

In collaboration
with Accenture



Industries in the Intelligent Age

The AI Playbook for Financial Services

INSIGHT REPORT

JUNE 2026

Contents

Foreword	3
Executive summary	4
Introduction	5
1 The landscape today	7
1.1 Context and industry outlook	8
1.2 AI spending trends and prioritization	10
2 Investment and maturity of the industry	12
2.1 Business-driven adoption models	13
2.2 Aligning AI to strategy	13
2.3 Measuring AI value	15
2.4 Pilots versus scaled programmes	16
3 Successful AI starts with a strong foundation	17
3.1 The foundation and platform architecture	18
3.2 Skills, hiring and culture change	19
3.3 Governance, risk and compliance management	20
3.4 Security, model risk and threat landscape	23
4 The role of agentic systems in unlocking value at scale	24
4.1 Why agentic AI is key	25
4.2 Human–AI partnership	25
4.3 Examples of agentic AI in financial services	27
5 Planning for the industry’s future	32
5.1 Emerging technologies and cross-industry innovation	33
5.2 Framework for scaled and sustained AI adoption	33
5.3 Leadership, sponsorship and accountability	35
Conclusion	36
Appendices	37
A1 Overview of risk framework adoption by region (non-exhaustive)	37
A2 Overview of AI-driven risks and corresponding risk management approaches	39
A3 Additional reading	40
Contributors	41
Endnotes	47

Disclaimer

This document is published by the World Economic Forum as a contribution to a project, insight area or interaction. The findings, interpretations and conclusions expressed herein are a result of a collaborative process facilitated and endorsed by the World Economic Forum but whose results do not necessarily represent the views of the World Economic Forum, nor the entirety of its Members, Partners or other stakeholders.

© 2026 World Economic Forum. All rights reserved. No part of this publication may be reproduced or transmitted in any form or by any means, including photocopying and recording, or by any information storage and retrieval system.

Foreword



Andre Belelieu

Head, Financial Services Industries; Head, Business Engagement, North America, World Economic Forum



David Parker

Global Industry Lead, Banking and Capital Markets, Accenture



Drew Propson

Head, Technology and Innovation in Financial Services, World Economic Forum

The artificial intelligence (AI) landscape of 2026 would have seemed implausible just a few years ago, with models that reason across complex problems, agents that autonomously execute multi-step workflows, and capabilities that are fundamentally redefining what is possible inside financial institutions. Driven by rapid advances in generative and agentic AI, the pace of change shows no sign of slowing.

AI has been a priority on the World Economic Forum's financial services agenda for more than a decade, resulting in a large body of research. Over that time, the conversation has transformed, from the exploration of a technology in its nascent stages to an era defined by breakthroughs that are testing every institution's ability to adapt and navigate change.

In 2024, recognizing that the latest wave of AI advancement demanded an intensive and structured response, the Forum and Accenture collaborated to launch a multi-phase AI in Financial Services initiative, bringing together a diverse set of stakeholders across the financial sector to share perspectives, successes and concerns as firms integrate this disruptive technology.

This report is the outcome of Phase II of the initiative, which has engaged over 150 senior leaders from more than 100 organizations. Building on the Forum's 2025 white paper, [Artificial Intelligence in Financial Services](#), which examined AI's impact on firm-level strategy and the broader financial ecosystem, this report offers practical guidance for organizations at every stage of their AI journey. It is grounded in case studies that illustrate real-world applications with measured results, and is enriched by insights from roundtables held in Hong Kong, New York, London and Singapore, alongside interviews and desk research. This work also connects to the Forum's cross-industry AI Transformation of Industries initiative, a group consisting of over 500 members catalysing the ethical and sustainable use of AI for the benefit of all.

We hope this report provides confidence and clarity as AI continues to evolve at a remarkable speed, and that the insights within, garnered through trusted dialogue among peers across regions, contribute to a financial sector that embraces AI's full potential while managing its risks responsibly.

Executive summary

As artificial intelligence reshapes financial services at an intense pace, institutions must move with both urgency and discipline.

Artificial intelligence (AI) represents a structural disruption, and its impact on financial services goes far beyond efficiency gains. It is reshaping how work is organized, how decisions are made, how value is created and how institutions compete and collaborate.

As the industry moves from experimentation to scaled deployment, the central challenge for organizations is fully integrating AI quickly and securely. The decisions involved, from data management to governance structures, are numerous and varied. This report aims to help financial services firms across all geographies and sectors work through the complexity.

The report's key findings are as follows:

AI must be a strategic leadership issue.

The priority has shifted from running pilots to embedding AI across the enterprise, quickly, responsibly and in ways that deliver lasting value. Leaders must define where to focus, what to automate and where human judgement remains essential, and how to weave AI into day-to-day operations to deliver better outcomes for customers.

Differentiation will increasingly come from customer relationships amplified by AI.

With analytical and generative capabilities now widely accessible, organizations that use AI to intentionally strengthen customer relationships will be better positioned to deliver superior experiences, deepen retention and drive growth. Central to this is ensuring customers trust how the technology is being used.

The biggest benefits will come from redesigning holistically.

While piecemeal deployment can bring some local benefits, firms will capture more value when they rethink workflows, organization design and technology architecture together, rather than layering AI onto fragmented legacy structures or processes.

A scalable enterprise intelligence platform is essential. Institutions need a unified layer that coordinates data, models, decisions and automation, underpinned by robust data management, identity controls, auditability and resilience. Without this intelligent platform, AI cannot be deployed consistently, securely or at scale.

People must remain firmly in the lead. Successful AI adoption depends on people embracing and driving the transformation. This will take the form of a hybrid workforce, where people and AI agents work together supported by evolving roles, skills-based architecture and continuous capability development. Human judgement and accountability remain essential to both safe adoption and lasting value.

Risk management and responsible AI need to be embedded at every layer. Institutions must define risk appetite early, validate models rigorously and ensure AI actions remain explainable, traceable and controllable. This means deepening model-risk disciplines and incorporating jurisdiction-specific compliance requirements into existing enterprise risk frameworks.

A two-speed implementation model sustains momentum. Institutions are increasingly pursuing two speeds simultaneously, using AI to deliver near-term gains in productivity, experience and decision quality, while building the enterprise foundations required for long-term transformation. That balance between practical delivery and structural reinvention is what will separate experimentation from scaled impact.

Although there is no single path to success and requirements will inevitably evolve, organizations that proactively define strategy, align leadership and build adaptive foundations will be well equipped to capture the opportunities AI presents.

Equally, as AI continues to reshape the financial system, greater collaboration across firms, policy-makers and regulators will be essential to managing shared risks and unlocking opportunities that no single institution can capture alone.

Introduction

As AI integration in financial services intensifies, experiences from early adopters are beginning to offer meaningful lessons.

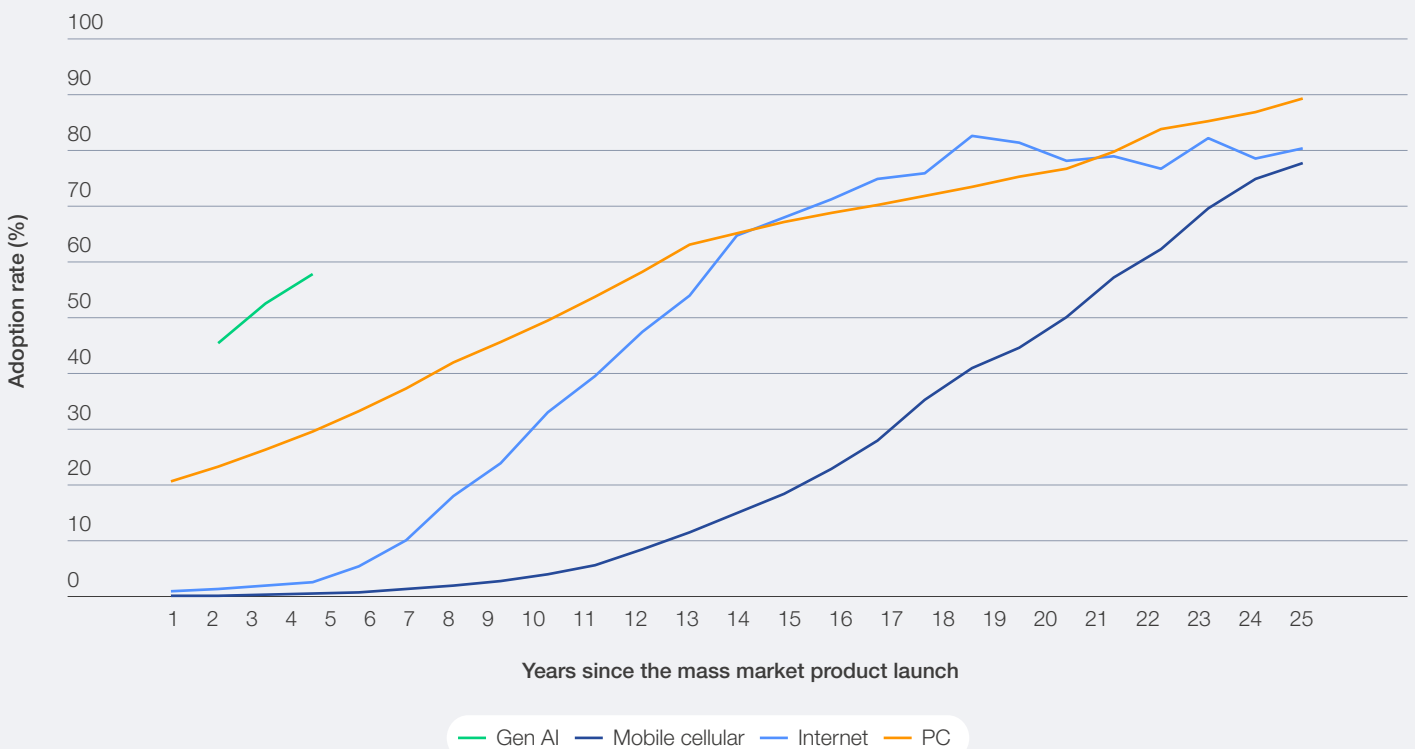
Artificial intelligence (AI) has taken centre stage in financial services discussions, as the technology touches every corner of the industry from banking and capital markets to insurance, asset management and payments. Reshaping the ecosystem across multiple dimensions simultaneously, AI is driving productivity improvements, redefining how firms manage risk and earn trust, opening the door to the reinvention of business models and prompting regulators to act.

Customer sentiment continues to evolve as well, with advances in AI elevating people's expectations of what their experiences could and should be. These trends are visible throughout the sector and in adjacent industries, such as retail and health. Figure 1 shows the relative rate of AI adoption compared to other technologies.

FIGURE 1 Adoption rates across technologies

AI adoption is outpacing previous technology waves

Adoption of AI, internet, PC and cellphone



Source: Accenture Research based on Real-Time Population Survey (RPS), Current Population Survey (CPS), United States Census Bureau and Our World in Data.

To date, AI has been shifting faster than many organizations have yet been comfortably able to absorb, deploy and govern it. For example, advances in reasoning capability and contextual understanding (commonly called “extended context”), along with the technology’s ability to

learn from users, are driving exponential increases in AI agents’ capabilities. Yet just one-third of respondents in a recent global survey of 1,320 C-suite executives and 4,560 employees across 20 industries and 12 markets said their talent strategy is fully aligned with their technology and AI strategy.¹

The financial sector faces a range of possible futures, with expectations of significant opportunities on one side and significant concerns about risk, trust and disruption on the other.

In response, financial services companies have intensified their focus on AI integration. They are honing their plans and prioritizing technology investments in generative AI (genAI), agentic AI and their underlying enablers, building on prior investments in traditional AI and digital transformation. After a period of experimentation, they are now beginning to expand the possibilities of what can be achieved by using a wider set of data products and more advanced AI models.

As they do, however, they face familiar constraints including fragmented data, outdated technologies and systems, and unclear regulatory requirements, coupled with uncertainty about how to capture sustainable value and advantage in uncharted territory.

To help firms navigate these challenges, this report provides a structured view of how AI is reshaping financial services, examining the current landscape, industry adoption trends and the evolving regulatory environment, while exploring how leading organizations are investing in and operationalizing AI across strategy, business models and value measurement. It also outlines the foundational capabilities required to scale AI responsibly, including data, talent, governance and risk management.

The result is a practical playbook, grounded in real case studies and emerging success stories, designed to help financial services organizations translate AI's potential into measurable, sustainable value. This report is equally intended as a resource for policy-makers and regulators as they work to keep pace with technological change and safeguard the stability of the financial system.

The decisions made now about where to act, how to build and how to govern will define the industry for years to come.

1 The landscape today

Recent AI advancements, combined with an abundance of data and language tasks across financial services, point to further industry transformation.





Private and public sector leaders navigating the AI landscape in financial services need a clear baseline, one that captures the current state of the technology, as well as its trajectory and the practices organizations are adopting today. What follows establishes this foundation.

1.1 Context and industry outlook

Context

While financial services companies have used AI extensively for decades, only recently has the technology crossed a critical technical threshold, becoming context-aware within non-sequential architectures. AI systems can now respond in natural language and maintain coherence across an entire interaction rather than performing isolated, linear steps. That is, the technology now possesses human-like reasoning abilities, allowing it to operate with far less reliance on people than before, and it continues to advance rapidly, both functionally and in real-world applications. The impact is far-reaching, as:

- Natural language is becoming a practical interface for work and services, **extending access to advanced capabilities** beyond a specialist team. More people can create what was previously out of reach more quickly.
- The **marginal cost of creation** is falling sharply, **trending towards zero** for many forms of content and analysis, enabling more rapid experimentation, scaling and innovation and getting new things done faster with less input cost.
- Agentic systems can also execute tasks end-to-end, **pushing down the marginal cost of executing tasks beyond creation** and **freeing people to do higher-value work** and accelerating business processes.

AI, therefore, decouples headcount from organizational capacity beyond what was previously possible with traditional AI and automation, unlocking what is generally accepted as the next major industry innovation and productivity shift. It is poised to disrupt financial services significantly, due to the industry's vast amounts of data, high concentration of language tasks and large volume of manual processes.

Industry outlook

The industry is beginning to fundamentally redefine what it means to deliver financial services.

Internally, advancing AI technologies are both enabling and compelling executives to reinvent their organizations end-to-end (strategy, people, process, technology). AI could ultimately allow financial services companies to operate as AI-orchestrated,

event-driven platforms. The back office may become fully automated, operating around the clock with minimal to no human intervention. The middle office would also be able to run continuously, with real-time risk and pricing enabled by AI.

On front office operations, organizations will be increasingly proactive, anticipating customer needs and acting autonomously with customer consent guidelines. However, they are also balancing the tension between what is technically possible and what customers feel comfortable with. Moving too quickly risks eroding customer trust and damaging the brand, but moving too slowly risks being left behind by more agile competitors.

For example, the technology can now create highly personalized experiences with contextual conversations, anytime-anywhere services at scale, essentially creating individual customer propositions (a market of one). In this vision of the future, customers will be able to experience financial services as a trusted, always-on capability embedded into their everyday lives and acting on their behalf. An agent (or an agentic system) anticipates their needs, evaluates options in context, executes actions proactively (with consent) and clearly explains outcomes. Customers won't need to navigate apps, fill in forms or handle fragmented decisions. Instead, they will interact with their financial services providers through digital assistants, AI-assisted operations centres or relationship managers.

