Your Customer (KYC) processes remain labor-intensive – and this is just two steps in a complex chain of interconnected banking workflows. Traditionally, human agents have needed to manually review and interpret customer documentation before navigating their way through fragmented systems to verify data and meet internal risk requirements.

If the agent needs additional customer documentation, an already long process becomes even more drawn out, creating bottlenecks at various points. Since the same customer case may be handled by multiple people, there's an added risk of inconsistent handling, missteps, and variable outcomes.

While human agents may carefully document their every action, limited lineage and explainability remain. Manual decision logs make it difficult to justify KYC outcomes and may not meet regulatory challenges.

In addition, KYC is currently event based rather than risk based. Dynamic risk changes on a near real-time basis are not captured and risk is only assessed again at the next event.



4 <https://legal.thomsonreuters.com/en/insights/articles/trends-in-synthetic-identity-fraud>



## The Agentic AI Opportunity

With autonomous AI agents in the mix, KYC orchestration goes from being a document chase to an intelligent, self-improving process. These agents extract, interpret, and validate information from diverse document types (passports, government IDs, bank statements, etc.), requiring no manual data entry or initial human analysis. Inconsistencies and anomalies are immediately flagged for further action, accelerating the process of requesting additional information.

Risks are assessed contextually, rather than using static thresholds that may be arbitrarily applied by each human agent. These AI agents dynamically incorporate data from real-time sanctions lists, database entries, and internal customer history to act on concerning information immediately.

Thomson Reuters research shows 95% of synthetic identities are not detected during the onboarding process at financial institutions.<sup>4</sup> Agentic AI orchestration addresses this by conducting continuous behavioral analysis and pattern recognition that traditional one-time KYC checks cannot achieve, creating a moving target for fraudsters rather than a single compliance checkpoint.

An AI-driven onboarding and KYC orchestration solution for a digital bank reduced new account approval cycle times by 45% and improved first-time-right verification rates, while strengthening auditability of onboarding decisions.

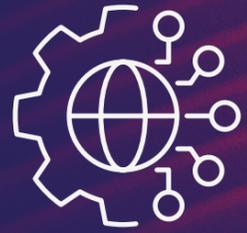


At the same time, exceptions are handled intelligently. The AI agent knows the difference between routine clarification (missing middle initial) and a genuine red flag (mismatched beneficial owner). Each escalated case comes with full documentation and explainability.

This unified data profile supports ongoing remediation and monitoring. Once a customer is successfully onboarded, their file stays updated as new information becomes available. Remediation actions, such as relationship reviews or exit decisions, are made in real time rather than solely at scheduled quarterly or annual reviews.



## Practical Actions for Banking Leaders



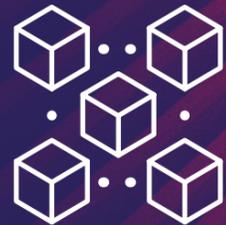
**Consolidate and govern customer data:** Integrate identity, verification, risk, and all third-party data into a **unified, auditable data fabric**. Ensure the data lineage is traceable for each and every KYC decision. This is critical for regulatory compliance and audit defense.



**Digitize and standardize onboarding rules:** Formalize all **KYC rules**, including risk thresholds and escalation criteria. Put into a machine-readable format, such as rule sets, decision models, and policy graphs that can be applied consistently and automatically by AI agents.



**Identify the best use cases for agentic onboarding and triage:** Carry out a structured **digital assessment** to identify and start with high-volume, lower-complexity onboarding tasks, such as retail customers and new deposit accounts. Prioritize those that can greatly accelerate case intake by reducing the most manual data entry hours.



**Create explainable risk scoring models:** Since black-box models are a regulatory liability, use knowledge bases to score risk in ways that agents and regulators can understand. Seek to explain the “why” behind a customer escalation or approval. Enlist the help of Sutherland’s knowledgeable AI frameworks and **digital assurance** offerings to build models faster and with proven results.



**Build human-in-the-loop controls:** Clearly distinguish between cases where automation is appropriate and those that require **human review**. Include system override recommendations if warranted, and provide feedback mechanisms to improve the agentic reasoning as it works continuously.



