



Unlocking Tomorrow:

A playbook that explores the power of agentic AI for the Financial Services industry



Table of contents

Executive Summary	3
01 Introduction	6
02 Sector spotlights	13
• Banking	14
• Insurance	18
• Capital markets (including private equity)	25
03 Becoming a Frontier organisation	28
04 Why PwC and Microsoft	39
Authors	43

Executive Summary



Financial services organisations are operating in a period of profound transformation. Macroeconomic uncertainty, rising compliance demands, increasing fraud risk, shifting customer expectations and a rapidly ageing workforce are all reshaping how banks, insurers and capital markets firms must operate. At the same time, regulatory frameworks such as EU's AI Act, DORA and GDPR are accelerating the need for transparent, secure and well-governed technology adoption. Against this backdrop, many financial institutions are turning to agentic AI to strengthen resilience, accelerate key processes, transform the customer experience, and unlock new opportunities for growth.

Agentic AI represents a step change from traditional analytics, automation and generative AI. Rather than completing isolated tasks, agents can reason, act and orchestrate multistep processes with minimal human intervention—accessing data, executing workflows and adapting to new conditions. When deployed responsibly, this enables financial institutions to improve efficiency and accuracy while maintaining essential human oversight for high-stake decisions.

The impact of agentic AI is already visible across the industry. In banking, significant value is already emerging in financial crime, regulatory compliance, customer engagement and lending. Agents can gather and validate Know Your Customer information, assemble source-of-wealth documentation and streamline onboarding. They enhance transaction monitoring and fraud detection processes while maintaining audit trails for regulatory reporting. In customer experience, agents support contact centre staff, power self-service journeys and deliver real-time personalised offers. Lending processes are being reshaped end to end, with document verification, risk assessment and workflow coordination increasingly automated.

In insurance, agentic AI is transforming underwriting, claims, customer service and workforce sustainability. Underwriters can base decisions on continuously updated risk signals drawn from policy data, claims activity, climate models, behavioural insights and geospatial imagery. Claims can be triaged and resolved nearly autonomously from first notice of loss through policy validation and recommended settlement, improving both speed and claimant experience. Fraud detection benefits from multi-agent collaboration across structured and unstructured data. For insurers facing significant demographic challenges, with many skilled workers nearing retirement and limited availability of new talent, agentic AI provides a critical mechanism for preserving institutional knowledge and maintaining operational capacity.

Capital markets organisations are leveraging agentic AI for high value analytical and operational activities. Multi-agent workflows compress time to first analysis from days to minutes by ingesting financial statements, earnings calls, market news and alternative data into structured research drafts. Agents can also support rapid interpretation of external shocks such as tariffs or geopolitical events and expand stress testing scenarios to model changing market conditions. In private equity and investment operations, agents accelerate due diligence, surface operational risks across portfolios and enhance client advisory with real-time, personalised insights.

Across all subsectors, a common pattern is emerging: agentic AI unlocks significant value when it transforms decision-heavy, high-complexity workflows, transforms customer experience, and connects previously disconnected parts of the enterprise. Yet scaling these capabilities requires more than isolated pilots. Financial institutions must build reusable, enterprise-wide AI foundations that connect strategy, governance, technology, talent and delivery. A Frontier organisation—one that combines human expertise with orchestration at scale—can operate with greater speed, precision and resilience in an increasingly demanding market environment.

Achieving this transformation requires a clear, actionable roadmap. Leaders should define a bold but responsible vision for AI; raise organisational literacy and confidence; redesign operating models around value chains, not isolated tasks; modernise data and cloud foundations; embed Responsible AI into every stage of deployment and extend these principles into real time monitoring and control; and cultivate partnerships that accelerate adoption safely. By combining these elements, financial services firms can move beyond experimentation and capture the full enterprise value of agentic AI.

Agentic AI is no longer a distant promise. It is a practical catalyst for transforming financial services today. Institutions that act now—scaling responsibly, governing effectively and empowering their people—will be the ones that set the competitive pace, strengthen customer trust and shape a more resilient, intelligent and inclusive future for the industry.

01



Introduction

Market situation and industry challenges

Financial services companies are facing a rapidly accelerating pace of change. Even though the industry has shown great **resilience in the face of recent shocks**, it is contending with many challenges that could increase risks in the near future. Among the challenges that are converging in a highly complex market environment: ongoing digital transformation, an ageing workforce, changing expectations and reduced loyalty from younger customers, shifting economic and geopolitical realities and emerging technologies such as AI. These winds of change are creating both new opportunities and new threats for all of the industry's subsectors, including banking, wealth and asset management, insurance, and capital markets.

Meanwhile, new regulatory requirements have come into force across different jurisdictions with the aim of safeguarding the stability and security of the global financial system. But these are also increasing compliance pressures and costs for banks, insurers and other institutions. PwC's **2025 Global Compliance Study** found that 90% of respondents in the financial services sector say compliance requirements have grown increasingly complex. And 84% are looking to technology to help them automate and optimise compliance and transaction monitoring activities.