Future digital assistants may not even be owned by the financial service provider, raising questions this report will explore in later sections, about what entity owns the relationship and how that affects customer-institution dynamics, and trust and responsibility concerns.

As an indicator of consumers' appetite for AI solutions, 71% of banking consumers globally say they would welcome an AI assistant in their primary bank's mobile app. However, 82% want to approve an agent's actions before they are executed.²

Societal, regulatory and prudential obligations

The rapid adoption of AI is already attracting heightened societal, regulatory and prudential scrutiny, introducing greater caution into how institutions direct and govern their AI initiatives.³

71%

of banking consumers globally say they would welcome an AI assistant in their primary bank's mobile app.

“ Most financial services companies are experimenting with internal process improvements, responsible AI governance aligned with existing structures, and the use of copilots and basic AI training.

For example, AI regulations may limit experimentation or force a change in direction, even as the technology evolves and organizations learn new ways to tap its potential. Hyperscalers, the large cloud service providers used by many companies to deploy advanced AI models or the data products they rely on, add another layer of complexity to the compliance landscape.

On the regulatory front, recent changes in the US include Executive Order 14365, *Ensuring a National Policy Framework for Artificial Intelligence*; the *Financial Services AI Risk Management Framework* published by the US Treasury; and *A National Policy Framework for Artificial Intelligence*. These are setting national-level expectations for AI governance, model risk and accountability while explicitly avoiding prescriptive constraints on deployment.^{4,5,6}

In the UK, the Financial Conduct Authority’s (FCA) *AI Update* and the launch of AI live testing formalize a principles-based supervisory approach, embedding AI oversight into existing frameworks such as the Consumer Duty regime and the Senior Managers and Certification Regime (SM&CR) rather than introducing standalone AI regulations.⁷

In Europe, the EU Artificial Intelligence Act, Regulation EU 2024/1689, introduces binding, risk-based obligations across financial services, with phased requirements from 2025 that materially raise expectations for governance, transparency, human oversight and supervisory

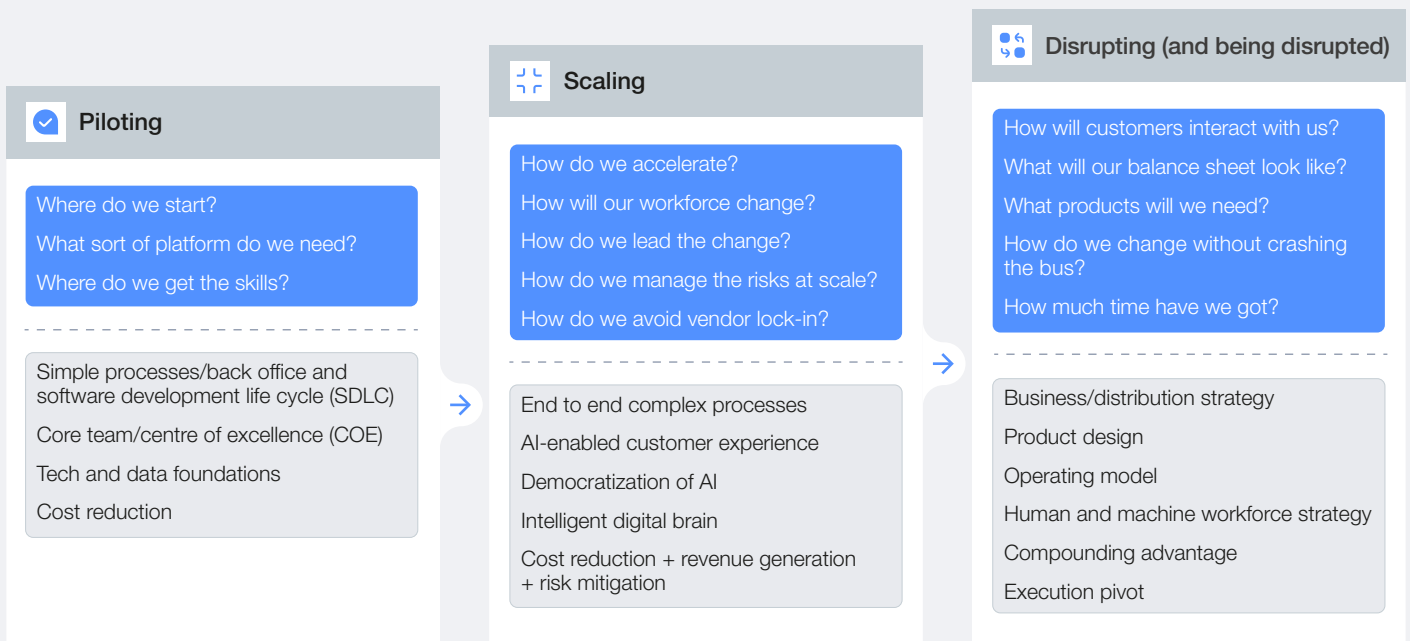
accountability, particularly for high-risk use cases such as credit, pricing and market activity.⁸

Across the Asia-Pacific region, regulators have long been converging on sector-specific AI governance frameworks, including the Monetary Authority of Singapore’s *Principles to Promote Fairness, Ethics, Accountability and Transparency (FEAT)*, their recently published *Artificial Intelligence (AI) Risk Management Toolkit*, and the Hong Kong Monetary Authority’s *High-level Principles on Artificial Intelligence*, signalling greater supervisory attention as AI moves from experimentation into core financial services operations.^{9,10,11} As another example, in Indonesia, the Financial Services Authority has issued AI governance guidance for banking,¹² while a presidential decree¹³ on the responsible use of AI is being prepared to provide overarching cross-sectoral guidance.

Investments

To date, most financial services companies have prioritized “no regret” AI investments. They are experimenting with internal process improvements, responsible AI governance aligned with existing structures, and the use of copilots and basic AI training. As experience with AI use cases and deployment grows, however, some are moving beyond isolated pilots, knitting together AI initiatives throughout and across functions and increasing AI’s presence in the front of the house. Figure 2 illustrates this progression.

FIGURE 2 Financial sector AI progression (illustrative)



Source: Accenture

1.2 AI spending trends and prioritization

“ IDC research finds that the global AI market is projected to rise to more than \$631 billion by 2028.

Global AI spend estimates must be treated with caution as figures vary considerably across analysts and methodologies and do not yet provide a comprehensive or consistent view across the emerging AI market in financial services. However, International Data Corporation (IDC) research finds that the global AI market is projected to rise to more than \$631 billion by 2028, with the software and information services, banking and retail industries leading in spending. The IDC also notes that genAI accounts for more than 19% of the total investment across these industries.¹⁴ Stanford University's *2026 Artificial Intelligence Index Report* reported that global corporate AI investment more than doubled in 2025.¹⁵

Recent research by NVIDIA, which surveyed financial services AI practitioners and executives, also indicates that financial services companies are increasing their AI spending in 2026. Most (83%) suggested spending would increase, and 44% suggested it would increase by more than 10%. The top three priorities for doing so were optimizing existing AI capabilities, exploring additional AI use cases, and building and providing access to AI infrastructure.¹⁶

Technology and service providers to the financial services industry observe that organizations are increasingly looking for localized ways to improve efficiency, with the goal of freeing up funds to invest in enterprise-wide AI foundations and pursue a more clearly integrated path to creating value with AI.¹⁷ They are seeing investment in scalable platforms, model management and enhanced compliance capabilities, as well as investment in AI across more complex mega-processes and value chains. Key call-outs on investment priorities are:

- 1 **Responsible AI guardrails and capabilities**, as organizations move to integrate ethical AI oversight into their governance frameworks and cadences, rather than running it alongside. In particular, they are increasingly using AI to support compliance mandates – for example, by deploying genAI agents to document AI decisions and produce audit-grade records. They are also using agentic supervisors to monitor other AI systems, validate policy compliance and trigger alerts when outputs reach guardrails.
- 2 **Foundational data and technical capabilities** to support AI systems in production. Financial services companies are prioritizing enterprise-grade data platforms, investing in modern data architectures that support real-time ingestion, large-scale unstructured data and governed access for AI models and agents. They are building AI platforms and coordination frameworks that manage model execution, agent routing, logging and life cycle controls, recognizing that scaling AI requires control planes and operating models beyond individual models or tools.
- 3 **Building AI skills and strengthening leadership and culture** to drive effective AI adoption. There is a surge in activity focused on enterprise-wide AI literacy and targeted upskilling, as leaders increasingly realize the importance of educating their employees to work confidently with and alongside AI agents and genAI. Only a portion of the workforce needs deep technical AI expertise to effectively use AI.



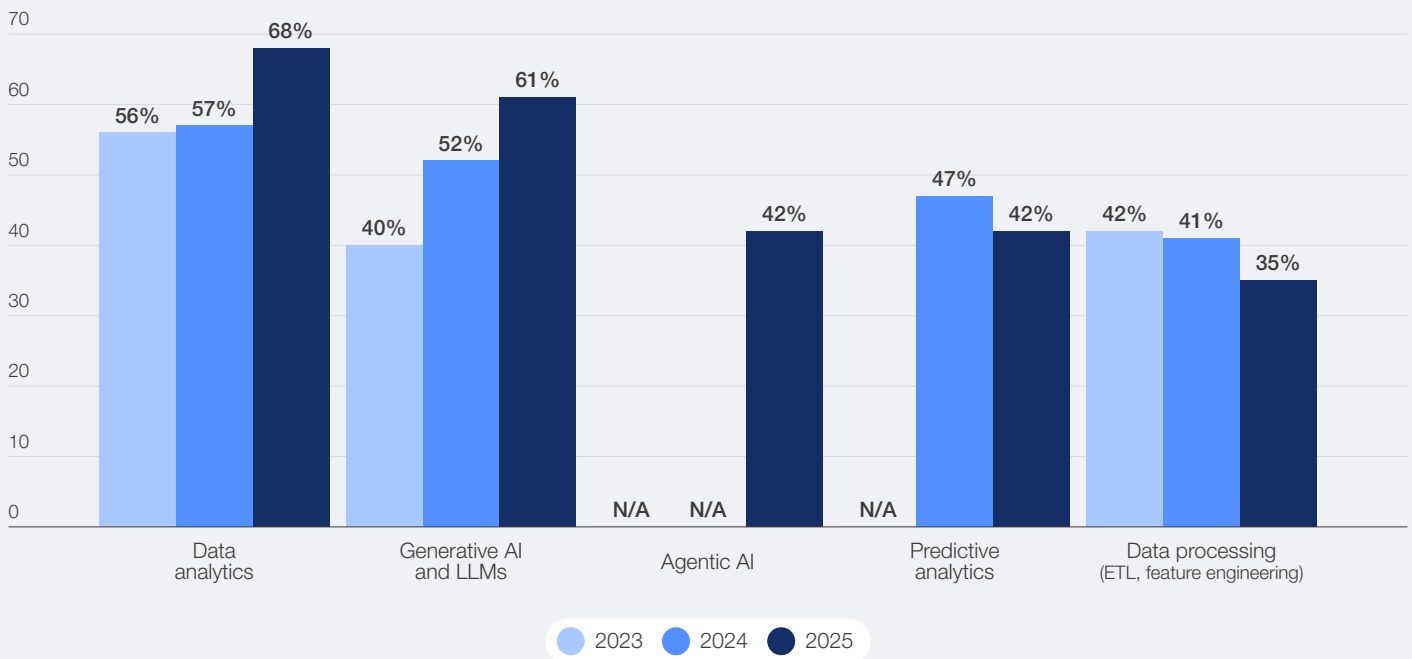
For business areas, the shift to front-office-facing AI is evident in NVIDIA's research, which suggests that most organizations deploy AI use cases in customer experience and engagement.¹⁸ Advanced tools, promising greatly enhanced experiences, will offer proactive advice, a frictionless experience and the ability to execute processes on a customer's behalf.

As Figure 3 shows, while financial services companies prioritized data analytics for use and assessment in

2025, genAI ranked second. Figure 4 highlights the three top AI focus areas by industry segment.

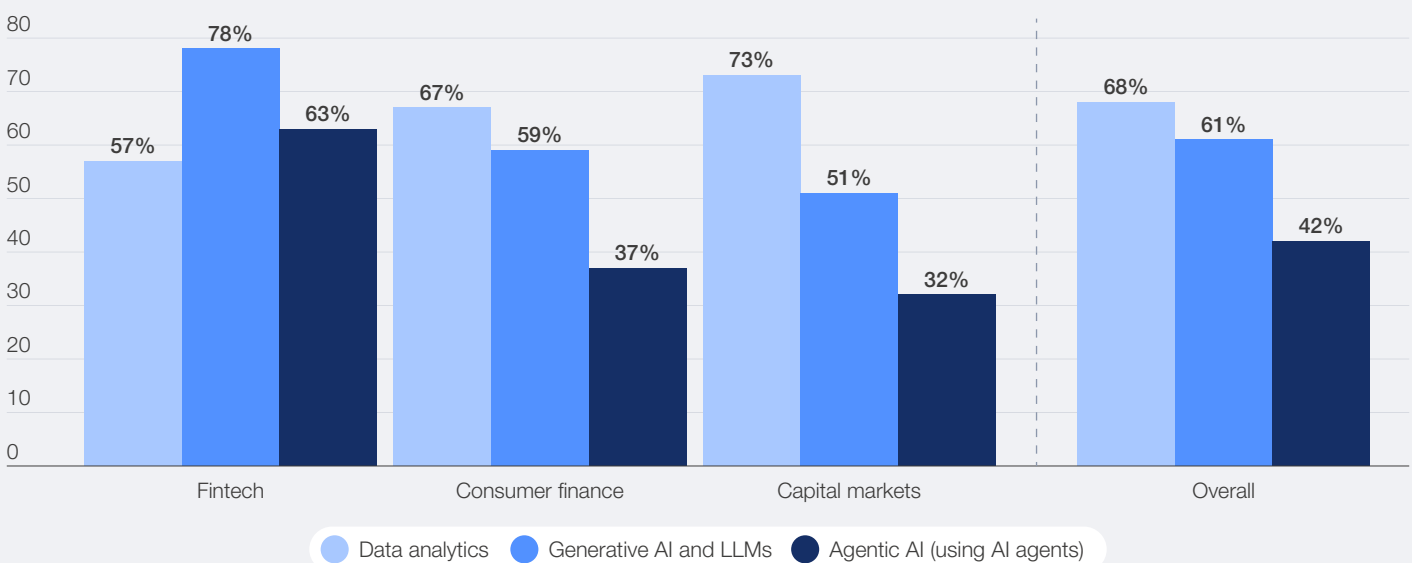
With capabilities, customer expectations and regulatory responses evolving, the landscape is ripe for further disruption. As new business models emerge, the value of AI as a driver of efficiency and incremental quality enhancements lessens as well. These qualities will become table stakes as disruption rewards bold moves into unfamiliar territory.

FIGURE 3 Top five AI focus areas for financial services institutions



Source: NVIDIA. (2026). *State of AI in Financial Services: 2026 Trends*.