## Opportunity 2: Autonomous Fraud Detection and Investigation Orchestration

### The Challenge

Each year, financial crimes are becoming more sophisticated, with fraud rings moving faster via synthetic identities and AI scams. The fallout has cost banks more money in restitution to customers, as well as compliance costs and regulatory fines.

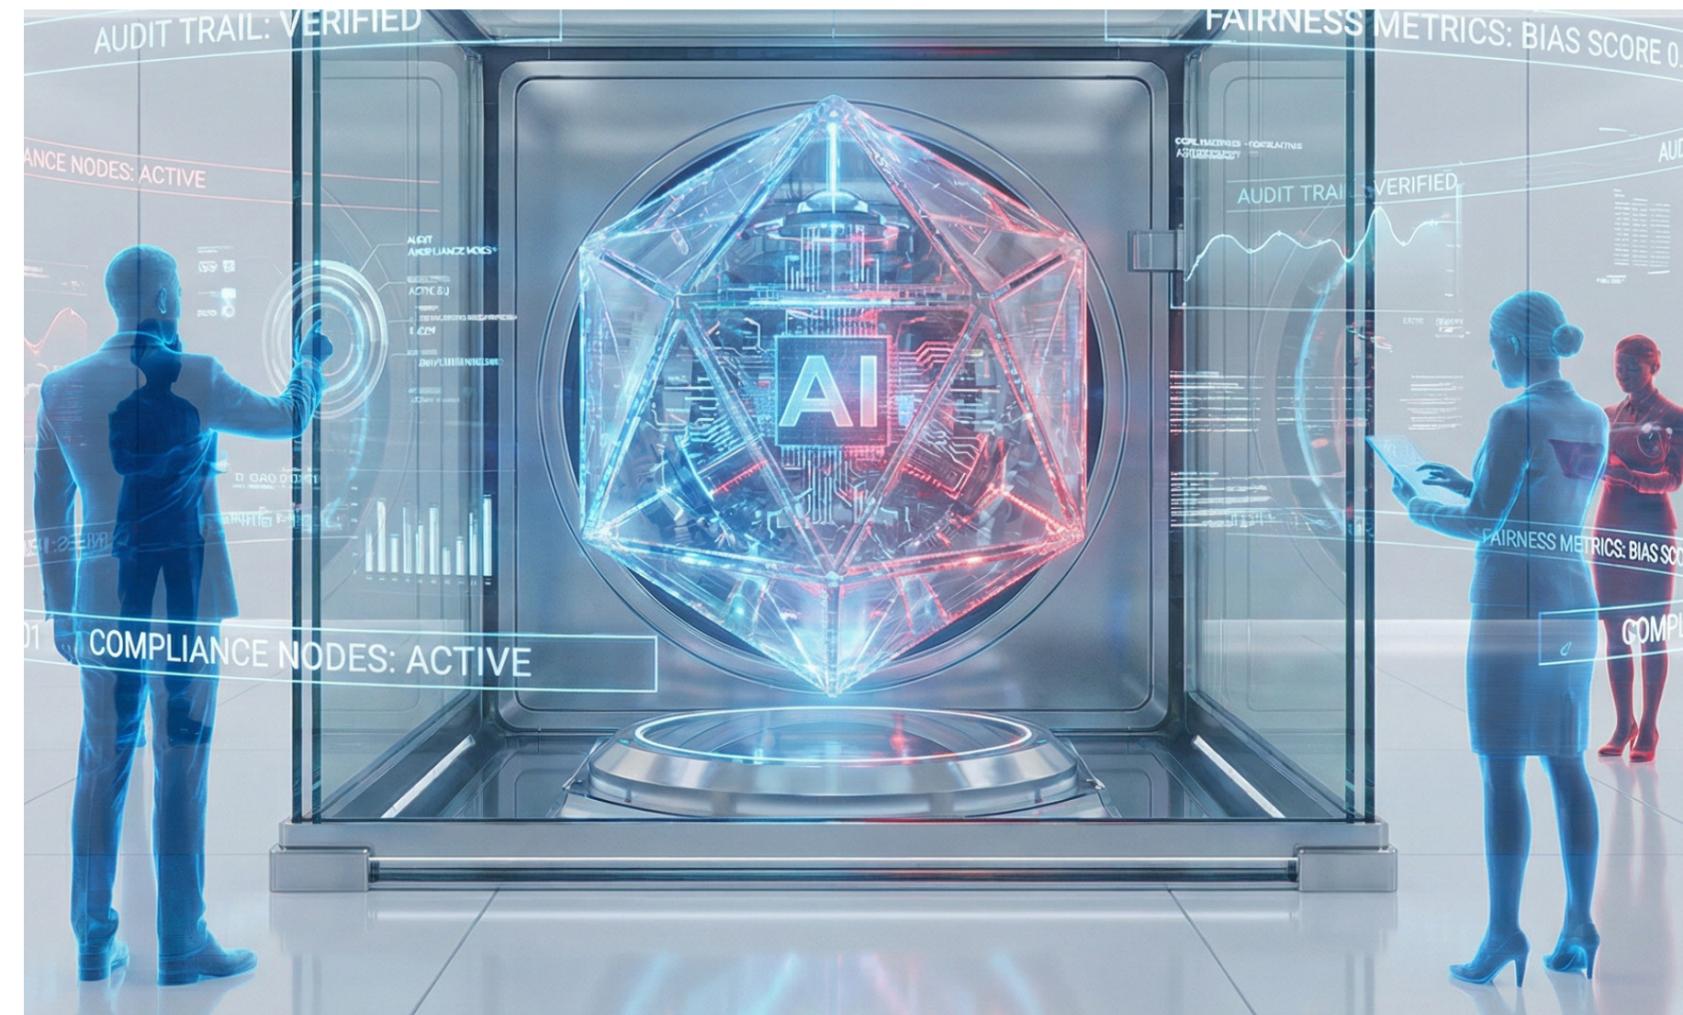
Now, more than 50% of fraud involves artificial intelligence, with 44% of financial professionals reporting deepfakes are used in fraudulent schemes, 60% citing voice cloning, and 59% identifying AI-powered SMS and phishing scams.<sup>5</sup> The global fraud management market was valued at \$11.6 billion in 2025 but is forecast to surpass \$80 billion by 2035, a 20% CAGR reflecting the escalating threat.