Many new regulations affecting the industry focus on cybersecurity and cyber resilience due to the growing risk of fraud, ransomware, hacking and other threats. For example, worldwide losses due to fraud in banking are projected to **reach \$58 billion** by 2030. Rising fraud and compliance costs, especially when combined with other trends such as dwindling customer loyalty, threaten to shrink the financial sector's current profit margins.

worldwide losses due to fraud in banking
are projected to reach

\$58 billion

by 2030

Given all these developments, it's clear that financial services organisations need to prepare for a future business environment that could look very different from today's. And they're increasingly looking to AI and agentic AI to help them meet these challenges.



What are AI agents and agentic AI?

AI and agentic AI represent different layers of capability. Before exploring how financial institutions use them, it's important to clarify what these systems are and how they differ.

Artificial intelligence (AI) refers to systems and technologies that use data and advanced algorithmic models—including machine learning and deep learning—to execute tasks traditionally requiring human intelligence, such as predicting outcomes, making or recommending decisions, automating processes and generating content through platforms like generative AI and large language models. Beyond its technical capabilities, AI serves as a catalyst for business transformation within financial services, enabling organisations to innovate, optimise operations and drive growth. Central to this advancement is the commitment to Responsible AI practices, ensuring ethical governance, transparency and accountability in the deployment of AI solutions.

Agentic AI goes beyond the capabilities of traditional AI, which is limited to identifying anomalies and scoring risks without being able to take independent action. It also outperforms generative AI, which can only create content in response to specific prompts but cannot execute a sequence of tasks autonomously. Furthermore, agentic AI is more advanced than automated systems that combine automation and AI, because these still require substantial human oversight to function.

In contrast, agentic AI can operate with minimal human intervention. While critical decisions in financial organisations will continue to require human accountability, AI agents can serve as digital colleagues, applying reasoning to tasks, adapting to new situations and using a variety of APIs and tools to achieve outcomes with much less human involvement.

Agentic AI is redefining what is possible for financial services institutions offering innovative solutions to longstanding industry challenges. Within the financial services industry, many companies are emerging as early adopters of AI, driven by margin pressure, regulation, and competition. These Frontier organisations are leading the shift to human-led, AI-operated models where agents scale decision-making, automation, and speed across the enterprise.

By automating complex tasks, enhancing customer experiences and strengthening compliance, agentic AI empowers banks, insurers, asset managers and other financial services organisations to operate with greater efficiency and agility. As regulatory demands and market expectations evolve, adopting agentic AI positions will enable financial services firms to lead with resilience and deliver sustainable value in a rapidly changing environment.



Business value of agentic AI in financial services

By deploying AI agents across many parts of their businesses, banks, insurance firms and other financial services institutions can reduce the need for human support across entire processes such as customer service, pricing, underwriting and claims settlement. This automates and speeds up many tasks that previously required human decision making based on person-to-person conversations or manual review of documents. In addition to automating insights and improving efficiency, agentic AI can help businesses to maintain or increase their levels of service despite staffing challenges created by talent shortages or large numbers of employees reaching retirement age. Using AI agents, organisations can also open new opportunities to increase revenues and grow the business.

30%

A 2025 PwC survey of financial services executives found that **30% have made agentic AI a top investment priority** for 2026.

But successfully implementing AI, especially agentic AI, is a multi-pronged challenge. AI applications can occasionally behave in unexpected or unpredictable ways, sometimes by generating false responses known as hallucinations, or by taking actions that might not always be in the best interests of customers. To prevent the likelihood of this, enterprises need to strictly align the behaviour of AI agents with their business rules and ethical norms. They should also adopt a human-in-the-loop approach that requires a human to give final sign-off on certain agent actions that exceed defined thresholds, based on risk factors, transaction values, or other considerations aligned to organisational policy.

In addition, it's critical that financial services companies invest the time and effort to manage AI-related challenges around data privacy and security, potential bias in AI models, regulatory compliance and customer trust.

It's worth noting that use cases can vary geographically depending on regional regulations, preferred business models, infrastructure maturity and market readiness. See **Regulatory considerations for the EU financial sector** for details about key regulations that are having the significant impacts on organisations across Europe.



Regulatory considerations for the EU financial sector

In force since January 2025, Europe's **Digital Operational Resilience Act** (DORA) was enacted to strengthen cyber resilience of organisations in the financial sector and harmonise existing regulations across the continent. Among its requirements: information sharing about cyberthreats, board-level responsibility for compliance failures and monitoring of risks related to third-party providers of critical ICT services.

The EU's **AI Act** establishes different levels of control for how organisations across sectors use AI applications based on risks. Adopted in June of 2024, it will be fully in force in 2026, with some provisions applicable sooner, such as its ban of AI systems that pose unacceptable risks, which started to apply in February 2025. That prohibition on unacceptable levels of risk includes applications for cognitive behavioural manipulation, social scoring, biometric identification and categorisation, and real-time and remote biometric identification. The AI Act also sets security and other requirements for high-risk applications, requires transparency for applications with limited risks and leaves minimal risk uses unregulated.

Since May 2018, the EU's **General Data Protection Regulation** has protected the privacy and security of people who interact with organisations that control or process their personal data. It requires such organisations to handle personal data in a lawful, fair and transparent manner; collect data only for specified, explicit and legitimate purposes; limit data collection to only what is adequate, relevant and necessary; ensure that data is up to date and accurate; store data for no longer than is necessary; and process data securely to protect it from unauthorised uses, loss, destruction or damage.