FIGURE 4 Top three AI focus areas by industry segment



Source: NVIDIA. (2026). *State of AI in Financial Services: 2026 Trends*.

② Investment and maturity of the industry

Capturing the full value of AI requires strategic clarity, disciplined measurement and sound judgement about what to scale and when.



Successfully moving to scale with AI requires challenging some long-held foundational assumptions and embracing new business model concepts.

Consider, for example, what happens as agents begin to sweep deposits across providers, moving money from one account to another to optimize

customers' finances. Firms will no longer be able to count on deposits as a primary relationship driver. The agents that manage customers' finances proactively become their digital stewards and, by proxy, their trusted intermediaries; they hold the intimate customer knowledge, the context and, increasingly, the relevant history.

2.1 Business-driven adoption models

Industry, business and government leaders, aware of the rising challenges, are elevating AI integration to the C-suite and board level. This shift is playing out in the actions of some early movers, where CEOs (chief executive officers) are working in concert with talent and technology leaders to reimagine how their business strategy can be transformed using AI. These leaders invest the time to understand its possibilities and

then set the tone for the whole organization on how to adopt it.

For example, JPMorgan Chase is recasting AI from a toolbox into the bank's operating system, with strong leadership support. Chairman and CEO Jamie Dimon responded to sceptics at the Forum's 2026 Annual Meeting in Davos by highlighting the importance of business leaders taking charge.



It is what it is. And you can hope for the world you want, but ... you're going to get the world you got, and your competitors are going to use it ... and if you put your head in the sand, you will lose.

Jamie Dimon, Chairman and Chief Executive Officer, JPMorgan Chase¹⁹

2.2 Aligning AI to strategy

Leaders need to consider a number of strategic drivers and trade-offs, including:

1. **When to own the primary customer relationship.** Firms need to decide whether to own the primary customer relationship directly or operate through another platform or interface, and when in-person interactions still add value. For non-retail finance organizations in particular, an individual delivering AI-augmented advice may offer more value to customers than any other sort of interaction.
2. **How to build trust.** Customers will learn to trust their agent (representing an agentic system) increasingly over time, as the agent uses its compounding knowledge to optimize outcomes. This could work for or against a firm, depending on whose agent the customer chooses to use and how frequently an external agent selects a single firm's offerings. Meanwhile, a corporate and investment banking player might extend services to a mid-market client base using a trusted digital investment advisory model. Financial technology (fintech) organizations may specialize in cross-industry ecosystem aggregation centred on the needs of a group of target customers, such as the newly retired, creating a digital concierge proposition that

optimizes those individuals' end-to-end financial and related non-financial goals. StockGro offers an example, detailed in case study 1, of how starting with solving an industry-wide customer problem can build trust.²⁰

3. **How to ensure visibility.** The number of agents and bots surpassed the number of people on the internet for the first time in 2024, accounting for 51% of total web traffic.²¹ Machine-driven traffic is still mostly one-sided and, all too often, driven by malicious intent rather than legitimate business use. However, it demonstrates how strategy must account for the terms by which an organization's services are discoverable, comparable and routable by other agents. There is a real risk that financial services companies will lose control of demand and margins.
4. **What to sell.** This includes, for example, where to seek opportunities and how to both bundle and set boundaries for new offerings. This could mean casting products as living services or deciding to provide value across a customer's life event, such as buying a car or moving to a new location. AI enables bundling and orchestration across ecosystems, shifting value from discrete products to outcomes; this changes proposition design, operating model, partners and metrics.

5. **Where to seek growth.** Organizations may use AI to access adjacent ecosystems and new markets, as the International Bank of Azerbaijan (ABB) demonstrates in case study 2; the bank has used AI to scale financial inclusion. Complementing these efforts, ABB is also harnessing conversational AI to simplify banking. Organizations are identifying and creating new value pools that were previously beyond imagination or simply unfeasible to pursue.
6. **How to monetize new AI-enabled services.** Pricing will need to evolve to align with models that capture continuous outcomes, automation value and variable usage (e.g. more subscription

or use-type fee schedules). For example, AI can update research and investment portfolio advice much faster than people can manually. Once matured and used by internal corporate bankers, this capability could be used by their clients' corporate treasury functions to optimize their capital management. The question is whether that move would serve as a brand differentiator (no fee) or as a new service element that attracts a fee. As another example, as agents and AI solutions become more valuable to customers over time (gaining insight into customers' preferences and behaviours as they gather data), how should organizations price their products to reflect this value compounding for customers?

CASE STUDY 1

StockGro Stoxo: using multi-agent AI to scale trust, clarity and conviction in retail investing

Retail investing in India suffers from a fundamental trust and access problem. Despite a population of 1.4 billion and over 470 million millennials with earning power, only around 45 million individuals actively invest in equities. Most retail investors rely on social media narratives, unverified "experts" or fragmented news, leading to guesswork rather than conviction. High-quality investment research remains expensive, complex and inaccessible to the mass market.

StockGro pivoted from a content-led investing platform to a research-first, AI-native model, embedding multi-agent AI directly into the investment decision journey. Drawing on its large, engaged user base and proprietary data, StockGro built Stoxo, an AI research platform where specialized AI agents work together to analyse markets, validate relevance and translate complex signals into clear, actionable intelligence. Key attributes of the solution include:

- A proprietary small language model (SLM) built in-house, powered by over 80 AI agents trained on inference, relevance and contextuality
- Training on proprietary conversational and behavioural data from 35 million users over five years, combined with expert trade calls and real-time market analytics

- Aggregation of community intelligence, expert insights and live market data that converts headlines into trade-ready guidance
- A radical degree of accessibility, pricing institutional-grade research 98% below traditional industry pricing

The solution, deployed at scale, delivers clear benefits including:

- A total of 2 million users onboarded within six months, serving over 25 million research queries per month
- Compression of more than four hours of manual research into under 10 seconds with high accuracy
- An 85% improvement in decision accuracy by reducing reliance on stale or unverified advice
- A total of 64% of users taking a subsequent trade action after receiving Stoxo insights, up from 43% with standard alerts

By redesigning investment research around multi-agent AI and proprietary data, StockGro has transformed AI from a generic assistant into a trusted, always-on research analyst for millions of retail investors.

CASE STUDY 2

The International Bank of Azerbaijan: using AI to scale financial inclusion

The ABB took an ethical approach to solve a problem for micro-business owners and entrepreneurs who struggled to secure finance due to their imperfect credit history and relatively high cost for banks manually underwriting micro loans.

Enabling responsible digital lending to largely non-banked customers requires automated, data-driven risk assessment capable of evaluating thin-file and first-time borrowers in real time.

A new segment for ABB, the bank built a machine learning (ML) application scorecard and embedded it directly into its digital lending platform to fully automate decisioning. The model enables real-time approval, limit assignment, and risk-based pricing with instant digital disbursement. Key components of the solution developed include:

- An ML-based application scorecard that uses internal transaction, credit bureau and other external data
- Cross-bank turnover data via automated PDF statement parsing

- Integrated real-time decisioning into the digital lending platform
- Automated approve/reject logic, limit allocation and risk-based pricing
- 100% end-to-end digital disbursement with zero manual intervention

As a result, the bank provided a valuable solution to customers that were previously precluded, which also provided a new segment of significant growth for the bank. The straight-through, standardized, data-driven risk model provided a workable, reliable solution for first-time borrowers that was scalable.

The solution generated a profitable revenue stream, with an average portfolio yield of 21% and non-performing loan (NPL) ratio under 4%.

In the future, the bank is looking to expand into other micro lending customer groups using data-driven learning to responsibly scale financial inclusion.

2.3 Measuring AI value

Not all AI initiatives deliver value, and the debate over AI's value will continue as applied capabilities mature. Early experiments, both failures and successes, are learning grounds, not rigid pass-or-fail tests. As such, a value-led approach to determining the return on investment (ROI) is essential. It may be that a pilot that was initially deemed a failure was badly defined, data products were skewed or the initiative was early in the adoption curve, where outputs lag but will compound. In addition to measuring outcomes, measuring inputs can be an important way to assess whether AI and AI agents are working effectively and serve as a stepping stone towards achieving business value.

Internally and across the industry, leaders are engaged in a robust, ongoing conversation about how to measure AI value. Several factors should be considered here, such as:

- **Economic value:** How much value does AI create, and what does it cost to run the technology at scale? Beyond cost reduction, leaders need to consider capital efficiency and revenue to keep AI tied to profit and loss (P&L). Example key performance indicators (KPIs) include AI-attributed revenue uplift; cost-to-serve reduction net of run costs; and capital efficiency impact, such as risk-weighted assets (RWA) intensity or economic capital per dollar of revenue.
- **Decision effectiveness:** Is AI improving judgement as well as speed? This measure might

include the rate at which humans override AI-supported decisions, the decision-loss rate when genAI is used compared to the baseline process or the degree to which AI-supported decisions diverge from established policies and standards.

- **Adoption and behavioural change:** Is the organization improving its operations as a result of AI integration? This measure might require evidence that workflows have changed enough for AI to operate alongside people, with humans providing oversight and handling exceptions. It might look for instances of people avoiding AI and, if so, why. Behavioural change could be evidenced by the share of cases completed using AI or the human intervention rate required to reach an acceptable outcome; the percentage of cases completed end-to-end with AI in the workflow; and active use of AI by role, coupled with those individuals' assessments of the tool's support.
- **Risk, control and sustainability:** What will it take to sustain AI over its life cycle at scale? To ensure that value is durable, an organization might measure the cost of control per model over time; the frequency of remediation events and the firm's ability to pause, roll back or retrain without business disruption; the model incident rate (weighted for severity) and the time to remediate incidents; and the control cost per model per month (including monitoring, validation and audits).

“ In addition to measuring outcomes, measuring inputs can be an important way to assess whether AI and AI agents are working effectively and serve as a stepping stone towards achieving business value.

2.4 Pilots versus scaled programmes

To integrate AI at scale effectively, companies are creating enterprise-wide foundations and reusable, sustainable capabilities. Early movers are conducting efforts to:

- **Create shared intelligence** through central AI platforms, developing the organization's enterprise AI platform. Doing so ensures access to shared models, orchestration and guardrails; persistent context, memory and feedback loops; and channel-agnostic decisioning that supports core goals and values.
- **Industrialize delivery** with high-throughput and repeatable pipelines. Doing so standardizes build (reducing the need for bespoke builds) and deploys and monitors AI life cycles. The organizations doing this are prioritizing use cases with clear service level agreements (SLAs).
- **Turn data into a compoundable asset** by restructuring it into owned, reusable products

that feed the enterprise AI platform. Doing this requires domain-based data ownership (customer, risk, transactions); continuous feedback loops from live interactions to improve data, real-time data accessibility and quality controls. Building the right data infrastructure is one of the greatest challenges identified by those who move to scaled solutions.

- **Ensure people remain at the helm** by redesigning roles and the workflows so that humans supervise AI systems as needed and handle exceptions. Clear accountability is also needed for AI-driven decisions, along with a commitment to providing people with increasingly personalized support to develop skills and advance their careers.

A recent collaboration between Allianz Partners and Taktile demonstrates how organizations are creating AI advantage at scale.

CASE STUDY 3

Allianz Partners and Taktile: using agentic AI to scale accuracy, trust and efficiency in global health claims

Processing health insurance claims on a global scale is inherently complex. Thousands of claims arrive daily in different languages, formats and levels of detail, each requiring a careful interpretation of medical invoices, coverage rules and eligibility criteria. This complexity creates friction for customers and significant operational cost for insurers. Allianz Partners identified health claims as a high-impact opportunity for AI, given that they are high in volume and variability, and require intense manual effort.

Allianz Partners pivoted from a predominantly manual approach to a human-in-the-loop (HITL), agentic AI model, partnering with Taktile to redesign the end-to-end claims workflow. Instead of treating automation as a single model, the team deployed multiple specialized AI agents, each responsible for a discrete step in the claims journey, with people overseeing exceptions and edge cases. Key attributes of the solution include:

- Sequential AI agents covering document classification, structured data extraction, treatment mapping and claim assessment
- Modular agent design, allowing capabilities to be reused and extended across products and geographies without rebuilding workflows

- HITL governance, ensuring full traceability, auditability and compliance through clear escalation and validation mechanisms
- Delivery on Taktile's decisioning platform, enabling collaboration between subject matter experts and technology teams

The solution, deployed at scale, delivers clear benefits, including:

- Cutting claims processing time from days to minutes, improving customer experience
- Reducing manual effort for insurance members and claims adjudicators by automating information extraction from documents and invoices
- Over 1.5 times better fraud and false payout detection through consistent application of advanced decision logic

By redesigning claims processing around agentic AI with strong human oversight, Allianz Partners has established a scalable foundation for AI-driven operations. The solution supports growth while preserving trust, transparency and regulatory confidence.

3

Successful AI starts with a strong foundation

Scaling AI sustainably requires deliberate focus across enterprise platforms, workforce capability, risk management and governance.