While banks can often detect these thousands of suspicious transactions each day through analytics and rules engines, it takes manual effort to investigate and act on them. It can take days or even weeks to fully resolve a legitimate fraud claim, requiring human agents to compare data across several platforms.

Legitimate fraud claims take their toll on enterprises, but the noise of false positives and their alerts adds to it. These unwarranted cases clog up already-stressed human agent workflows and limit customers' accounts until they can be resolved.

Meanwhile, customers often grow dissatisfied waiting to use funds or essential services to work, live, and play in an increasingly stressful economic climate. When a decision is finally made, it may be inconsistent, depending on who handled it and how long it sat in limbo.

The challenge, therefore, is twofold: reducing preventable fraud losses at the point of detection and dramatically accelerating post-detection investigation and resolution.



<sup>5</sup> <https://www.feedzai.com/resource/state-of-ai/>



## The Agentic AI Opportunity

Agentic AI in fraud operations is not about replacing core fraud detection engines. It is about orchestrating detection, triage, investigation, and resolution in a unified, accountable workflow.

Instead of bouncing from system to system to manually assess fraud risk, automated AI agents look at fraud signals in the context of a customer's history and spending patterns.

The technology also uses merchant intelligence and advanced behavioral analytics to instantly distinguish between edge cases (such as a customer traveling on business) and true high-risk scenarios. It pulls in data from many systems, including geolocation data, to analyze before sending to a human investigator. When vetted and likely cases do arrive on an investigator's desk, they are sound and with ample data to support them.

AI agents also categorize on likely fraud outcomes (confirmed, suspected, false positive) and recommend the right action for the human agent to take – if it doesn't take the action itself. Human agents can always override to meet regulatory and best customer service practices. Each human action also helps the AI agent learn and approach new cases with added context and applicability.

As new attack vectors or fraud tactics emerge, AI agents learn to identify and surface them. They adapt detection rules automatically and dynamically to grow their knowledge base.

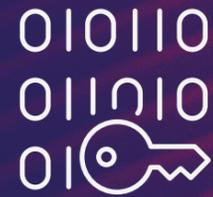
Institutions can then stay ahead of fraud evolution while leaving a proper documentation trail of each decision to satisfy regulators.