Agentic AI can play a critical role in helping financial institutions stay compliant by continuously monitoring regulatory requirements, automating evidence collection and ensuring that controls are applied consistently across business processes. And bespoke AI tools that incorporate regulatory requirements are also already helping financial services organisations on tasks like analysing and optimising contracts with technology partners, helping financial organisations to better manage third-party risks as required by regulations such as DORA. Such tools can also generate documentation to comply with mandates for transparency and auditability.

As financial services organisations expand their use of agentic AI, they will need to ensure that they have sufficient oversight of their agent ecosystems to meet the industry's strict regulatory requirements. That's likely to require the use of a control tower or control plane, such as **Microsoft Agent 365**. A control plane is a foundational infrastructure for agentic AI that acts as a "system of record" and "system of control" to record and monitor the activities of agentic workflows.



02



Sector spotlights

With the right foundation in place, agentic AI will benefit both companies and customers across the financial services industry, from banking to insurance to capital markets. The next section outlines some specific impacts on each of these sectors.



Banking

Across the banking industry, significant immediate value from agentic AI is emerging in three domains: financial crime and regulatory processes, customer engagement and personalisation, and lending and credit decisioning. These areas combine high regulatory pressure, large volumes of complex data and multistep workflows—conditions where autonomous, reasoning driven agents deliver rapid, measurable impact.

Financial crime and regulatory processes

Agentic AI is reshaping the financial crime lifecycle by orchestrating data collection, classification and analysis across multiple systems. Banks automate Know Your Customer (KYC) and onboarding as agents compile, summarise and validate source-of-wealth information that previously required extensive manual work. For example, one major Asian bank reduced the task of creating source-of-wealth memos from roughly ten days to around one hour using agentic workflows, with final validation conducted by relationship managers. This reflects a wider trend: banks are adopting a hybrid model in which agents handle the bulk of the investigative and documentation work, while humans provide oversight at the key decision points regulators expect. These capabilities extend to anti-money laundering and fraud detection, where agents analyse realtime transactions to surface suspicious patterns and reduce false positives. The advantage comes not only from improved detection, but from full process orchestration—retrieving evidence, preparing case narratives, applying policy rules and triggering human review where required.

In other regulatory compliance areas, agent-driven workflows support **COREP**, **FINREP** and **DORA** reporting by generating structured evidence packs, maintaining audit trails and proactively monitoring control gaps. Banks are also building “control tower” governance structures to supervise agent behaviour, embed regulatory rules into workflows and ensure agentic systems remain fully explainable and compliant.

Darren Ketteringham, Partner, Financial Services Leader, PwC UK, stresses the speed of change being driven by agentic AI:



What’s striking about agentic AI in banking is the pace. These systems don’t just shave a few minutes off a task, they radically streamline end-to-end processes. KYC, transaction monitoring, fraud reviews, contact-centre journeys: agents can now handle the heavy lifting across all of them, with better quality and far more consistency than traditional automation ever delivered.



Customer engagement and personalisation

Agentic AI elevates every stage of customer interaction. In contact centres, agents summarise conversations, recommend next steps in real time and update records automatically—supporting consistent service and freeing colleagues to focus on complex cases.



~30%

Agent driven self-service can resolve ~30% of routine requests without human intervention, improving satisfaction and lowering operating costs.

Across marketing and relationship management, agents analyse behavioural, transactional and demographic data to refine segmentation and surface personalised offers. They can also use these insights to tailor existing services or launch new products that target the needs of specific types of customers. In wealth and asset management, multi-agent workflows help advisors synthesise signals from portfolios and client preferences to deliver highly tailored outreach and timely product suggestions. That increases engagement with higher-value customers and creates new opportunities for targeted growth, a must in an environment where many organisations are facing the “**profitability paradox**” – growing portfolios coupled with shrinking margins.



Lending and credit decisioning

Agentic AI is transforming lending from a traditionally manual, paperwork-intensive process into a dynamic, end-to-end digital journey. By automating document gathering, verification, and risk assessment, agentic systems enable faster, more accurate lending decisions and streamlined coordination among all parties involved. This shift is streamlining loan origination and servicing, enhancing transparency, and opening up new opportunities for personalised customer experiences in both consumer and commercial lending.

Agentic workflows automate document preparation, verification and risk assessment for both consumer mortgages and commercial loans. Emerging multi-agent platforms help banks orchestrate complex loan-book operations such as drawdowns, rollovers and repayments and enhance liquidity and hedging decisions by connecting disparate datasets.

Looking ahead, banks are experimenting with agent-assisted end-to-end journeys for high-stakes consumer transactions, bringing advances in lending process workflows together with customer experience and personalisation. For example, in the near future, a homebuyer could be supported from property search and offer negotiation through multi-lender applications, mortgage approval and insurance activation—all through orchestrated agentic workflows.

Patrice Witz, Partner, Microsoft Alliance Leader, PwC Luxembourg, on the implications of agentic AI:



For banks, the real breakthrough comes when agentic AI is treated not as an add-on, but as a trigger to rethink end-to-end value chains. Incremental automation of a loan process is not transformation on its own. It requires a fundamental redesign of operating models for an agentic world. Those who delay will see the gap widen rapidly: while the technology is advancing fast, the cultural, organisational and skills shifts required take time to develop, so companies should start early.



Insurance

In insurance, agentic AI can make a significant impact across underwriting and risk decisioning, claims and fraud management, customer servicing and retention, and workforce resilience. These areas reflect the sector's reliance on expert judgement, its use of unstructured data and its sensitivity to customer expectations.