Many firms experience a lag between investing in AI and generating scaled results or returns. Common reasons for this are:

- A lack of data and technical readiness
- Insufficient agility to exploit, adopt and adapt to the technical possibilities of AI
- Concerns about data security, model output and third-party risk
- Underprepared AI governance, risk management and compliance across all levels

Financial services organizations are investing in four foundational areas to address the challenges of scaling AI. While each organization's path differs

based on its strategic priorities, these areas are consistently critical across the industry:

- Building resilient, scalable and adaptive AI platforms underpinned by secure, reliable and available data
- Upskilling and empowering the workforce to drive adoption and continuous improvement that reinvents work tasks rather than automating what happens today
- Integrating AI-specific risk management capabilities into the enterprise risk management spine
- Moving beyond broad principles to embed responsible AI firmly into appropriate governance structures

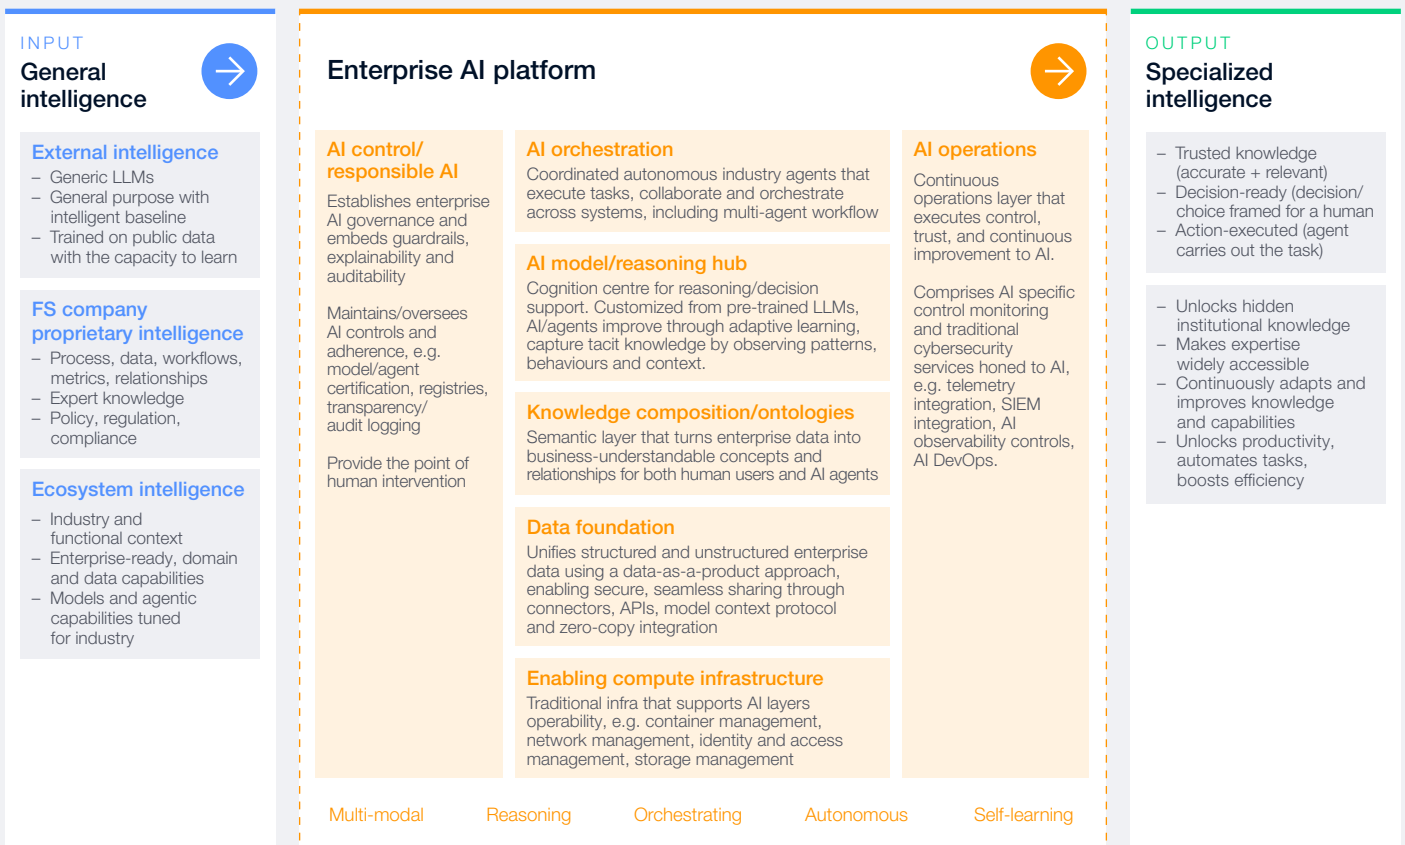
3.1 The foundation and platform architecture

After early experimentation with proofs of concept (POCs), most financial services firms seek to scale AI through an integrated, enterprise-wide intelligence platform. This platform emulates key aspects of human cognition; it can perceive, reason, act and learn across multiple modalities in

dynamic environments. Crucially, it does this in a secure, scalable and cost-effective way.

Such platforms typically share a common architecture and include common elements, as shown in Figure 5.

FIGURE 5 Components and architecture of an enterprise-wide intelligence platform



Source: Accenture synthesis of AI platform desk research (MS 2024, AWS 2025, Google Cloud, MAS, FSB), client experience and SME interviews.

“ Data products ensure consistent discovery and consumption, which allows agentic systems to reason and act with greater accuracy.

Any such platform is only as good as the data that underpins it. The intelligence platform can consolidate structured and unstructured enterprise/ecosystem data and enables secure sharing via connectors, application programming interfaces (APIs), Model Context Protocol (MCP) and zero-copy integration.

To prepare data for AI and human consumption, the platform needs to structure data into products: reusable, business-ready datasets that are packaged and managed with a defined purpose, ownership and life cycle. Data products ensure consistent discovery and consumption, which allows agentic systems to reason and act with greater accuracy. In turn, organizations provide semantic frameworks, such as business concepts, relationships and context, so that systems can interpret data correctly.

Buy, grow, partner

Most firms are using a mix of build, buy and partner to scale AI. Hyperscalers, other cloud data vendors and model providers bring powerful capabilities that evolve rapidly. This is evident in a recent global survey which revealed that 63% of financial services firms and 65% of regulators use external foundation models for internal workflows.²² Many of these organizations also customize or develop some AI systems in-house.

The challenge for leaders is deciding which parts of the architecture to build in-house and which to buy or partner for. The goal is to retain intellectual property (IP) and agility while avoiding over-dependence on any single external provider through long contracts or proprietary systems. Additionally, in some markets, sovereignty considerations are becoming a key factor for governments, regulators and financial institutions in their build-versus-buy-versus-partner decisions.

3.2 Skills, hiring and culture change

Companies throughout the sector are also enhancing human and organizational foundations they need to scale AI effectively. This matters because genAI and agentic AI will likely affect a larger share of tasks than in many other industries. Much financial services work is language-based and depends on unstructured information, cross-team coordination and multiple verification layers. For example, in banking, it's estimated that 73% of employees' time could be affected by genAI, with 39% through automation and 34% through augmentation.²³

The predominant emerging vision is to manage human and digital workers as an integrated talent ecosystem aligned with purpose, data and performance outcomes. Figure 6 shows the basic building blocks of such a system. Automating more execution through agents should allow humans to focus increasingly on applying judgement and oversight and creating value. For humans to do this well, they need clear, hybrid workforce models that define skills, decision rights and ways of working.

FIGURE 6 Building blocks of a hybrid workforce



Source: Accenture.



To fulfil this vision, work needs to be completely redesigned. While automating existing processes offers short-term value, that approach loses potency as people learn more about what the technology can do and as the technology, in turn, learns from the people who work with it.

This entails baselining processes and KPIs, and some teams are creating digital twins to imagine and model radically different scenarios enabled by AI. Today, skilled, innovative professionals do this work. Some organizations, however, are experimenting with AI agents to help those individuals reinvent their workflows and processes.

As a result, organizations may need to shift to a skills-based workforce architecture. Starting with individual skills and tasks, they will reconfigure people's roles, supported by AI, and, in turn, drive new talent, skill and capacity requirements. They will also adjust learning paths, moving much of

the training into the flow of work itself to address evolving skill needs for both people and agents.

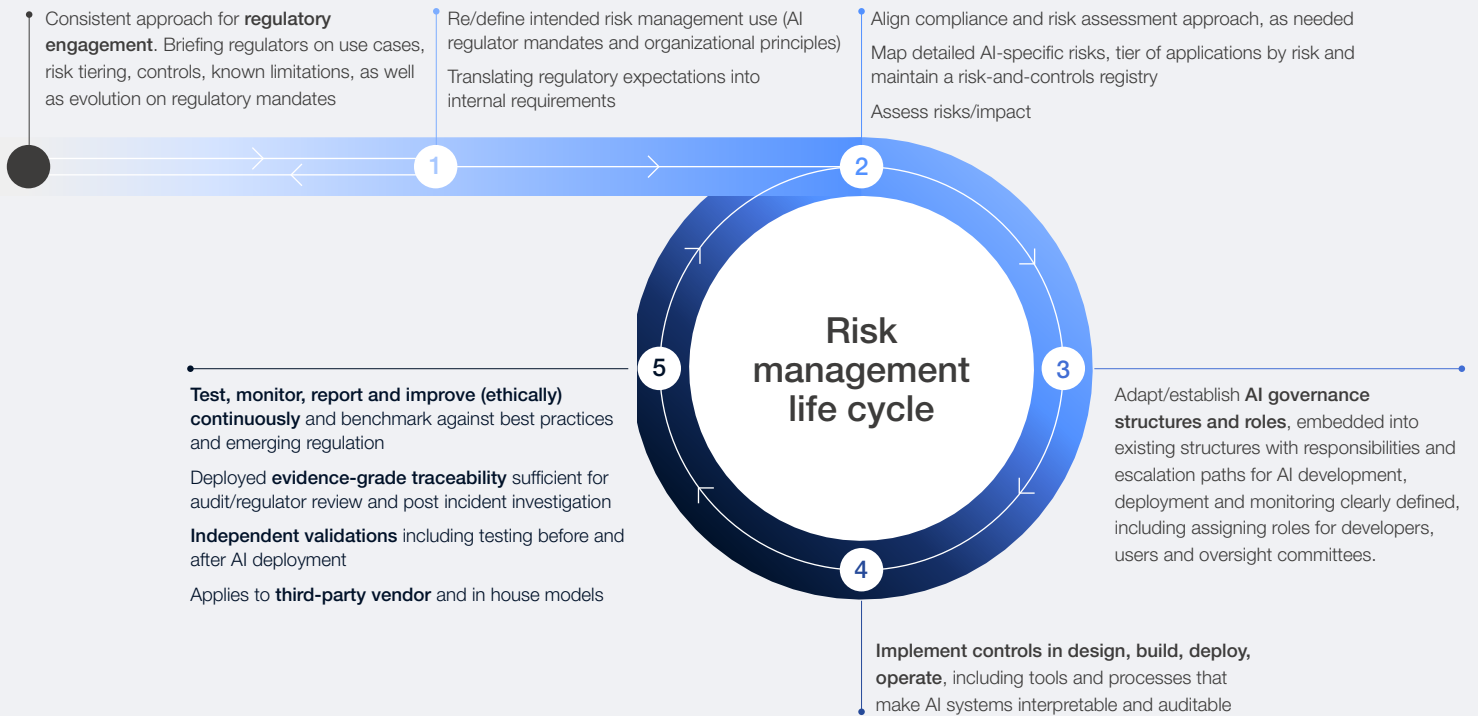
While 95% of employees say they see value in working with genAI, adoption still depends on trust, learning and experience. Only 19% say they work in teams that experiment with AI together, and only 17% feel comfortable speaking up or sharing ideas.²⁴ The World Economic Forum estimates that 59% of the global workforce will require reskilling by 2030, with 170 million jobs created and 92 million displaced.²⁵ It will be crucial for financial services firms to focus on improving culture and on better understanding and supporting people's needs as they increasingly work with AI. They will need to encourage safe experimentation, demonstrate prudent management of AI-related risks, and reward effective adoption. To that end, some firms have already revisited their employee value propositions to reflect this, in light of a future that will likely include managing a team of digital co-workers and rewarding them for joint outcomes.

3.3 Governance, risk and compliance management

As highly regulated entities, financial services organizations are enhancing AI governance, risk and compliance foundations to ensure AI is used as intended, and to maintain financial system stability. This requires adapting enterprise-wide risk management frameworks for AI while aligning with evolving regulatory standards.

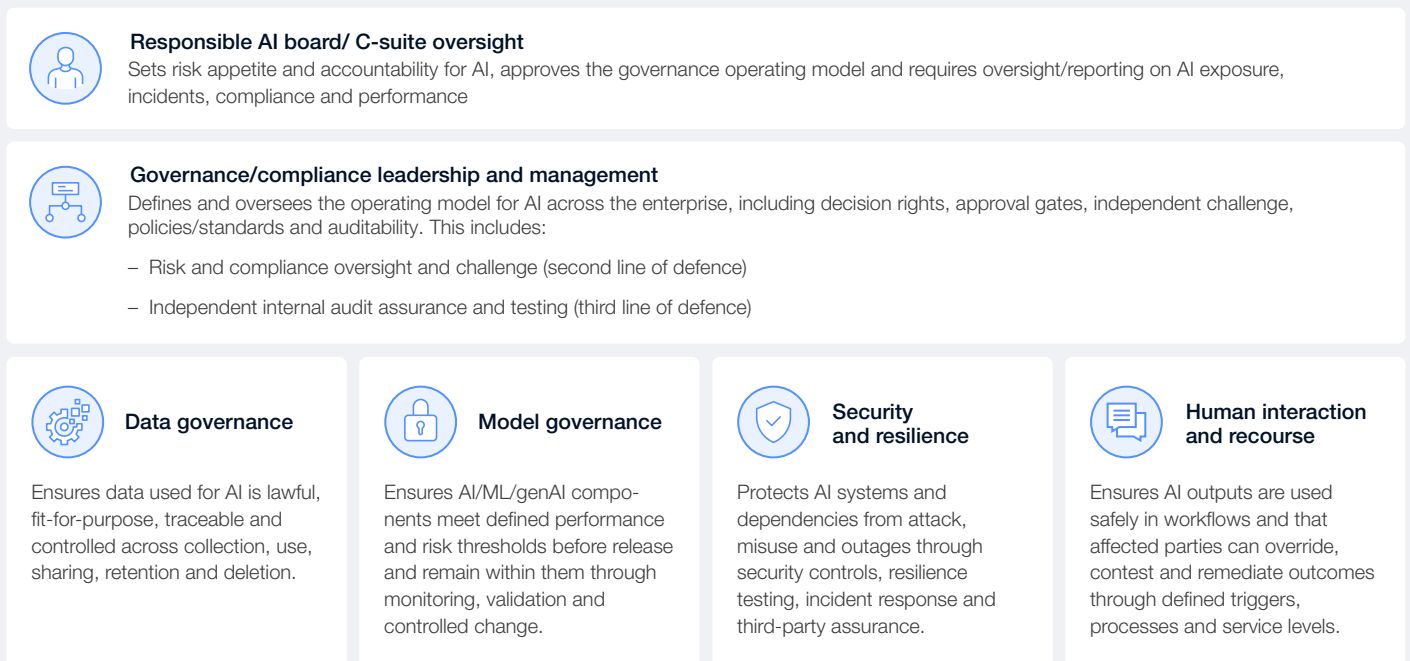
Across jurisdictions, industry risk frameworks aim to turn broad AI principles (fairness, transparency, privacy, robustness, accountability) into operational, embedded controls. Figure 7 explains the common life cycle these frameworks follow.

FIGURE 7 | Adapting the risk management life cycle to AI



Risk frameworks also converge on common control/governance domains, as shown in Figure 8.

FIGURE 8 | Integrated risk management – control domain taxonomy for AI governance



Risk frameworks differ mainly by their level of prescription and their primary risk lens. Standards like the International Organization for Standardization's (ISO) [ISO/IEC 42001](#) emphasize a certifiable management system and a policy-to-audit discipline. They work best when an organization needs repeatability and assurance across many use cases. The National Institute of Standards and Technology's (NIST) [AI Risk Management Framework](#) (AI RMF), a practical risk management handbook, is designed to map to existing enterprise risk and security practices. Prudential model-risk regimes, such as the Federal Reserve's [SR 26-2](#)²⁶ and the Bank of England's Prudential Regulation Authority's (PRA) supervisory statement (SS) [1/23 – Model risk management principles for banks](#), are narrower but deeper. These prioritize independent challenge, validation and evidence for models used in regulated decisions, such as credit, capital, pricing and fraud.

The EU's [Artificial Intelligence Act](#) is compliance-by-risk-tier and introduces explicit obligations

including documentation and governance duties that apply beyond classic models to broader AI systems. Legal obligations, for example, include clear ownership of AI cycles across the full lifecycle, defined escalation paths for AI-related risks and incidents, and senior management accountability for high-risk AI systems. APAC frameworks, such as the Monetary Authority of Singapore's [FEAT principles](#), tend to be sector-pragmatic and execution-oriented, focusing on customer-impact outcomes and supervisory expectations rather than a single cross-economy statute. See Appendix A for a more detailed overview of risk framework adoption by region/countries.