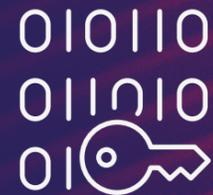
Using agentic AI, companies can see up to a 25% reduction in fraud losses and a 70% drop in false positives.<sup>6</sup>



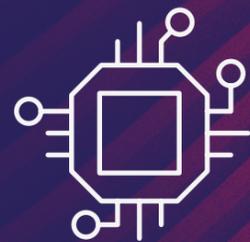
## Practical Actions for Banking Leaders



**Unify fraud intelligence sources:** Consolidate alerts from payment networks, internal rule engines, third-party fraud platforms, and customer behavioral tools into a single intake. Ensure network data, device intelligence, and merchant risk scores are based on **high-quality data**.



**Prioritize investigation orchestration.** Focus first on high-volume alert categories and customer-reported **fraud claims** where cycle time reduction and false-positive management will deliver measurable ROI.

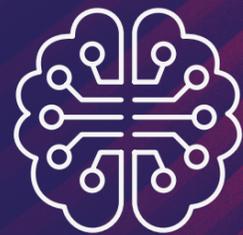


**Map the current FinCrime tech stack.** Compare your existing operating model against the capabilities of industry leaders, and flag overlaps, gaps, and systems you can retire. Use the map for a target ecosystem blueprint with clear integration points, APIs, and ownership that can plug into Sutherland's **unified solution** without adding a separate tool.





**Enable multi-agent coordination:** Build agent ecosystems where fraud agents coordinate with customer agents for timely customer notification, compliance agents for regulatory reporting, and risk agents for pattern analysis. Nurture a culture of bringing in the best talent for the risk and situation at hand.



**Harness continuously updated AI risk intelligence:** Inventory current sanctions, adverse media, and transaction-monitoring data sources with a comprehensive **digital assessment** to identify gaps in coverage or update frequency. Create a single internal view of these risk signals and define owners, refresh cycles, and quality checks so your teams can trust the data that feeds fraud and automated decisions.