Underwriting and risk decisioning

Agentic AI is reshaping underwriting from a periodic, document-driven activity into a responsive, continuously learning risk function. Rather than relying on scheduled updates or manual file assembly, underwriting decisions can draw on a steady flow of real-time signals that reflect what is happening across the portfolio.

AI agents automatically gather and organise the information required for a risk assessment, extracting details from application materials, historical policy and claims data, and relevant external sources. They flag missing or inconsistent elements and generate structured risk files that can be validated quickly—or, for simpler lines of business, approved without human involvement below defined thresholds. This reduces manual effort, improves consistency and enables straight through underwriting in areas where it was previously impractical.



A significant shift comes from the ability to update risk views continuously. Signals emerging from claims activity, behavioural patterns, weather models, imagery analysis and mobility data feed directly into underwriting models. Scores can adjust as exposure conditions evolve, whether due to predicted weather events, changing traffic risks or new behavioural indicators. Underwriting becomes a living, adaptive system rather than one constrained by periodic recalibration.

Agentic AI also expands the types of data that can be used operationally. High-resolution climate forecasts, drone and satellite imagery, and hyperlocal mobility patterns allow risk to be assessed at the level of individual properties, neighbourhoods or routes. This supports finer segmentation and enables coverage models that were previously too costly or complex to administer.

Because these agents surface emerging patterns automatically, insurers can design more targeted or event-responsive products without adding to underwriter workloads. Automated preparation and continuously refreshed insights make it feasible to experiment with microsegments and new pricing approaches.

Dorota Zimnoch, Worldwide Director for Innovation and Digital Strategy, Financial Services, Microsoft, describes the "Frontier shift":



AI is transforming how insurers build and evolve their product portfolios—but Frontier insurers are going further. By applying AI to unified, trusted data, they move from static products to continuously learning, adaptive offerings that surface new pockets of growth and respond in near real time to changing risk and customer needs. At the same time, more precise and cost-effective AI-driven fraud detection and risk intelligence are changing the economics of insurance, making it viable to cover risks that were previously uneconomic or off-limits. This is the Frontier shift: insurers moving from simply transferring risk to actively understanding, preventing, and dynamically pricing it—unlocking entirely new business models and sustainable competitive advantage.

Claims and fraud management

Modern AI agents no longer simply assist adjusters—they can now potentially orchestrate entire claims journeys autonomously, from first notice of loss through to payment initiation. These systems can capture evidence from customers via chat, voice, or digital channels; extract and verify policy information; detect inconsistencies; and generate structured case evaluations within seconds. Where human support is still required, agents surface key facts, highlight anomalies, and propose recommended next steps, substantially reducing manual review time.

AI-driven interaction also improves claimant experience: conversational agents maintain consistency, handle emotional contexts reliably, operate around the clock in multiple languages, and remove friction by replacing outdated phone-tree processes.

Looking forward, agentic systems may enable insurers to reduce some types of avoidable claims by acting on early-warning signals. Advanced models can extend weather forecasts with hyperlocal predictions of storm-related property damage, identify vulnerable assets using satellite or drone imagery, and send targeted alerts to customers with guidance on how to protect their homes. In motor insurance, agents continuously interpret granular road traffic patterns to detect emerging risks and adjust cover or routing suggestions accordingly.





These insights feed directly into underwriting, creating a continuous learning loop: real-time claims signals inform risk selection and pricing decisions without waiting for quarterly reviews. This helps insurers minimise loss ratios while offering more personalised products.

Fraud-focused agents collaborate to analyse behavioural signals, policy histories, claimant sentiment, document authenticity and cross-case patterns. By combining multiple detection strategies, they surface subtle anomalies that traditional rule-based systems miss—expanding the scope of cases insurers can review without adding friction for legitimate claimants.

As generative tools also empower fraudsters, these multi-agent frameworks provide the necessary defensive capability to keep pace, reducing investigative cost and broadening coverage.


The shift to agentic claims and fraud management is widening the performance gap between early adopters and slower movers. Insurers that modernise gain clear advantages in cost, speed, customer experience and risk accuracy, while those that hesitate face mounting pressure from rising claims volatility, more sophisticated fraud and tightening margins.

At the same time, underwriting, claims, prevention and fraud detection are becoming parts of a single, connected system in which signals generated in one area strengthen decisions in the others. This creates a self-improving loop that enhances portfolio resilience, improves pricing precision and accelerates decision cycles—leaving organisations that fail to adopt these architectures increasingly exposed in a market that is evolving more quickly with each cycle.


Customer servicing and retention

Agentic AI enables insurers to provide always-on, multilingual, omnichannel servicing. Agent-powered systems are improving insurers' ability to deliver round-the-clock capability across multiple regions and languages, so they can provide support when and where their policyholders need it.

Agents can also handle policy questions, claims updates and administrative requests in a unified flow, reducing unnecessary handoffs between departments. When escalation is needed, agents transfer the case with full context so human colleagues can resolve it quickly. This helps provide customers with an improved omnichannel experience.




54%



of organisations deploying agentic AI have seen improvements in customer experience.

57%



are using or soon plan to use AI agents for customer service.