In many jurisdictions, regulatory positions are still emerging. Financial services players are sharing their experiences with regulatory and industry working groups to inform policy. DANA offers an example of this approach, as outlined in case study 4. Fintech has proactively contributed its know-how to Indonesia's national AI roadmap.

CASE STUDY 4

DANA: meaningful participation in creating a national AI roadmap for Indonesia

DANA engaged early with the Ministry of Communications and Digital Affairs to contribute its insights into Indonesia's national AI roadmap, balancing the needs of digital financial services and national security and stability objectives.

As Indonesia embarked on defining its national AI roadmap, digital financial services providers were initially underrepresented in the debate. DANA proactively petitioned to contribute valuable input into the framework. It submitted an assessment of its existing internal policies for implementing and using AI, as well as showcasing how it instilled a strong responsible AI culture via its "AI Everywhere" initiative.

DANA's contributions positively impacted the AI industry agenda:

- Proactive alignment with anticipated AI regulation reduced policy-adjustment lead time by an estimated 30–40%, avoided potential remediation costs and accelerated compliant AI deployment (with an estimated benefit of \$500,000 to \$2 million in annualized economic value/opportunity gains)
- Enhanced the financial services sector narrative, ensuring that payment system risk protection, consumer protection and data governance considerations were embedded early
- Demonstrated that early, structured industry engagement increases policy relevance and reduces cross-sector regulatory friction, particularly between digital governance and financial sector supervision

Practically, the most scalable path to implementing responsible AI risk management and governance is three-tiered:

- 1 **Run** a single internal control taxonomy – a cross-enterprise baseline – to standardize roles, controls and assurance across all AI. ISO/IEC 42001 and NIST's AI RMF frameworks are often used for this purpose.
- 2 **Map** it to model risk requirements and use a highly stringent model risk management (MRM) framework, such as SR 26-2 or PRA SS1/23, for high-impact predictive decision models.

- 3 **Align** with jurisdiction-specific compliance modules and legal obligations, such as the EU Artificial Intelligence Act and local prudential requirements.

The baseline covers end-to-end AI system risk; the MRM framework provides regulator-grade challenge for decisioning models; the modules ensure compliance while preventing over-built controls where they are not required.

3.4 Security, model risk and threat landscape

“Data privacy and protection ranks as the leading AI-related risk in financial services, cited by 74% of industry players and 80% of regulators recently surveyed.”

AI – particularly genAI and agentic AI – elevates risk across many existing organizational risk categories and requires an integrated risk management approach. See Appendix B for a detailed breakdown of AI-driven risks and corresponding management responses. Financial services organizations commonly report three risk-oriented barriers to scaling AI within their organizations: data security and privacy, unreliable model output, and third-party vendor risk. In response, firms are strengthening their foundational risk management capabilities.

Data security and privacy

Data privacy and protection ranks as the leading AI-related risk in financial services, cited by 74% of industry players and 80% of regulators recently surveyed.²⁷ Gen AI models rely on large training sets, which may include externally sourced data. Training may involve the AI model ingesting sensitive or proprietary data, and outputs may reveal or infer personal information that should have been protected. In turn, these AI-related security vulnerabilities may also increase the likelihood of attacks that compromise confidential data. Additionally, model providers may retain AI prompts without permission, either inadvertently or to improve model performance.

Organizations can take a number of steps to mitigate these risks, including:

- Adding or enhancing data filters to remove or obscure sensitive data
- Deploying access-controlled architecture that segregates client data from model providers
- Conducting penetration and adversarial testing to identify and remediate vulnerabilities
- Conducting data protection impact assessments when using sensitive data
- Restricting user and agent access and conducting regular audits
- Deploying security disciplines for AI using DevOps

Model risk

GenAI models rely on massive datasets, making consistent accuracy checks challenging. Output can be unreliable or biased, and traditional deterministic testing methods are often less effective. In response, organizations are adopting multiple mitigation approaches for model risk. Human review is the simplest option, but it can be expensive and may introduce new risks, such as human error. Other mitigations include:

- Deploying multi-agent systems to review the accuracy of outputs and verify them against defined standards and other reliable data sources
- Using SLMs and fine-tuning using proprietary datasets to increase output accuracy
- Scoring outputs based on confidence, so that people can prioritize review and/or multi-agent systems can check low-confidence results
- Automating attributions so that outputs can be rapidly explained and assessed
- Monitoring operational performance to verify model outputs against a specified range of measures/criteria. This approach is particularly important when introducing new and upgraded model versions.

Third-party vendor threats

Most financial services organizations use third-party vendors for genAI models and solutions; doing so is more cost-effective than developing the needed internal expertise and capacity. However, standard audit processes often cannot trace vendor data and algorithms end-to-end. Vendors may also update software without consultation.

Teams typically access, customize and integrate genAI models in different ways, each raising risks. To mitigate these risks, organizations are adding third-party gateways, checkpoints and firewalls between different AI systems. They are also enabling independent testing and setting more stringent auditing arrangements.

4

The role of agentic systems in unlocking value at scale

As agentic AI is adopted across financial services, its value lies as much in collaboration with humans as in what it can automate.



Agentic AI is a genuine breakthrough for handling complex end-to-end processes and unlocking value. What makes it different is its ability to act with a degree of agency. Instead of stopping at a single answer, an agentic system taps a variety of individual agents to move work forward. An agentic system can request information, trigger downstream steps, escalate exceptions and complete complex jobs. It also evolves rapidly, learning and improving as it builds on existing predictive AI and automation solutions and as people provide feedback.

Financial institutions are adopting it rapidly for good reason. Organizations need traditional automation for repeatable processes, rules engines for deterministic policy decisions, ML for prediction and pattern detection, and genAI to contextualize, synthesize and generate outputs. Agentic AI can help organizations connect these steps more effectively.

The challenge is scaling it. This includes identifying where agentic AI is appropriate, where it should remain assistive only, and where it should not be used without strict human approval.

4.1 Why agentic AI is key

With its ability to plan, decide, act and learn autonomously to achieve defined goals, rather than simply respond to prompts or execute fixed rules, agentic AI essentially becomes an active participant in work, with several distinct capabilities:

- **Goal-directed autonomy:** agents initiate and complete tasks end-to-end within defined objectives.
- **Reasoning and decision-making:** they evaluate options, make trade-offs and select actions.
- **Continuous learning loops:** performance improves through feedback and outcomes over time.
- **Multi-step execution:** agents break work into tasks, coordinating sequences of actions across multiple agents, systems and processes.

- **Context awareness:** they adapt behaviour based on changing data, environment and intent.

Together, these capabilities represent a leap forward in how work is done, creating a compounding effect in which speed, quality and learning reinforce each other over time.

In practice, agentic AI works as a coordinated ecosystem of specialized agents. Some focus on analysis, others on execution, monitoring or optimization. Together, they form a clear hierarchy of responsibility and control.

Orchestrator agents set goals, break down tasks and manage workflows. Other agents execute specific tasks, such as synthesizing inputs, understanding intents, making recommendations and executing actions. People evaluate responses and make decisions at critical points in the workflow, as needed.

4.2 Human–AI partnership

Most commentary on digital and human workers focuses on delegation, risk management and supervision in the human–AI relationship. For example, accuracy-critical tasks often require maintaining a HITL workflow, in which humans are actively involved at each step or at regular intervals. An increasingly common variation of this is an autonomous AI system with a human overseeing the process and offering feedback or vetoing actions where required. However, with significant advancements in automation, a growing concern is whether humans can effectively supervise multiple agents. Relatedly, institutions must ensure the ability to intervene, override, or shut down AI systems where necessary.

While these are critical considerations, they tell only part of the story. Equally important is how agentic AI can unlock human potential and shape satisfying,

productive experiences and outcomes for workers. For many roles, it can be a true collaboration, where humans and AI exchange ideas, co-create solutions and push each other beyond individual limits. Successful collaboration in agent-human teams could offer:

- **A shared purpose.** Humans and AI are aligned around the same goals, and AI can generate ideas, suggestions or recommended actions towards an objective. Collaboration is outcome-driven, not always task-driven. Examples include a claims leakage target or fraud loss target.
- **Iterative co-learning.** Humans learn from AI patterns, predictions or insights. AI adapts to human feedback, preferences and decisions. Both improve through repeated interaction, creating a dynamic learning loop in the flow of

work. For example, this could happen in a trade anomaly review or when fine-tuning credit limits.

- **Complementary creativity.** AI generates options, scenarios or simulations that humans might not imagine. Humans integrate judgement, intuition and context. Together, they co-create solutions that neither could produce alone, such as creating hardship script variants or creating a market campaign.
- **Fluid, real-time interaction.** Humans and AI exchange inputs continuously rather than sequentially. Decisions and ideas emerge organically from the partnership. The process feels like a conversation or a collaborative workshop, not a handoff. Examples include resolving customer disputes or quoting.
- **Mutual influence.** AI can inform or influence human thinking by surfacing patterns or possibilities. Humans shape AI behaviour

through feedback, guiding exploration. The partnership is bidirectional. An example is the AI noticing a suspicious pattern surfacing and the adjuster asking for more details.

To collaborate effectively, people must be comfortable using AI and committed to adopting it. They also need to continuously increase their knowledge of the tools they are using or guiding. They need to know what they can ask of their new tools and how to learn from and teach those tools in the flow of work. Moreover, they must feel that they can do so safely and securely.

Many organizations provide training and access to AI. Case study 5, on the Kasikorn Business-Technology Group (KBTG), shows how a democratized, responsible approach to upskilling the workforce and creating space to experiment effectively can support a collaborative workforce model.

CASE STUDY 5

Kasikorn Business-Technology Group: embracing a human-first AI collaborative culture with strong results

KBTG's transformation intentionally adopted a human-first approach to AI. Within the first year, the organization trained its workforce on the AI tools being integrated into daily work. Leadership also democratized access to AI and provided a safe experimentation environment, where employees could explore ideas and build their own agents.

Key features of the approach include:

- An enterprise AI council was established as a formal AI governance and strategy body to guide, govern and scale AI adoption across KBTG in support of Kasikornbank, with a strong emphasis on responsible, human-first AI
- 100% AI literacy as a first-year initiative, focusing on upskilling all employees in the fundamentals of AI and productivity AI, while empowering staff to use AI responsibly
- A dedicated AI playground environment to support experimentation, solution creation and best practice sharing – this led to the emergence of early adopters who later developed AI agents within their teams
- Employee-driven innovation, enabling staff to build personal AI agents and minimal viable product (MVP) projects to solve their own challenges and those of their teams

- Selective scaling of high-impact use cases into enterprise-grade platforms with built-in quality control, security and responsible AI governance

As a result, employees embraced AI and achieved measurable outcomes demonstrating effective human AI collaboration:

- Over 200 innovation ideas generated, leading to 60 MVP solutions, with eight high-impact use cases scaled and deployed as enterprise platforms
- 20–59% productivity improvements for administrative tasks, documentation and preliminary analysis
- Approximately 30,000 workdays saved through operational improvements after prototype AI agents were scaled into production

KBTG continues to scale enterprise AI by building a centralized, vendor-agnostic AI platform, avoiding dependency on any single technology stack. This platform enables human–AI collaboration where employees and AI agents work together across workflows to transform productivity and innovation at enterprise scale.



4.3 Examples of agentic AI in financial services

Agentic capability is not a single construct. It spans a spectrum of autonomy, from assisting humans to independently planning and executing outcomes. It is important to ensure clarity around how decisions are made, how clearly responsibilities are divided and articulated for both people and machines, and how actions are triggered. Equally important is how much human oversight is needed and expected, and how people can escalate a situation when they spot an exception or another problem surfaces.

- **Assistive agents** support human decision-making by generating suggestions, drafts and recommendations, while the human remains accountable for final decisions and actions.
- **Semi-autonomous agents** execute defined tasks within clear guardrails, operating independently for routine activities but escalating exceptions or edge cases for human approval.
- **Autonomous agents** plan, decide and act against defined goals with minimal human intervention, managing end-to-end task execution within agreed boundaries.
- **Multi-agent systems or orchestrator agents** coordinate multiple specialized agents, each focused on one task, such as planning, execution or review. These agents collaborate to solve more complex, multi-step problems.
- **Tool-using agents** extend beyond reasoning by interacting with external systems, APIs and data sources to take real-world actions and complete tasks.

Figure 9 provides an example of agent use across the value chain for banking.

FIGURE 9 | Agent use across the banking value chain



Together, these agent types can provide a scalable model for increasing autonomy, enabling organizations to progressively shift to machine-orchestrated outcomes while maintaining control and accountability.

Agentic AI is still a nascent technology. Even though it has moved quickly to capture leaders' attention, organizations have only relatively recently deployed it to drive operational efficiency, enhance decision-making and deliver personalized client outcomes for financial services. Agents are used to enhance **customer engagement and advisory**, for example, where they enable proactive financial guidance, policy recommendations, personalized pricing and targeted product suggestions. In **operations**, agents autonomously handle loan origination, claims processing, transaction reconciliation and onboarding. In **decision and risk intelligence**, they provide portfolio optimization, trading signals, credit scoring, stress testing and underwriting recommendations, augmenting human judgement with scalable insights.

At a **strategic level**, agentic AI supports decision-making and anticipates credit defaults, claim surges and liquidity risks.

As agentic use cases emerge, examples that reinvent the customer experience and enhance operational efficiency and accuracy illuminate what's already possible.

Customer engagement and advisory

A key feature of agentic-enabled workflows is that they can deliver services at scale that firms previously offered only to private clients at a high cost. Agentic AI overcomes this constraint, and early forms of digital personal (financial) assistants act with customers' interests in mind. Candidly's experience, explored in the adjacent case study, demonstrates this concretely, illustrating how an organization can deliver services affordably to help individuals address a complex life stage issue.

CASE STUDY 6

Candidly: scaling personalized financial guidance with a multi-agentic AI assistant

Approximately 42.7 million Americans hold \$1.84 trillion in student debt²⁸ and must navigate over 140 federal programmes,²⁹ many with frequently changing repayment requirements. Without personalized guidance, borrowers make costly or no decisions, leaving billions in relief unclaimed. As a result, 25% face delinquency or default,³⁰ 42% sacrifice basic needs³¹ and lingering debt cuts retirement savings by 30%.³²

Human coaching can't help students at scale in an affordable way. Generalized genAI partially fills the gap but also introduces new risk. Frontier large language models (LLMs) fail roughly 60% of realistic financial analyst tasks, and 43% of their answers to student loan questions are incorrect or misleading.³³

To address the issue, Candidly built Cait, a domain-specialized, multi-agentic AI assistant. Cait encodes expert-level financial reasoning as auditable, human-authored knowledge modules that ensure accuracy, compliance and deterministic calculations. The solution comprises:

- The Candidly Intelligence Center (CIC), a five-layer intelligence stack combining LLM reasoning, API tools, human-authored skills, deterministic code execution and retrieval-augmented generation
- Specialized agents for student loan optimization, college planning, readiness for retirement and employee benefits – with cross-agent handoffs preserving full conversation state

- Configurable guardrails: input validation, compliance screening and complete decision provenance for every recommendation
- Launch across enterprise partners with white-label, co-branded and API deployment options, including zero-data-transmission partner-cloud instances

A large majority of Cait's users (71%) take a meaningful step following its insight; almost half (49%) of conversations reach the action stage. Almost a third (31%) of conversations happen outside business hours. Meanwhile, 41% of conversations continue for more than six messages and 61% of returning users explore new financial topics.³⁴

Candidly's approach triples the rate at which users complete their first meaningful action on the platform, validating that domain expertise, not model scale alone, drives accuracy in regulated domains.³⁵

Candidly is expanding Cait into retirement planning, consumer debt and equity plan guidance – deploying a new specialized agent roughly every three weeks. The architecture supports MCP, enabling partners to integrate Candidly's financial intelligence directly into their own AI assistants.