## Opportunity 3: Seamless Multi-Channel Customer Servicing and CX Orchestration

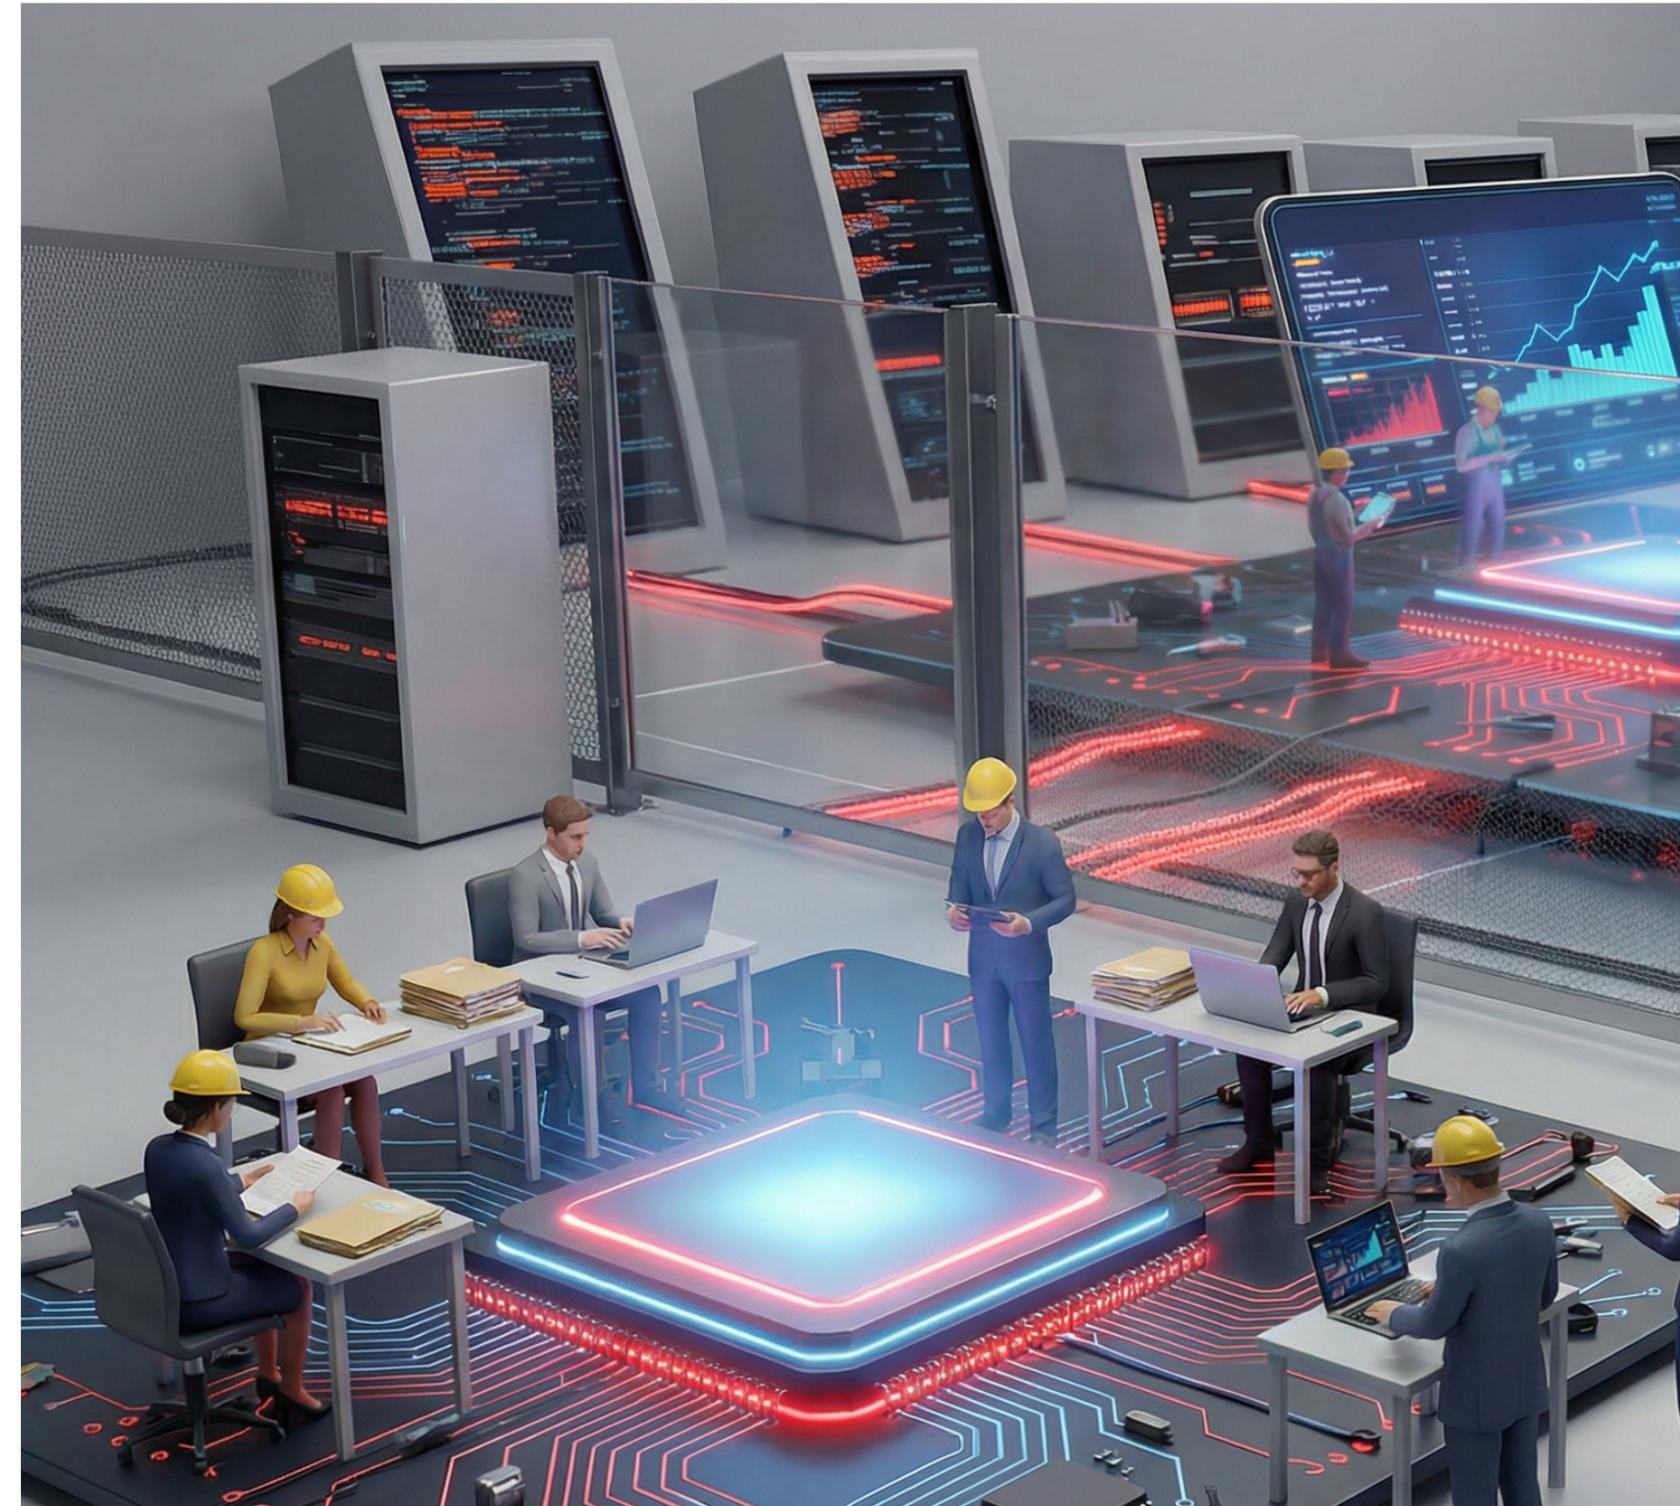
### The Challenge

Chatbots and 24/7 banking features have supported banking customers' expectations of instant, seamless servicing, but these features lack the ability to be truly cross-channel and offer fragmented experiences. Agents service routine requests through disconnected platforms that don't communicate with each other or provide context for an optimal customer experience outcome.

Simple requests, like "add a beneficiary to my account," require agents to navigate multiple systems, confirm policy rules, and document each step. If the process is disrupted, customers often have to start from the beginning with a different agent. Different agents may apply policies in different ways, depending on experience, knowledge level, and channel comfort.

This disjointed customer workflow is demanding on the customer and also places strain on human agents, who may feel existing workplace pressures from high turnover and continual context switching. High average handling times may lag further when human agents don't feel successful in customer interactions or aren't confident in the decisions made.

As new agents are onboarded to customer experience teams, the knowledge required can span multiple complex technologies, tools, and systems, with only static documentation to support them. New best practices may not be documented in time to transform how new human agents work, leaving them with less satisfaction from the start of their employment.



## The Agentic AI Opportunity

AI agents work, learn, and adapt within a unified layer that human agents can interact with in more intuitive, rewarding ways than traditional systems offer. The intelligent interface understands customer intent, based on case history, nuanced language models, and relevant data from all systems, and guides the next steps for either automated resolution or further escalation to the most qualified agent for the situation.

Agentic systems also autonomously retrieve and apply guidance from process documents, procedure manuals, and policy repositories in real time. Instead of manually searching knowledge bases, agents receive context-aware instructions embedded directly into the workflow, ensuring decisions align with the latest approved processes.

This doesn't simply move the case along, however; it also surfaces relevant data to support the human agent, giving them recommended actions and resources for the best possible outcome.