PwC's **AI Agent Survey** found that 54% of organisations deploying agentic AI have seen improvements in customer experience, with the number of customer-facing use cases rising quickly. Out of all companies surveyed, 57% are using or soon plan to use AI agents for customer service.

Workforce resilience

Along with providing powerful tools to streamline workflows in areas like claims and underwriting, Agentic AI is transforming workforce resilience. Many insurers face a structural capacity gap: decades of expertise concentrated in a shrinking workforce and limited new talent entering fields like claims, operations or actuarial work. **Recent benchmarking** shows that more than one-third of employees in the European insurance sector are over the age of 50. Agentic systems help preserve critical expertise by capturing how experienced staff make decisions and translating that knowledge into operational models that can continue to perform reliably as teams contract.

As this shift accelerates, a new hybrid model of work is emerging. Employees interact with agents throughout day-to-day processes, supervising outputs, guiding improvements and resolving exceptions. These interactions form a continuous training loop that strengthens system performance over time. The result is a workforce that can maintain—and in some cases expand—its throughput even as headcount declines.



Successfully adopting this model also requires cultural change. Insurers are traditionally risk averse, and many organisations are still early in redefining roles and embedding agent-supported workflows. Those that move quickly will be better positioned to offset demographic pressures and maintain operational resilience, while slower adopters risk falling behind as demand rises and available talent continues to diminish.


Andreas Hufenstuhel, Partner, FS Big Data & Advanced Analytics Advisory, PwC Germany, describes a major workforce trend in insurance:



Companies are exploring the use of AI to get the same amount of work done, but with fewer people. This reduces the need to replace people who are retiring. A few leading companies are going further and exploring how agentic AI can help increase market share and redefine their USP as well. There's a real opportunity in the next 2 to 5 years to shift towards a new model where humans work hand in hand with AI agents to both keep the day-to-day business running smoothly and drive new opportunities for growth.

Just as insurers unlock new value in risk and claims processes, capital markets and private equity firms are applying agentic AI to their own high-stake analytical and operational workflows — accelerating insight generation, enhancing due diligence quality and strengthening real-time risk management.





Capital markets

(including private equity)

In capital markets and private equity, including both financial exchanges and financial infrastructure providers, the impact of agentic AI is emerging across research and analytics acceleration, deal and due diligence operations, client advisory and personalisation, and risk, liquidity and compliance management. These areas depend on synthesising diverse information under tight time constraints — an ideal environment for multi-agent orchestration.

Research and analytics acceleration

Agentic AI accelerates investment research by transforming multiday analytical workflows into fully orchestrated processes that run in minutes. Agentic systems ingest financial statements, earnings transcripts, market news and alternative data to produce structured analytical drafts within minutes. Portfolio managers and analysts can then focus on judgement-driven evaluation and scenario exploration rather than manual synthesis.

When external shocks occur, such as a sudden policy announcement or the introduction of a new tariff, agents rapidly assess likely portfolio impacts by tracing exposure pathways across positions and sectors. This enables investment teams to respond quickly and confidently to market-moving events.

Marc Pfeiffer, EMEA Capital Markets Leader, Microsoft, describes how AI helps portfolio managers accelerate their research:



For portfolio managers, the breakthrough is getting an intelligent first view and seeing instantly how a market event flows through their holdings. Multi-agent systems can pull structured and unstructured sources together, deliver a first cut in minutes, and show how a sudden shock such as a new tariff could cascade through a portfolio. That allows portfolio teams to spend their time on judgement rather than on stitching data together.

Risk, liquidity and compliance management

Agentic AI fundamentally reshapes deal and due diligence operations by automating the time-intensive steps of data ingestion, verification and risk identification. Agent-driven orchestration accelerates due diligence cycles by automating data room ingestion, checklist management and risk flagging. In private equity, multi-agent workflows track operational KPIs across portfolio companies and surface improvement opportunities — enabling deal teams to focus on value creation rather than document handling.

Client advisory and personalisation

Agentic AI elevates client advisory by generating continuously updated, personalised investment insights drawn from real-time market, portfolio and behavioural signals. Agentic AI synthesises market developments, portfolio exposures and client behaviours to surface timely investment opportunities and generate custom insights. Advisors benefit from real-time signals indicating when client outreach is warranted. As firms increasingly operate within wider data ecosystems, agents can draw on both internal and external sources such as partner, supplier and macroeconomic feeds to contextualise opportunities and tailor advice more precisely.



Deal and due diligence operations

Agentic AI strengthens risk, liquidity and compliance management through continuous exposure monitoring, dynamic scenario expansion and automated regulatory workflows. Agents support continuous monitoring of exposures, liquidity positions and concentration risks. They also expand stress testing capabilities by modelling a wider range of macro financial scenarios, adjusting inputs dynamically and routing results for human validation. Beyond risk, agent-driven workflows handle recurring compliance tasks, improving transparency and reducing manual reporting burdens across capital and liquidity management processes.



Across banking, insurance and capital markets, a common pattern is emerging — agentic AI delivers significant value when it reshapes decision-heavy, high-complexity processes, transforms customer experience, or strengthens the organisation's ability to make faster, more consistent and more transparent choices.