Operational efficiency and accuracy

Efficiency remains a top priority for AI reinvention and financial services organizations are increasingly using agentic AI to reinvent operations processes. Case study 7 shows how Mastercard improved a slow, manual process, not only increasing efficiency and quality but also delivering better outcomes for merchants and their customers.

CASE STUDY 7

Mastercard Ethoca Consumer Clarity: using genAI to scale trust and transparency in digital payments

Unclear transaction descriptions are a primary driver of consumer confusion, disputes and chargebacks in electronic payments. These create unnecessary friction and cost. With hundreds of millions of transactions flowing through its network daily, Mastercard saw an opportunity to turn scale into continuous feedback loops that drive smarter decisions, faster optimization and compounding value over time.

Historically, enriching raw transaction data required manual cleansing and attribute matching by specialists, creating a bottleneck that constrained global expansion, increased operational costs and limited coverage despite growing transaction volumes.

Mastercard pivoted from a human-led enrichment model to a collaborative model of multiple AI agents with people in the lead. The team rebuilt the workflow, so that genAI now performs primary attribute extraction, with people overseeing exceptions. Key attributes of the solution include:

- Multiple sources of customer transaction data across Mastercard's data warehouse, Snowflake and Microsoft Azure to create a "merchant database" that allows the AI engine to enrich customer transactions
- Designing and deploying all ranges of AI from classical ML, such as XGBoost, to modern genAI or LLMs into one product to encourage both revenue growth while meeting customer satisfaction

- Embedding Mastercard's AI governance process of efficacy, fairness and transparency through results explainability with a scoring system, tracking regional variations, manual validation and AI judging agents
- Delivering with agility through Microsoft's AI product suite such as Copilot, Microsoft Foundry and Azure Machine Learning

Deployed at scale, the system provides clear benefits, including:

- 15% higher data accuracy than traditional benchmarks, improving merchant data quality
- 20% more records enriched from the same datasets, rapidly scaling coverage expansion
- 87% reduction in processing time and 92% lower cost per record, improving unit economics

Now, teams manage genAI as a continuous capability and performance depends on high-quality structured inputs, continuous feedback loops and systematic retraining. Confidence scoring and human review are critical to maintaining trust.

Next, Mastercard will embed this capability directly into clearing data, enabling the company to deliver enriched transaction details to all banks in its network.



Strategic decision-making

As an example of how AI supports broader strategic decision-making at scale, consider The European Bank for Reconstruction and Development's (EBRD) case. The EBRD used agentic AI to make its extensive evaluation knowledge more readily available, help shape stronger projects and deliver real benefits to communities.

Agentic AI is already demonstrating meaningful impact across workflows, yet it remains nascent, with capabilities advancing rapidly. Agents will continue to evolve from executing narrow tasks to complex, end-to-end workflows, coordinating across systems and adapting in near real time.

CASE STUDY 8

EBRD LessonsBot: using genAI to scale institutional learning, trust and evidence-based decision-making

For multilateral development banks (MDBs), institutional learning is critical to improving project outcomes and maximizing impact across countries and regions. Since 1991, the EBRD has accumulated extensive evaluation knowledge that captures what works, what doesn't and why. However, as collaboration across the MDB ecosystem increased, so did the challenge of making this knowledge easily accessible and actionable at scale.

Historically, the EBRD's evaluation insights were spread across more than 300 technical documents and over 20,000 pages, stored in multiple repositories. Locating relevant lessons often required hours or days of manual searching and interpretation, limiting reuse, slowing decision-making, and reducing the effectiveness of institutional learning – particularly for new countries of operation and cross-institution collaboration.

EBRD pivoted from a document-centric knowledge model to a genAI-enabled, conversational access model, developing LessonsBot as its first bespoke genAI solution. Built by the AI Centre of Excellence in close collaboration with the internal evaluations department, LessonsBot enables staff to query authoritative evaluation evidence using natural language, with built-in citations and governance. Key attributes of the solution include:

- A cloud-based, retrieval-augmented generation (RAG)-enabled chatbot, grounded exclusively in the EBRD's official evaluation documents to ensure accuracy and trust

- Consolidation of over 300 evaluation reports into a single AI-powered knowledge base spanning more than 20,000 pages
- Use of GPT-4o and Microsoft Foundry, enabling secure deployment, explainability and scalable access
- Iterative, agile development with evaluators, ensuring usability, relevance and continuous improvement
- Phased rollout with strong governance, expanding access from pilot users to all staff, the board and evaluators across 10 partner international financial institutions (IFIs)

Deployed at scale, the system delivers clear benefits, including:


- At least 20% time savings for evaluation managers and analysts
- A 66% increase in internal lessons-learned reports
- 100% of surveyed users reporting increased productivity, with 75% citing reduced mental effort and improved insight quality
- Adoption across EBRD and 10 other IFIs

Now, teams manage genAI as a shared institutional capability rather than a one-off tool. Performance depends on high-quality structured inputs, citation-based retrieval, continuous feedback loops and human oversight to maintain trust.

5

Planning for the industry's future

Realizing AI's full potential in the long term demands forward-looking frameworks and leadership commitment.



The future of financial services will be connected, intelligent and ecosystem-driven. As this unfolds, leadership across the C-suite will play a critical role in embedding AI more deeply into their organizations. Simultaneously, ecosystem players will need to stay ahead of emerging technologies as they continue to evolve.

5.1 Emerging technologies and cross-industry innovation

“ The massive quantum-enabled increase in compute power promises to further supercharge the development of AI solutions, creating significant gains.

The financial sector is continually adopting emerging technologies that are reshaping how institutions operate. Cloud platforms enable scalable infrastructure and data management. Open banking allows secure integration with fintechs and third-party services. Blockchain and digital assets facilitate faster, transparent and programmable transactions. Advanced cybersecurity protects digital channels while enabling trust in increasingly automated systems.

Organizations can deploy AI as an orchestrator across this emerging technology ecosystem and catalyse a reinforcing cycle of value. AI amplifies the value of investments in cloud, open banking, blockchain and cybersecurity. Cloud infrastructure enables scalable AI models and data pipelines, while open banking APIs feed AI systems with rich, integrated datasets for personalized insights and decision-making. Blockchain and tokenized assets provide structured, auditable data that AI can analyse for fraud detection, risk assessment and optimization of financial flows. Meanwhile, though advanced AI models can now identify and exploit vulnerabilities in software, these same models can, and must, be used to enhance cybersecurity and help drive resilience across digital channels.

As other technologies progress, such as quantum computing, further transformation of financial services will take place. Quantum has recently been receiving increasing attention, primarily as a security concern. Once quantum computers sufficiently advance, they hold the potential to break the encryption standards that currently protect financial data and transactions, posing a significant future risk to the security of the global financial system.³⁶ In response, bodies such as NIST have been working to get ahead of the threat. Through a multi-year process involving industry, academia and governments, NIST released its first post-quantum cryptography standards in 2024, and provides resources and guidance to help financial institutions begin migrating their systems to quantum-resistant cryptography.³⁷

Beyond the security threat, quantum also has a transformative upside. The massive quantum-enabled increase in compute power promises to further supercharge the development of AI solutions, creating significant gains. In particular, the technology promises to improve model optimization, reducing training time and improving the quality of outputs. It may also exponentially accelerate certain core subroutines used by ML, which will enable scaling in domains with vast and complex state spaces, such as financial analysis.

5.2 Framework for scaled and sustained AI adoption

Leaders need a course of action based on factors including their organization's AI maturity, ambition and financial resources. Intent matters, and there is a clear difference between using AI and truly transforming with it. Transformational AI requires a strong commitment from the CEO, recognizing that

AI can reshape the operating model, encouraging experimentation to find practical uses, and completely redesigning workflows rather than limited integration.

Leaders typically guide their organizations through several transitional stages, illustrated in Figure 10.


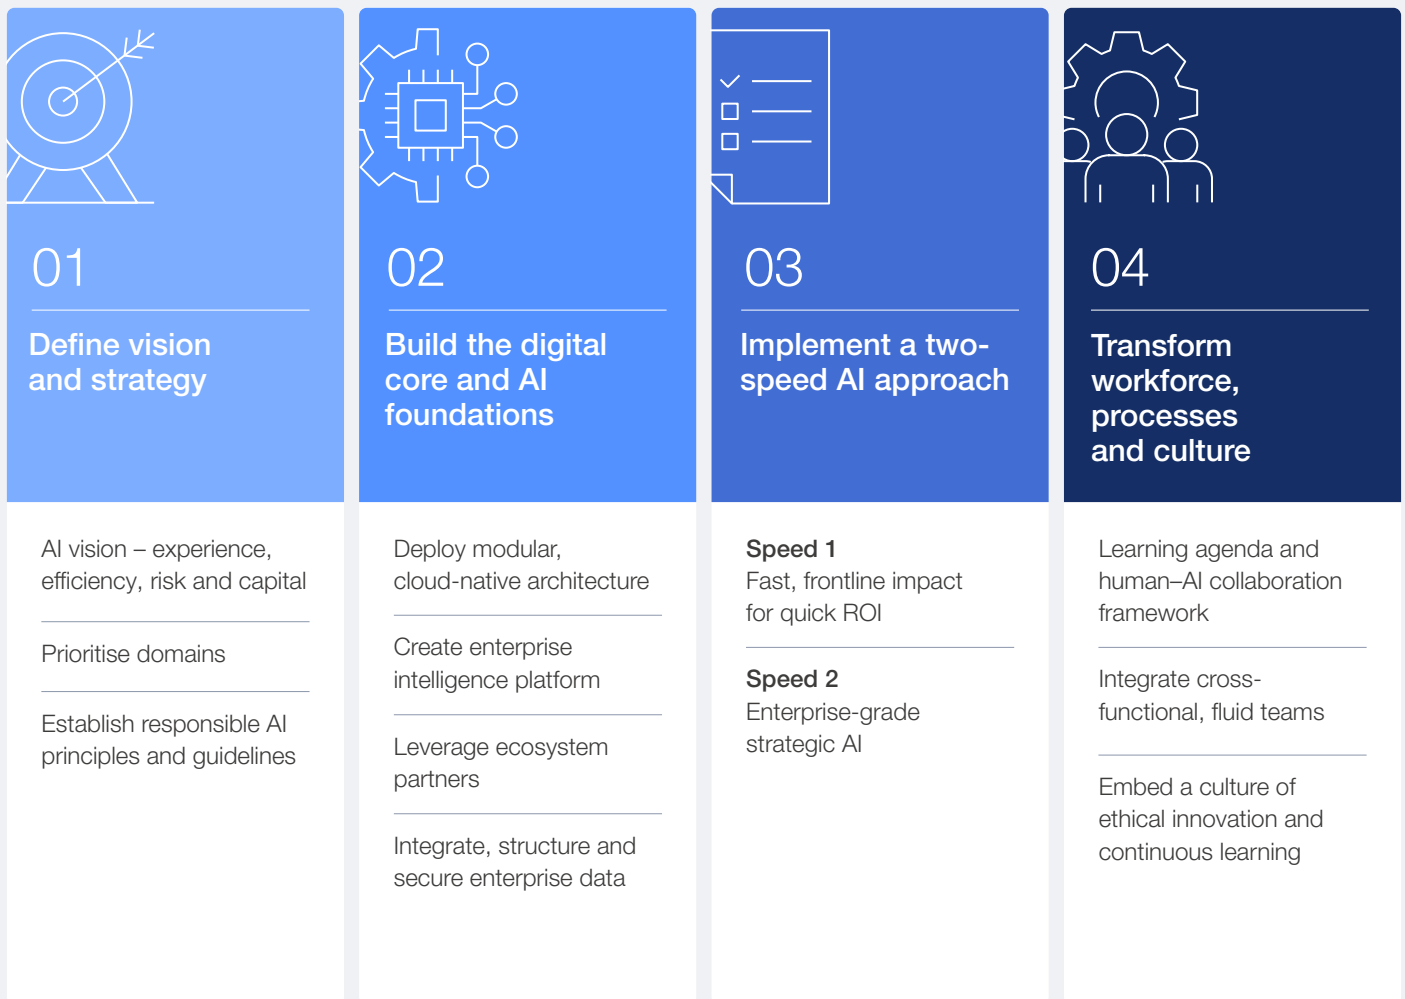


FIGURE 10 | Framework for transformational AI



1 Define vision and strategy

This stage sets direction and prevents fragmented AI adoption. It defines where AI will create value, which workflows matter most, and where autonomous action must stop. During this stage, organizations establish a baseline, clarify the relationship driver, assess data and model readiness, and identify the highest-value agentic opportunities. The goal is board-level alignment on priorities, measurable outcomes, and the guardrails for risk, regulation and responsible AI.

2 Build the digital core and AI foundations

This stage creates the conditions for AI to scale reliably. It connects data, updates identity and consent, and puts in place interoperable architecture and a strong enterprise platform. Organizations move successful pilots into production by embedding AI into workflows, orchestration layers, and multi-step decisions. They also establish continuous oversight of drift, bias, explainability, and ethical compliance so deployment stays controlled, scalable and resilient.

3 Implement a two-speed AI approach

This stage balances near-term results with long-term advantage. One speed focuses on rapid customer and operational gains, while the other builds strategic AI that learns, improves, and scales across the enterprise. Organizations expand AI across business units, automate repeatable decisions, and refine models using enterprise feedback and outcome measurement. This can create immediate efficiency and experience gains while building sustained value through growth, cost reduction, and stronger risk and capital management.

4 Transform workforce, processes and culture

This stage embeds AI adoption and turns it into enterprise value. During this stage, work, decision rights, incentives, and leadership roles are redefined and embedded into how the organization operates. Organizations implement skills-based structures, institutionalize AI-driven decision-making, and build reusable capabilities that can scale across the business. The result is a new operating model and culture in which AI supports continuous improvement, stronger adoption, and more ambitious innovation.

5.3 Leadership, sponsorship and accountability

Embedding AI into an organization requires engagement from leaders across the C-suite. Every leader contributes a distinct perspective, whether in strategy, operations, finance, risk management, technology or human resources, that collectively cultivates innovation, governance and trust. Achieving success relies on shared responsibility and effective collaboration among all parties involved.

It also requires new leadership models. For example, IT and HR leads are aligning more closely

as AI integration drives changes in work, workflows and governance needs across the organization. The chief AI officer role is increasingly common, sometimes reporting directly to the CEO or jointly to technology and HR leaders. More variations, such as chief work officer or chief workforce transformation officer, may emerge to coordinate digital and human resources. These new roles would not replace existing positions. Rather, they would synchronize strategies for work, workforce and technology as the company scales agentic AI.