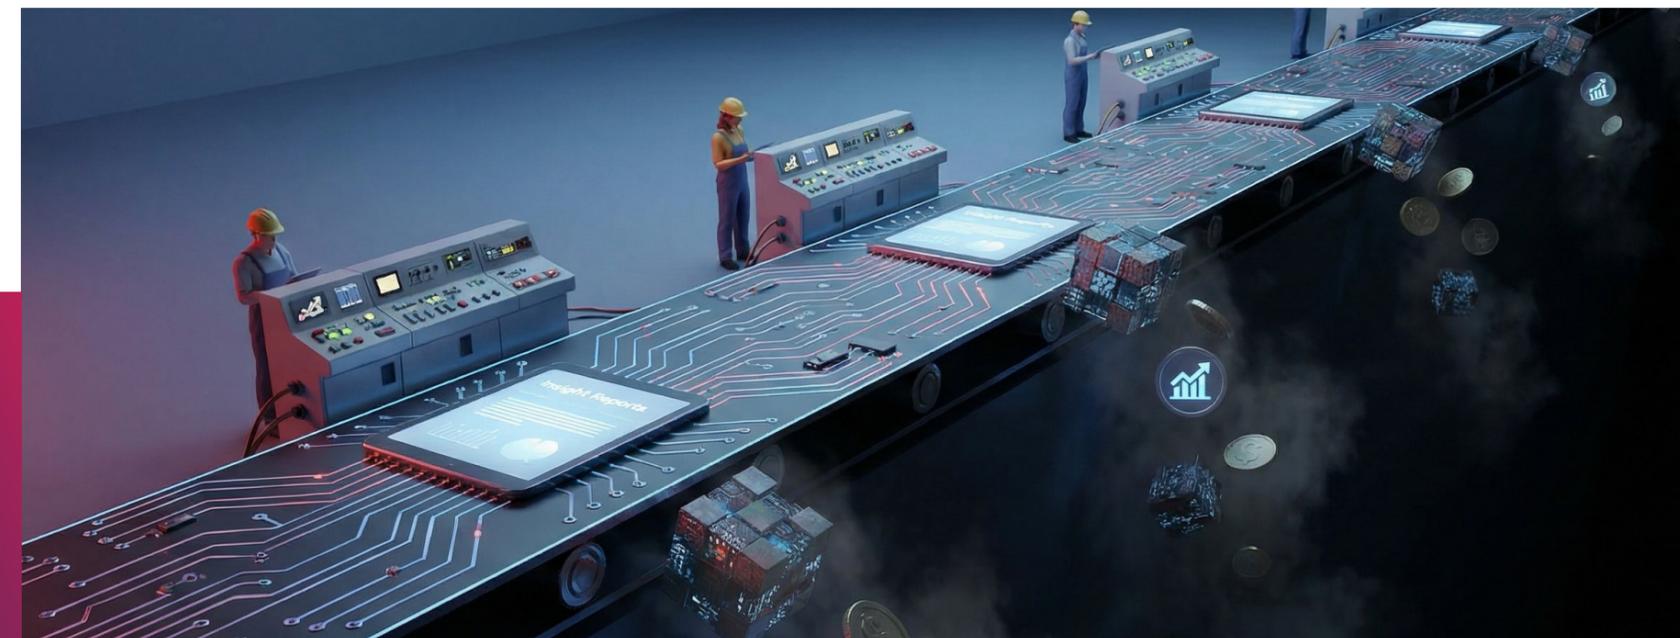
For example, a customer calling about a mortgage refinancing can speak or type their question, then be routed automatically to the right expert – and skip

several generic interactions with agents not qualified or authorized to help. Dispute inquiries go to the dispute team without the customer having to repeatedly request “speaking to the manager.”

Agentic AI can also resolve transactions and account actions that don't need human interaction. Examples include address updates, card reissues, or even credit disputes (depending on product line, regulations, and internal policies). With agentic AI, these cases never make it into the human agent pipeline, reducing overall volume and breaking peak-season bottlenecks before they ever form.

As new human agents onboard more quickly, they may see a more positive overall work experience – reducing attrition and building more legacy teams that know and understand the company culture. Simultaneously, customers receive faster, more personalized resolutions from AI and human agents, regardless of which touchpoint (phone, email, online account access) they use.

For a top 10 US card issuer, deploying 120+ AI-powered bots equipped them to handle credit disputes with 99.5% accuracy and cut manual intervention by 60%.<sup>7</sup>



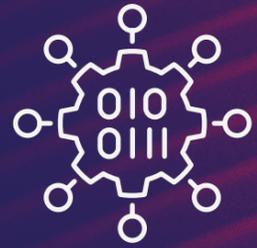
## Practical Actions for Banking Leaders



**Map and prioritize end-to-end customer journeys.** Define the highest-impact customer journeys across channels, such as dispute resolution, loan servicing, and more. Identify friction points and bottlenecks, then use **journey mapping** to pinpoint where autonomous resolution, knowledge retrieval, and intelligent routing will deliver measurable impact.