03



Becoming a Frontier organisation

A leadership roadmap for scaling agentic AI with trust and speed

Agentic AI presents a significant opportunity for financial services institutions to drive transformation, deliver exceptional service and create lasting enterprise value. Microsoft describes such organisations as “Frontier Firms” — businesses operating at the leading edge of AI adoption, where human expertise and agentic AI collaborate to deliver on-demand intelligence at scale. Moving toward this model requires a clear strategic vision, a practical roadmap and robust governance.

This report has explored a wide range of high-value applications for agentic AI — from transforming claims handling and underwriting in insurance, to enhancing fraud detection, customer experience and regulatory compliance in banking, to accelerating research, due diligence and risk management in capital markets. The challenge is not a shortage of opportunities, but selecting and implementing those that can scale across the organisation. Transformation begins with taking a process-first mindset, identifying the core business processes that drive value and anchoring AI initiatives around them, rather than starting with technology or tools. Prioritisation should then focus on capabilities that strengthen core value chains, connect previously siloed activities and create reusable foundations for enterprise-wide transformation. Interconnected value chains amplify total impact, and organisations that invest in robust, repeatable agentic capabilities realise far greater benefits than those deploying isolated, one-off use cases.

Lilia Christofi, Partner, EMEA FS Data and AI Leader, PwC UK, highlights the importance of building a solid cloud and data foundation:



By developing blueprints for AI solutions, companies can design them for scalability across the organisation. This approach enables enterprises to extend a capability that has proven effective in one business area by retraining it on new datasets and redeploying it in other contexts. Such reuse is crucial, given the significant investment required to build AI solutions.

Sustained success comes from building enterprise-wide AI capabilities that reinforce each other over time. A coherent roadmap ensures that ambition is matched by execution. This begins with a strategic vision supported by foundational pillars: delivering personalised and proactive service; embedding responsible, transparent AI standards; empowering colleagues with advanced tools and insights; and investing in scalable, future-ready platforms. Together, these pillars guide the effective deployment of agentic AI across the organisation.

To translate ambition into outcomes, high-level goals must be anchored in practical principles that govern every stage of the transformation. Connecting purpose-driven pillars with operational rules ensures that adoption remains both innovative and responsible — strengthening trust, capability and governance while accelerating value creation.

When developing an AI transformation strategy, financial organisations need to answer six key questions: What is our vision and appetite for risk? What knowledge and skills do we need to promote? What will our transformed operating model look like? What is our required technology foundation? How do we ensure Responsible AI? And which partners will help us achieve our goals?



Six pillars for agentic AI success

1

Vision and risk appetite

AI will power the vision for modern financial services, anchored in customer value, trust and a clearly defined risk appetite that enables bold innovation with prudent guardrails

Ambition with assurance: AI that's fearless on value, disciplined on risk

2

Literacy and skillsets

We will raise organisation-wide AI fluency through role-based learning paths, communities of practice and hands-on adoption

Support an AI-confident workforce: skills that turn intent into impact

3

Operating model

A federated, product centric AI operating model, central standards, local delivery and measurable outcomes

One way to build, many ways to win: governed, agile, value-led AI

4

Technology

A secure, scalable data + AI platform (lakehouse + MLOps + agentic capabilities) with a predictable cost model

Engineering for trust at scale: AI that is powerful, practical and privacy-first

5

Risk and governance

Responsible AI by design: clear accountability, transparent controls and ethical standards that protect customers and society

Support ethical decision-making governance that moves at the speed of AI

6

Partnerships

Ecosystem partnerships across companies, fintechs, academia and regulators to accelerate safe innovation and community benefit

Together, faster: partnerships that turn mutual purpose into AI progress



From vision to action: prioritise value chain impact

The roadmap begins by turning strategy into selection. Use cases should be prioritised where they are embedded in value chains and can demonstrably lift outcomes across customer, risk and operations — not as isolated proofs of concept. Multiple interconnected value chains can amplify overall value, provided they are sequenced with a clearly defined risk appetite and transparent guardrails. Reuse is key: develop blueprints so each success becomes easier and faster to replicate across the organisation.

Prepare your people: raise AI fluency and confidence

Truly transformative AI deployments require a workforce that's prepared to use new tools and work in new ways. Strong technology is just the beginning; successful agentic AI deployments require change management and support for your organisation's people through education and access to the right resources. PwC's **recent workforce research** shows that employees are adopting AI faster than many leaders expect, and are generally positive about using technology that removes repetitive work. However, they want clarity on how AI will affect their roles and support to build the skills they will need in an AI-enabled enterprise.

Making the case for change should include clearly describing the purpose behind agentic AI deployment, including the specific business problem being addressed and the ways in which the solution aligns with your organisation's transformation roadmap. Clarifying these aspects helps secure the broad organisational support necessary for successful AI adoption.



You also need to be prepared to adapt your existing business processes to the new capabilities that AI can bring. Agentic AI, for example, can automatically manage workflows that might have previously required different teams from different parts of your organisation, as previous examples show. As a result, teams might need to be restructured. And it might be necessary to reimagine work more comprehensively.

For example, if you're a bank looking to use AI agents to streamline lending, you might need to re-examine the entire loan process, rather than simply look for ways to improve a few steps with the help of AI. Subsequently, value chains can end up looking very different from the ones that you're used to today. A process that once had 20 steps, for instance, might require only three steps in an agentic world.