TABLE 1 Emerging areas of accountability for leaders

Role	Accountability
CEO Chief Strategy Officer (CSO)	<ul style="list-style-type: none"> – Set the strategic intent and sponsor AI transformation. – Communicate to staff, shareholders and regulators how AI is integrated into the overall strategy. – Lift AI to a board-level priority, build stakeholder confidence.
Chief Operating Officer (COO) Business heads	<ul style="list-style-type: none"> – Select processes with established baselines, identifiable risks, accessible data and clear benefits. – Define SLAs for blended digital workers and manage risks. – Start adoption in controlled, high-value impact areas.
Chief Financial Officer (CFO)	<ul style="list-style-type: none"> – Track costs before and after, assess financial impact once quality is confirmed, and allocate expenses including compute and licensing. – Develop reporting that reflects digital capacity as part of the workforce. – Embed digital workers in financial management.
Chief Risk Officer (CRO) Chief Compliance Officer (CCO)	<ul style="list-style-type: none"> – Define risk appetite. Require validation, continuous monitoring, anomaly detection and rollback options. – Ensure agent actions are explainable and logged. – Protect from compliance breaches and provide regulators with confidence.
Chief Information Officer (CIO) Chief Technology Officer (CTO) Chief AI Officer (CAIO) Chief Data Officer (CDO) Chief Information Security Officer (CISO)	<ul style="list-style-type: none"> – Publish a reference architecture covering identity, entitlements, data governance, tool access and telemetry. Define agent certification requirements. – Plan for vendor and in-house agents. Ensure interoperability, and security, across platforms. – Provide the technical backbone to scale safely and avoid vendor lock-in.
Chief Human Resources Officer (CHRO) Chief People Officer (CPO)	<ul style="list-style-type: none"> – Partner closely with business to redesign roles and reskill employees. Identify hybrid roles such as agent supervisors. – Update HR frameworks so workforce planning and performance management includes digital workers. – Ensure employees view digital workers as collaborators.

Conclusion

AI marks a structural technological shift that some experts are comparing to the internet, electricity or the introduction of rail. The implications are profound for financial services and the public institutions that shape trust, stability and inclusion. This report offers industry and public sector leaders across all geographies practical examples and guidance to support the development or refinement of AI strategies and to translate ambition into sustained, governable execution.

The case studies, provided by members of the World Economic Forum's AI in Financial Services community, are central to that intent. They move the conversation from what AI could mean to what leaders are learning as they build foundations, set guardrails, redesign work and scale new capabilities in live environments. These contributions make the guidance more concrete, credible and immediately useful for peers navigating similar choices.

No organization has a complete playbook. Capabilities, risks and supervisory expectations are evolving quickly, and strategies will need to be revisited as agentic systems mature and as new operating models emerge. However, as the financial services industry moves from digitization to cognition-enabled economies, one thing is clear: guided by trust and governance, comprehensive AI adoption will define the future of financial services and underpin the global economy.

That is why knowledge-sharing is so critically important. Exchanging practical lessons, patterns that work and approaches that fail safely accelerates progress for everyone, raising the floor on governance and trust while helping the industry capture value responsibly. The Forum supports this ongoing dialogue and welcomes further contributions, comparisons and collaboration to advance the sector's collective progress.

Appendices

A1 Overview of risk framework adoption by region (non-exhaustive)

Region	Sub-region/ country	Common risk/governance frameworks used by financial services for AI	Application/key call-outs
Global	Cross-jurisdiction	<ul style="list-style-type: none"> – NIST AI RMF 1.0 (+ genAI profile) – ISO/IEC 42001 – ISO/IEC 23894 – ISO/IEC 27001/27002 (security controls) – OECD AI Principles – UNESCO on Ethics of AI 	Used as a common baseline to harmonize responsible AI requirements across business lines and countries, then mapped to local regulator obligations. ISO 42001 is often adopted as the management-system wrapper; NIST/ISO 23894 provide practical risk processes; OECD/UNESCO are used most often for principles, policy language and external stakeholder alignment.
Americas	US	<ul style="list-style-type: none"> – Federal Reserve SR 11-7/ OCC 2011-12 (updated to SR 26-2) (Model Risk Management); NIST AI RMF; sector cybersecurity baselines (e.g. NIST CSF as a control reference) 	Extends model risk management to AI/ML/genAI (inventory, tiering, validation, monitoring, independent review), with NIST AI RMF used as overlay for broader non-model risks (human impact, governance, transparency). Strongest applicability is banking (prudential focus), but insurers and fintechs commonly align as it provides evidence-grade governance expectations for regulated decisions.
	Canada	<ul style="list-style-type: none"> – OSFI guidance/ expectations (model risk, operational risk, third-party risk) – ISO/NIST overlays 	Typically treated as part of enterprise risk management. AI is governed through model risk (where predictive), operational risk and third-party risk, with ISO/NIST used to standardize practices across groups. Practical call-out: regulator focus is often on robustness, controls and accountability rather than a single AI-specific statute.
	Latin America	<ul style="list-style-type: none"> – Local financial supervisor requirements and guidance (varies by country) – Local data protection laws – Central bank/payments regulator rules for digital finance – Voluntary standards (ISO/IEC 42001; ISO/IEC 23894; NIST AI RMF) 	In Latin America banks, insurers and fintechs typically implement AI governance through domestic banking/ insurance supervisors, central banks and data protection authorities, rather than a single region-wide AI statute. MNIs in the region usually overlay ISO/NIST to standardize controls and evidence across countries, and local firms also adopt these standards selectively to meet partner/outsourcing assurance needs, cross-border group requirements and investor expectations.
Europe, the Middle East and Africa	EU	<ul style="list-style-type: none"> – EU AI Act – GDPR – DORA – EBA guidance on ML for IRB (banks) – ISO/IEC 42001/23894 	EU AI Act is the dominant compliance driver (risk classification + obligations), with GDPR shaping data/ automated decisioning controls and DORA driving resilience and third-party management for AI services. Banking internal-model use cases often add EBA/ ECB supervisory expectations; ISO 42001/23894 are commonly used to operationalize a management system and risk process across the organization.
	UK	<ul style="list-style-type: none"> – PRA SS1/23 Model Risk Management principles – FCA expectations (conduct/consumer outcomes) – ISO/NIST overlays 	UK approach is commonly “MRM-first” (especially for banks): firms extend PRA model risk principles to AI/ genAI and use ISO/NIST to broaden into enterprise governance (policy, controls, audit evidence). Compared with the EU AI Act, the UK has historically relied more on regulator principles and supervisory expectations than a single cross-economy AI statute, so mapping and documentation discipline matters.

Region	Sub-region/ country	Common risk/governance frameworks used by financial services for AI	Application/key call-outs
Europe, the Middle East and Africa	Middle East	<ul style="list-style-type: none"> – National AI strategies/ethics principles (vary) – ISO/IEC 42001 – NIST AI RMF – Sector regulator guidance (where issued) 	Many financial services firms use ISO/NIST as the dominant operational baseline and then tailor to national digital/AI strategies and data laws. A common differentiator is strong emphasis on national digital transformation agendas, with governance often driven by regulators and central banks for systemic stability and cyber/third-party risk.
	Africa	<ul style="list-style-type: none"> – National AI strategies/data protection regimes (vary); ISO/IEC 27001; ISO/IEC 42001; NIST AI RMF 	Often firms lean on international standards (ISO/NIST), especially where local AI-specific regulation is still emerging. Focus is often on privacy, financial inclusion/consumer protection and operational resilience, with regulators/central banks influencing expectations.
Asia-Pacific	Singapore	<ul style="list-style-type: none"> – MAS FEAT Principles – MAS AI Model Risk Management – IMDA Model AI Governance Framework/AI Verify – ISO/NIST overlays 	Singapore is a leading reference point for FS-specific fairness/ethics expectations (FEAT) and practical model risk governance for AI, with tooling/assurance approaches promoted via AI Verify. MAS guidance used as a “gold standard” overlay on ISO/NIST, especially for customer-impacting decisions (credit, pricing, fraud, advice).
	Hong Kong	<ul style="list-style-type: none"> – HKMA High-level Principles on AI – ISO/NIST overlays 	Used primarily as supervisory expectations for governance, accountability, and controls, commonly implemented through model risk management and third-party risk processes. Compared with Singapore’s FEAT, the HK approach is more principles-led and typically relies on firms’ internal frameworks to evidence implementation.
	Australia	<ul style="list-style-type: none"> – APRA prudential expectations (operational risk, CPS 230/outsourcing, information security, model risk practices) – ISO/NIST overlays 	AI governance is commonly implemented via existing prudential risk disciplines, operational resilience, third-party risk, security and (increasingly) model risk rather than a single AI rulebook. Tends to be most stringent for banks and insurers regulated by APRA; fintechs align to meet partner and customer assurance expectations.
	Japan	<ul style="list-style-type: none"> – Japan AI governance guidelines (cross-industry) – ISO/IEC 42001 – NIST AI RMF 	Firms typically align to national governance guidance and then operationalize via ISO/NIST to create auditable controls across the AI life cycle. Compared with the EU, the approach is less compliance-by-risk-tier and more governance-by-principles, so internal control design and monitoring are central.
	India	<ul style="list-style-type: none"> – National AI guidance/initiatives (evolving) – Sector regulator expectations (RBI/IRDAI where applicable) – ISO/NIST overlays 	Financial services firms commonly implement global-standard governance (ISO/NIST) and tailor to local data rules, outsourcing, and sector supervisory expectations as they evolve. Given the fintech/neobank ecosystem and extensive third-party dependencies, third-party risk, data governance and consumer protection controls are often emphasized.
	China (Mainland)	<ul style="list-style-type: none"> – Cybersecurity Law (CSL) – Data Security Law (DSL) – Personal Information Protection Law (PIPL) – CAC algorithm governance rules – NFRA data security rules (banking and insurance) 	Local Chinese financial services firms typically implement AI governance as part of a broader compliance stack, focused on data localization, security, algorithm filing and assessments, and strong content/safety controls MNCs operating in China (including foreign banks, insurers with PRC entities/ branches/JVs) generally run a dedicated “China compliance” variant of their global AI framework to meet data residency/cross-border transfer constraints and local supervisory expectations, often requiring separate tooling, vendors and governance evidence tailored to Chinese regulators.


A2 Overview of AI-driven risks and corresponding risk management approaches

Risk category	Detail	GenAI/agentic impact	AI-specific foundational capability requirements
Reputational risk	Risks arising from loss of trust and credibility with customers, regulators and the public when AI-assisted actions are wrong or misleading	<ul style="list-style-type: none"> – Unreliable model outputs may bypass review – Human challenge may weaken, as over-reliance and automation bias reduce escalation and second-line checks – External deception outpaces controls, as genAI increases fraud and disinformation 	<ul style="list-style-type: none"> – Quality controls that detect and block high impact misstatements, with thresholds for escalation – Customer visible transparency for AI involvement and a fast path to human resolution when confidence is low or impact is high – AI incident response, for example: misstatement, deepfake and public misinformation incidents
Legal and regulatory risk	Risks arising from laws, regulation, supervision, legal enforceability, conduct obligations and financial crime	<ul style="list-style-type: none"> – Regulators increasingly expect demonstrable controls and governance for AI use (not just principles) – GenAI raises financial crime risk (identity fraud, deception, disinformation) intersecting legal obligations – Credit scoring and similar use cases trigger strict governance, documentation and human oversight expectations 	<ul style="list-style-type: none"> – Compliance-by-evidence audit trails that capture what the system did and why, including log, monitoring artefacts, model/agent governance records – Codify human accountability for regulated decisions (who approves, who can override, how intervention is recorded) – GenAI-enabled financial crime control uplift (prevention, detection, response, reporting)
Financial risk	Risks that primarily impact capital, earnings, liquidity or valuation through financial exposures	<ul style="list-style-type: none"> – Similar models/data across market participants can increase correlation in price moves/hedges, worsening value at risk/stress losses – Faster AI-driven adjustment in markets can amplify intraday volatility, widen bid-ask at moments liquidity might be needed – Automation can make limit changes/risk actions happen faster, steepening P&L swings when regimes shift 	<ul style="list-style-type: none"> – AI crowding and correlation stress testing (model similarity, correlated unwind scenarios, ties outcomes to buffers) – Intraday liquidity risk governance that matches AI speed (key risk indicators, triggers, escalation) – Automated action circuit breakers that pause or constrain model driven limit changes and rebalancing during stress
Operational resilience risk	Risks that threaten the bank's ability to deliver critical operations within tolerance (end-to-end service continuity)	<ul style="list-style-type: none"> – Agentic workflows may reduce human buffers and failures can propagate across multiple steps and teams quickly – Critical services may depend on additional AI providers and outages or loss of access can impair service delivery – Agent decision paths and dynamic execution may complicate root cause analysis, rollback and controlled restart 	<ul style="list-style-type: none"> – Resilience testing for agent failure cascades and mass exception events (not just classic IT outages) – Define safe fallback modes and service tolerances and that keep critical services operating in degraded states – Third party continuity and exit controls for AI-critical suppliers with evidence, recovery playbooks and concentration risk management
Technology and cyber risk	Risks arising from technology engineering and operation, as well as hostile compromise of systems/ information	<ul style="list-style-type: none"> – New AI interfaces and data flows may create more opportunities for compromise and attackers also use AI – Reliance on a small set of AI/cloud/model providers may increase exposure if one is disrupted or compromised – Deepfakes/voice cloning weaken human trust signals used in approvals and servicing 	<ul style="list-style-type: none"> – Agent identity and privilege management a tier-0 security (tool access control, segregation of duties, agent action monitoring) – AI interface security controls (new attack surfaces, trust attacks, including deepfake-resistant verification for approvals and servicing) – AI supplier assurance integrated into cyber third-party risk, including incident coordination and provider concentration review
Data and model risk	Risks arising from data integrity/use and model behaviour, including AI/genAI and decision automation	<ul style="list-style-type: none"> – Complexity, limited explainability and opaque training data may complicate data quality assessment and model validation – Hallucinations and non-determinism may create new accuracy and reproducibility challenges – Frequent updates to models/prompts/tools may increase drift and control burden 	<ul style="list-style-type: none"> – Inventory and materiality for AI systems so controls scale proportionately to risk – Independent validation and continuous monitoring (including hallucinations, non-determinism, frequent changes to models) – Unstructured data governance (traceability, quality checks, provenance)



Cross-industry

Impact on industrial ecosystems



Organizational Transformation in the Age of AI: How Organizations Maximize AI's Potential



AI in Action: Beyond Experimentation to Transform Industry



From Paradox to Progress: A Net-Positive AI Energy Framework



Artificial Intelligence and Cybersecurity: Balancing Risks and Rewards



Regional-specific

Impact on regions



Asia's Human-led AI Opportunity: A Framework for Transformation



Industry- or function-specific


Impact on industries, sectors and functions

Advanced manufacturing and supply chains



Intelligent Industrial Operations Outlook 2026

Consumer goods




Unlocking AI Value at Scale: Workforce Adoption A Practical Guide for Consumer Industries