**Centralize and structure enterprise knowledge to support those journeys.** Consolidate product information, policies, pricing, compliance guidance, process documents, and decision rules into a single, structured, version-controlled, and machine-readable **knowledge repository**. Align knowledge architecture to prioritized journeys and embed rule-based workflows directly into AI-assisted servicing.



**Revisit and establish refreshed CX metrics:** Track and measure AHT, first-contact resolution (FCR), digital containment rate, and Net Promoter Score (NPS) before deploying agentic systems. Measure and refine as they are deployed. Use metrics to identify gaps and guide continuous improvement, utilizing established best practices as guidance.



**Conduct an AI readiness and change impact assessment:** Bring together CX, operations, risk, and IT leaders with a digital assessment to identify where AI-assisted decisions will change scripts, approval paths, and handoffs. Document which roles, skills, and incentives need to evolve. Use those findings to define a phased change plan, so frontline teams are prepared to adopt and trust the new workflows.

# The Road to the Agentic Banking Enterprise

The banking and financial services sector may look very different by Q4 of 2026, and many organizations haven't waited to take their cue on agentic AI. In fact, some fintechs and hyperscalers are already operating agentic architectures, unburdened by legacy systems or organizational inertia.

While typically not as nimble, traditional banks and large financial institutions can choose to follow suit or cede any market advantage to their faster-moving competitors. The competitive gap is widening. More than 50 of the world's largest banks announced over 160 agentic AI use cases in 2025 alone, with early implementations showing 30-50% reductions in manual workload and cycle times reduced by up to 80%. McKinsey warns that agentic AI could put \$170 billion in global banking profits at risk for institutions that fail to adapt their business models.<sup>8</sup>

Three distinct adoption tracks are emerging: incumbent banks focusing on back-office efficiency (68% reporting substantial returns), digital-native fintechs embedding agents across customer-facing experiences, and hyperscalers leveraging agents as foundational operating architecture.



<sup>8</sup> <https://www.mckinsey.com/industries/financial-services/our-insights/global-banking-annual-review>



In every era of banking and financial services transformation, the organizations that invested in rethinking risk and customer experience gained a lasting advantage. It is no different now. The distinction is that innovation isn't a new product, channel, or pricing strategy. It's a new mode of operating, where intelligence becomes operational and adaptive across every process.

Banking enterprises that thrive in 2026 and beyond will be the ones driving agentic change, not reacting to it. They will move beyond pilots and proofs-of-concept to operationalize agentic intelligence at scale. Instead of a plug-and-play suite of disparate tools built around a legacy system, they will redesign workflows around autonomous agents with an anticipatory approach to future regulation and AI threats.

This shift from "AI as insight" to "AI as action" embeds autonomous decision-making into every critical process, bringing in humans as needed for accountability and warmth. Early adopters have seen real change, not just planned for theoretical wins. They represent real operational leverage, measurable cost savings, and competitive positioning.

The competitive window is open for BFS organizations willing to lead the transformation.

**Agentic capability relies on the right foundational elements. Explore Outlook 2026: The Road to the Agentic Enterprise for the fundamentals every organization needs.**



**Disruption is inevitable. Make it intentional.**

Artificial Intelligence. Automation. Cloud Engineering. Advanced Analytics. For Enterprises, these are key factors of success. For us, they're our core expertise.

We work with global iconic brands. We bring them a unique value proposition through market-leading technologies and business process excellence. At the heart of it all is Digital Engineering Services – the foundation that powers rapid innovation and scalable business transformation.

We've created 363 unique and independent inventions, 250 of which are AI-based and rolled up under several patent grants in critical technologies. Leveraging our advanced products and platforms, we drive digital transformation at scale, optimize critical business operations, reinvent experiences, and pioneer new solutions, all provided through a seamless "as-a-service" model.

For each company, we provide new keys for their businesses, the people they work with, and the customers they serve. With proven strategies and agile execution, we don't just enable change – we engineer digital outcomes.