Put structure around delivery: establish your operating model and CoE

A critical prerequisite for scaling AI successfully is the creation of a robust AI Centre of Excellence (CoE) that provides the structure, standards and capabilities needed to industrialise delivery. An effective CoE establishes the foundations — data readiness, technology, security and architecture — while embedding AI Trust by Design to ensure responsible and compliant use. It enables repeatable, scalable solution delivery through common patterns, platforms and blueprints, supported by clear value realisation frameworks and performance metrics.

Equally important, the CoE drives workforce enablement by building skills, fostering adoption and supporting continuous learning. Together, these elements create a governed, efficient and high-impact AI engine that balances speed, risk and cost, and consistently delivers meaningful outcomes across the organisation.

An AI Centre of Excellence plays a central role in governing modern data platforms — setting standards for ingestion, lineage, quality checks and model deployment so that agentic systems operate on trusted, well-defined data. This unified governance ensures that complex, multi-agent workflows can be deployed consistently across business units without duplicating infrastructure or creating parallel data silos.

Technology — creating a modern, scalable foundation for agentic AI

As enterprises gain experience with agentic AI, their technology foundations must evolve to support more sophisticated applications. Modern AI systems increasingly rely on coordinated workflows in which multiple agents carry out different steps of a process — and, with the right guardrails, these orchestrated workflows can manage complex journeys from end to end. This shift is already visible in use cases such as accelerated credit card issuance or rapid personal loan approvals, where several agents work together to intake information, verify documents, assess eligibility and complete decisions. Tools such as Microsoft Copilot Studio, with its library of prebuilt agents, can help organisations accelerate these early steps.

Yet orchestrated agentic workflows require a stronger technical backbone than many legacy environments can support. Integrating agentic AI brings new operational and regulatory risks, particularly in financial services, where organisations must demonstrate the security, traceability and resilience of their data and systems. Preparing for adoption therefore begins with strong data governance — role-based access controls, encryption, monitoring and guardrails to detect anomalous behaviour — supported by modern cloud infrastructure capable of running AI workloads consistently across business functions.

PwC's **2025 EMEA Cloud Business Survey** found that agentic AI capabilities are now decisive in organisations' cloud provider decisions, even though many are still early in scaling agents.



Prafull Sharma, Partner, Technology & Data FS Leader, PwC Switzerland, highlights the importance of building a solid cloud and data foundation:



Every AI transformation succeeds or stalls on the strength of its data. Governance turns fragmented information into dependable insight, and a well-architected cloud platform turns pilots into platforms. Together they let you scale AI responsibly, across functions, markets and regulatory environments.

Modern data platforms built on lakehouse architecture — such as **Microsoft Fabric** — are becoming essential foundations for scaling agentic AI. These platforms bring an organisation's data together in one secure, governed environment, replacing fragmented legacy systems. They make it far easier to manage information at scale, improve data quality and accelerate the development of new digital products. It's important to note that well-designed data warehouses are still an essential prerequisite. Ensuring that data lineage can be tracked, traced and documented is also critical. **Microsoft Fabric now includes agents** that make it easier for non-technical users to get precise, context-rich answers to questions about data.

Employing modern engineering practices around machine learning (MLOps) and AI (AIOps) help ensure that AI models and agent workflows remain reliable and up to date, creating the stability enterprises need as they begin to automate more decision making.

Against this backdrop, one global financial market infrastructure provider worked with Microsoft to consolidate dozens of data systems into a unified Fabric environment. The result: improved trust in its data and significant reductions in product development timelines.

When combined with automated processes for maintaining and updating AI models, robust data platforms provide the control required for advanced agentic capabilities — from coordinating several agents to complete multistep tasks to embedding AI directly into the tools colleagues use every day. These platforms also offer more predictable and flexible cost models, allowing organisations to scale AI usage up or down as needed without committing to always-on infrastructure. Together, these foundations create a technology environment that is flexible, resilient and ready for enterprise-wide agentic AI.

Embed trust — Responsible AI, governance and human-in-the-loop

Financial services organisations looking to deploy agentic AI must update their governance and ethical frameworks to ensure that they build and maintain trust with both customers and regulators.

Organisations can reap significant and sustainable benefits from AI when they adopt a robust approach towards Responsible AI that addresses risks in a transparent, consistent and accountable way. This involves considering risks around data, AI models, infrastructure and legal requirements, as well as how AI processes are integrated into the business and how AI tools might be intentionally or accidentally misused. PwC research shows that such an approach not only brings quantifiable benefits but enhances resilience as well.

Building a Responsible AI programme starts by assessing how well your organisation's processes, policies and operating model align with best practices and support the goals identified in your Frontier Firm transformation roadmap. From there, your business can develop an end-to-end governance framework that considers risks and necessary controls at every stage of the AI journey.

For agentic AI systems, it's critical to have strong, real-time controls. Preventative and reactive controls support safe use, and corrective controls are also vital to ensure that identified issues are resolved quickly. New faster delivery models where humans supervise and validate coded controls and work in mixed teams with non-human agents are also expanding what controls mean. It's important to go beyond technical issues and also monitor that outcomes align to business intent. Controls should include strict user authentication, safeguards for AI prompts and outputs, transparency regarding content sources, early detection and protection for vulnerable users, and centralised interaction logging.



For financial services institutions, it's critical to ensure that any decisions made or actions taken with the help of AI are transparent, explainable and well-supported. Use of AI needs to be unbiased, fair, secure, legal and ethical. And it's important to build accountability into AI systems by keeping a human in the loop to verify consequential decisions or actions.

Addressing data security, privacy and compliance

Data lies at the heart of any AI deployment. To ensure the best possible results, you need to make sure that your organisation's data is accurate, complete and up to date. It is critical to maintain privacy and security as required by regulators — and, where appropriate, to exclude certain data types to preserve trust and compliance.

When deployed responsibly, AI can help meet requirements around security, data privacy, money laundering, fraud prevention and more. Documenting compliance can be time and labour-intensive, but AI agents can help automate and accelerate many processes that once required manual effort.

Anisha Kaul, EMEA Banking Leader, Microsoft, explains the importance of strong data for building trust:



In financial services, trust is earned through actions. Microsoft strengthens that trust by investing deeply in cybersecurity, data privacy and Responsible AI. If AI is the plane, then data is the fuel — and no aircraft gets off the ground without the highest-quality fuel. A robust data foundation supports AI outcomes that are responsible, scalable and truly valuable for customers.

Connect the ecosystem — work with partners to accelerate safely

Ecosystem partnerships across financial institutions, technology providers, academia and regulators can accelerate safe innovation and deliver community benefits. They help shorten time to value, spread best practice, and provide the platforms and methods required to scale responsibly.

Once you've established the roadmap and the foundations required to become a Frontier organisation, the next step is partnering with teams that can help turn this ambition into reality. Scaling agentic AI demands both deep sector understanding and proven technology.



04



Why PwC and Microsoft

The right foundation

Financial institutions need both strategic insight and trusted technology at scale. PwC and Microsoft offer complementary strengths. PwC brings companies deep financial services expertise, regulatory understanding and transformation experience. Microsoft provides secure, enterprise-grade cloud and leading AI platforms designed for highly regulated environments. Both are important for companies looking to accelerate AI adoption with clarity, control and measurable business impact.

Patrice Amann, EMEA FSI Leader, Microsoft, describes the strengths that Microsoft brings to AI deployments in financial services:



For global financial services leaders, the next competitive Frontier is not AI alone, but enterprise intelligence — intelligence that is unified, trusted, compliant and operational by design. Microsoft has developed its IQ architecture which brings this to life by connection how people work (Work IQ), how the business reasons over its data (Fabric IQ), and how AI agents are grounded and governed (Foundry IQ). This unified intelligence layer enables organisations to scale innovation with confidence, turning trust, transparency, and regulatory discipline into a strategic advantage at board level. Our sharp focus on security, compliance, and trust are built in by design, and we are outcomes-driven, supporting customers with accelerators, strong industry data models, deep integration with core systems, and a powerful partner ecosystem, including firms such as PwC.

PwC and Microsoft are both experienced at helping financial institutions deploy agentic AI responsibly and at scale. PwC's Agentic AI Value Chain can be integrated with Microsoft technologies such as Copilot Studio, Agent 365, Fabric, Foundry and IQ architecture to help organisations automate complex processes, generate deeper insights and improve customer and colleague experiences. Both PwC and Microsoft are committed to Responsible AI principles. These support innovation which is transparent, explainable and aligned with regulatory expectations.




PwC acts as a long-term transformation partner to the financial sector, supporting organisations as they navigate market shifts, regulatory change and evolving client demands. We help leaders design AI-enabled operating models that drive resilience, competitiveness and growth. As one of the largest global adopters of Microsoft technologies such as Copilot Studio—with deployments across more than 100 countries—we apply first-hand implementation experience to our client work, ensuring that guidance is practical, tested and grounded in real enterprise-scale adoption.

Mauro Xavier, Partner, EMEA Microsoft Alliance Leader, PwC Spain, describes how PwC built its own AI capabilities:



PwC is our own Client Zero: we deployed Copilot across 230,000 users in more than 100 countries, making us one of the largest adopters globally. That first-hand experience, combined with the security and scale of Microsoft technology, equips us to help clients adopt AI responsibly, govern it effectively, and deliver through proven patterns.



Agentic AI is no longer an experimental frontier but a practical catalyst for transformation across banking, insurance and capital markets. Organisations that anchor their ambitions in strong governance, invest in reusable capabilities and empower their people to work confidently with AI will define the next era of financial services. The firms that move first, scale responsibly and collaborate across the ecosystem will not only navigate the turbulence ahead but also shape a more resilient, intelligent and customer centric industry.

PwC and Microsoft are both ready to support your agentic AI journey. Reach out to learn more.

Authors

Microsoft



Patrice Amann

EMEA FSI Leader
Microsoft



Anisha Kaul

EMEA Banking Leader
Microsoft



Marc Pfeiffer

EMEA Capital Markets Leader
Microsoft



Dorota Zimnoch

Worldwide Director for Innovation
and Digital Strategy, Financial
Services
Microsoft

PwC



Mauro Xavier

Partner, EMEA Microsoft Alliance
Leader
PwC Spain



Lilia Christofi

Partner, EMEA FS Data and AI Leader
PwC UK



Darren Ketteringham

Partner, Financial Services Leader
PwC UK



Andreas Hufenstuh

Partner, FS Big Data & Advanced
Analytics Advisory
PwC Germany



Prafull Sharma

Partner, Technology & Data FS Leader
PwC Switzerland



Patrice Witz

Partner, Microsoft Alliance Leader
PwC Luxembourg