Healthcare



The Future of AI-Enabled Health: Leading the Way

Life science



Strategic Choices in the Age of AI: Shaping the Future of Life Sciences

Media, entertainment and sport



Artificial Intelligence in Media, Entertainment and Sport

Telecommunications




The Strategic Role of Telecom Providers Across the AI Value Chain

Transport



Intelligent Transport, Greener Future: AI as a Catalyst to Decarbonize Global Logistics

Financial services



Upcoming

Contributors

World Economic Forum

Andre Belelieu

Head, Financial Services Industries; Head, Business Engagement, North America

Drew Propson

Head, Technology and Innovation in Financial Services

Accenture

Francesca Caminiti

Global Head of Research, Financial Services

David Parker

Global Industry Lead, Banking and Capital Markets

Astrid Schlunzen

Associate Director, Financial Services

Corrine Vitolo

Head of Research, North America Financial Services

Acknowledgements

The project team thanks the following individuals for their contributions and participation in roundtables, workshops and interviews (in alphabetical order):

Usman Ahmed

Senior Director, Payments Market and Partnerships, Corporate Strategy, PayPal

Emina Ajvazoska

Coordinator, Investors and Financial Services Industries, World Economic Forum

Ximena Aleman

Co-Founder and Co-Chief Executive Officer, Prometeo

Sultan Alhamidi

Chief Executive Officer, Social Development Bank

Hisham Alrayes

Group Chief Executive Officer and Member of the Board, GFH Financial Group

Nada AlSaeed

Chief of Strategy, Bahrain Economic Development Board

Gustavo Alves

Founder and Chief Executive Officer, Nagro

Agshin Amirov

Deputy Chief Executive Officer and Chief Risk Officer, The International Bank of Azerbaijan (ABB)

Seif Amr

Co-Founder, Thndr

Omair Ansari

Co-Founder and Chief Executive Officer, Abhi

Marc Antaki

Deputy Chief Strategy and Risk Officer, Mubadala Investment Company

Michelle Auchincloss

Chief Strategy Officer, State Street

Aizal Baharuddin

Head, Strategic Research, Corporate Strategy, PETRONAS

Wachi Bandara

Head, Artificial Intelligence, DRW Holdings

Rweyemamu Barongo

Head, AI and Data, Innovation Hub, Bank of Tanzania

Maria Basso

Head, AI Applications and Impact, World Economic Forum

Ahmed Bawareth

Senior Executive Director, Service Excellence Office, National Development Fund (NDF)

Anke Bivens

Global Client Partner, Cognizant

Niall Byrne

Chief Financial Officer, Qatar Investment Authority

Lawrence Calcano

Chairman and Chief Executive Officer, Institutional Capital Network

Morgan Camp

Coordinator, Financial and Monetary Systems, World Economic Forum

Teresa Carlson

Chief Executive Officer, General Catalyst

Jesse Caruso

Manager Financial Services, Accenture

Komes Chandavimol

Principal AI Evangelist, KASIKORN Business Technology Group (KBTG), Kasikornbank

Rohit Chauhan

Executive Vice President of AI and Security Solutions, Mastercard

Tony Chen

Executive Vice-President and Chief Digital Officer, Fubon Financial

Caecilia Chu Wing Tak

Co-Founder and Chief Executive Officer, You Technologies Group Limited

Federico Cohen Freue

Senior Vice-President, AI and Data Operations, Mastercard

John Collison

President and Co-Founder, Stripe

Joseph Cordeira

Chief Data and Analytics Officer, Global Wealth Management Americas, UBS

Jennifer Cosco

Global Head of Government Relations and Regulatory Strategy, London Stock Exchange Group (LSEG)

Adrian Cox

Managing Director, Thematic Strategist, Deutsche Bank Research, Deutsche Bank

Stefano D'Ambrosio

Director, Digital Strategy and Innovation, Cornerstone United Holdings Jamaica (CUHJ)

Edward Dandridge

Executive Vice-President; Chief Marketing and Communications Officer, American International Group (AIG)

Sebastien Danloy

Chief Business Officer, Euroclear Group

Jonathan Davis

Data Science and AI Lead, Zurich Insurance Group

Bhavesh Dayalji

Chief AI Officer, S&P Global and CEO, Kensho Technologies (S&P)

Peter De Caluwe

Co-Founder and Chief Executive Officer, Thunes

Miroslav Dimitrov

Chief Executive Officer, Growth Protocol

Dante Disparte

Chief Strategy Officer and Head, Global Policy, Circle

James Dodds

Vice-President, TD Bank Group

Kris Dorr

Managing Director and Head of Rates, Guggenheim Partners

Denise Dourado

Director, Digital & Application Innovation and Data & AI, Microsoft

Ariane Dwyer

Executive Director, JPMorgan Chase

Lisel Engelbrecht

Executive Head, Data Engineering, Analytics and AI, Standard Bank

Nathan English

Senior Manager, Financial Services, Accenture

Mark Fairless

Group Chief Executive Officer, ClearBank

Naïma Al Falasi

Senior Vice-President, AI Strategy and Transformation, Mubadala Investment Company

Hisham Al-Falih

Co-Founder and Chief Executive Officer, Lean Technologies

Song Feifei

Chief Executive Officer, FESCO Adecco

Virginia Folgueiro

Chief Executive Officer, Menta

Robert Forsythe

Chairman, Growth Protocol

Aaron Friedman

Deputy U.S. Head of Government and Regulatory Policy, Citadel

Dennis Gada

Executive Vice-President and Global Head of Banking and Financial Services; Member of the Executive Committee, Infosys

Julie Gerdeman

Global Head, Data and Analytics, The Bank Of New York Mellon Corporation (BNY)

Glenn Goldman

Co-Founder, Plurall

Federico Gomez

Founder and Chief Executive Officer, Plurall

David Griffiths

Chief Technology Officer, Citi

Marcela Grover

Senior Legal Director, Transformation and Governance, London Stock Exchange Group (LSEG)

Ahmad Hammouda

Co-Founder and Chief Executive Officer, Thndr

Claus Harder

Head, Group Strategic Steering, Danske Bank

Marc Hauert

Chief Strategy Officer, Santander UK

Rob Heyvaert

Founder and Managing Partner, Motive Partners

Janine Hirt

Chief Executive Officer, Innovate Finance

Jill Hoang

Initiatives Lead, Applied AI, Centre for AI Excellence, World Economic Forum

Rob Hornby

Co-Chief Executive Officer, AlixPartners

He Huajie

Group President, Tread

Hu Dianming

Chief Executive Officer, SenseDealAI

Kim Huffman

Senior Vice-President, Chief Information Officer, Workiva

Vincent Henry Iswaratioso

Co-Founder and Chief Executive Officer, DANA Indonesia

Jon Jacobson

Chief Executive Officer and Co-Founder, Omnicient

Zhang Jiwei

Editor-in-Chief, Caixin.com, Caixin Media

Jean-Marc Joris

Head, Organization and Information Technology; Member of the Executive Board, Banque Cantonale de Genève

Tiernan Kennedy

Chief Executive Officer, Umba

Sudeep Kesh

Chief Innovation Officer, S&P Global Ratings, S&P Global

Ram Komaraju

Managing Director, Technology, CLS Bank International

Eleni Kouletaki

Group General Counsel, The Olayan Group

Bijoy Koyitti

Senior Editor, Accenture Research

Suresh Krishnasamy

Head of Enterprise Strategy and Chief Financial Officer of Technology of Operations, Bank of Montreal

Rob Krugman

Senior Vice-President, Digital Strategy and Innovation; Chief Digital Officer, Broadridge Financial Solutions

Pravina Ladva

Group Chief Digital and Technology Officer, Swiss Re

Ajay Lakhotia

Chief Executive Officer, Assetgro Fintech (STOCKGRO)

Kenny Lam

Chief Executive Officer, Asia-Pacific, Two Sigma

Juliette Xue Lascoux

General Manager, Shanghai, Skandinaviska Enskilda Banken (SEB)

Philippe Laurensy

Chief Executive Officer, Asia-Pacific, Euroclear

Kelvin Lee

Co-Founder and Chief Executive Officer, Alta

Chen Leiming

Senior Vice-President, Ant Group

Nellie Liang

Director, Office of Financial Stability Policy and Research, Federal Reserve Board, Federal Reserve System of the United States

Lily Ma

General Manager, North Asia, ADP

Sarah Macfarlane

General Counsel, Technology and Operations, London Stock Exchange Group (LSEG)

Romana Maione

Director of Product Management, Head of Regulatory Reporting Solutions, Americas, Nasdaq

Francis Malige

Managing Director, Financial Institutions, European Bank for Reconstruction and Development (EBRD)

Vishal Marria

Chief Executive Officer and Founder, Quantexa

Regina Maruca

Senior Editor, Accenture Research

Joseph Masri

Chief Risk Officer, General Retirement and Social Insurance Authority (GRSIA)

Linda Maxwell
Operating Partner, DCVC Management

Sarkis Mazmanian
Head of Fintech Risk, Deutsche Bank

Bob McCooley
Vice-Chairman, Nasdaq

Willard McLane
Group Head, Strategy and Corporate Development, HSBC

Mario Magalhães Carvalho Mesquita
Global Head, Macroeconomics and Research, Itaú Unibanco

Jo Miyake
Head, Banking, Asia and Middle East, HSBC

Nicolas Moch
Head of SEBx, Skandinaviska Enskilda Banken (SEB)

Shain Mohammed
AI Generalist, LuLu Financial Holdings

Zia Mohzani
Vice President, Private Markets, Permodalan Nasional (PNB)

Amy Mollin
Vice-President, Product Management Applications, Guidewire Software

Yao Morin
Chief Technology Officer, JLL

Daniel Murphy
Industry Communities Lead, World Economic Forum

Daniel Murray
Deputy Chief Information Officer, EFG International

Zebon Mwashambwa
Head of AI Strategy Implementation, Bank of Tanzania

Kivashan Naidoo
Head of Strategy, COO Portfolio, The Standard Bank

Monsinee Nakapanant
Co-President, Ascend Group

Mounir Nakhla
Founder and Chief Executive Officer, MNT Investments

Asad Nasir
Chief Executive Officer, JS Group

Osama Nasr
Group Chief Digital and Technology Officer, GFH Financial Group

Ng Swee Khiang
Head, Primary Market Section, Capital Market Department, Employees Provident Fund (EPF)

Margaret Nienaber
Chief Operating Officer, The Standard Bank

Yasuyuki Nishihara
Managing Director and Head of Information Systems Planning Division, Mitsubishi UFJ Financial Group (MUFG)

Jennifer O'Neil
Managing Director; Head, Corporate Strategy, BlackRock

Paulo Ossamu
Senior Managing Director, Technology Lead, Asia-Pacific, Accenture

Muzzaffar Othman
Group Chief Operations Officer, Permodalan Nasional (PNB)

Jorgen Ouaknine
Group Head, Innovation and Digital Assets, Euroclear

Pinar Ozcan
Professor, University of Oxford

Sebastian Petric
Head of FX Strategy, LGT

Alistair Phelps
Managing Director, Wholesale Chief Strategy Officer, Nomura Securities International

Ian Phoenix
Director of Intelligence and Digital, Financial Conduct Authority

Julieta Picorelli
Manager, Capital Markets, Accenture Research

Thadpong Pongthawornkamol
Managing Director and Distinguished Visionary Architect, Kasikorn Business Technology Group (KBTG), Kasikornbank

Ruangroj Poonpol
Group Chairman, Kasikorn Business-Technology Group (KBTG), Kasikornbank

Johnna Powell
Managing Director; Head, Technology Research and Innovation, The Depository Trust and Clearing Corporation (DTCC)

Ritwik Priya
Managing Director and Head of EMEA Strategy, Nomura Securities International

Jakša Puljiz
Acting Director, Institute for Development and International Relations (IRMO)

Philipp Raether
Group Chief Privacy and AI Trust Officer, Allianz

Farheen Rahimtoola
Strategy Manager, Executive Director,
JPMorgan Chase

Clairelle Rakipovic
Chief Executive Officer, Pipe

Archie Giridhar Ravishankar
Co-Founder and Chief Executive Officer, Cogni

Rahul Rekhi
President, Rogo AI

Francesca Hopwood Road
Centre Head, BIS Innovation Hub – London,
Bank for International Settlements (BIS)

Vincent Robert
Chief Operating Officer, Arab Bank Switzerland

Darrick Rochili
Chief Innovation Officer, DANA Indonesia

Jessica Rusu
Chief Data, Information and Intelligence Officer,
Financial Conduct Authority (FCA)

Anoop Sagoo
Chief Executive Officer, South-East Asia, Accenture

Omar Saleh
Chief Executive Officer, Khazna

Giridar Prasad Sankaran
Executive Vice-President, Software and Platform
Engineering, Cognizant

Santitarn Sathirathai
Adviser, Future Economy, Thailand Development
Research Institute (TDRI)

Rene Saul
Co-Founder and Chief Executive Officer, Redem

Martha Sazon
President and Chief Executive Officer, Mynt

Roland Scharrer
Partner and Managing Director; European Data
and AI Leader, Kearney

Jeffrey Schumacher
Founder, Growth Protocol

Aaron Schumm
Founder, Chief Executive Officer and Chairman,
Vestwell

Nihat Senyuva
Chief Information Officer,
The International Bank of Azerbaijan (ABB)

Srinivasan Seshadri
Chief Growth Officer; Global Head,
Financial Services, HCLTech

Andrey Severyukhin
Chief Executive Officer, Sumsud

Mina Shahid
Founder and Chief Executive Officer, Numida

Roshan Shetty
Head, Banking, Financial Services and Insurance
(BFSI) and Public Sector, Tech Mahindra (Americas)

Cecilia Skingsley
Governor of Stockholm County, City of Stockholm

Eric Solvet
Managing Director, Compagnie Financière Tradition

Ram Srinivasan
Managing Director, Consulting, Work Dynamics, JLL

Jirayut Srupsrisopa
Founder and Group Chief Executive Officer, Bitkub

John Stackhouse
Senior Vice-President, RBC Financial Group

David Sun
Senior Executive Vice-President,
Cathay Financial Holdings

Sun Shoudong
Member of the Board, Dalian Pengsheng

Madras Seetharaman Suresh
Global Associate Vice-President and Head,
Sales, Infosys (China)

Sam Swartz
Head of Central Policy, Stripe

Lukasz Szpruch
Programme Director for Finance and Economics,
The Alan Turing Institute

Taleh Tahirli
Deputy Chief Executive Officer and Chief Strategy
Officer, The International Bank of Azerbaijan (ABB)

Abby Tan
General Manager, Brunei, Yidu Tech

John Tang
Managing Director, Private Equity,
Office of the Chief Investment Officer, GIC

Laurel Taylor
Founder and Chief Executive Officer, Candidly

Mark Thumser
Chief Financial and Strategy Officer, Americas,
Mitsubishi UFJ Financial Group (MUFG)

Ravi Tiwari

Group Corporate Development, Strategy and Investor Relations Director, Lloyds

Luiz Carlos Trabuco Cappi

Chairman of the Board, Banco Bradesco

Calvin Tse

Vice-President, Business Development and Strategy, WorldBridge Group of Companies

Greg Ulrich

Chief AI and Data Officer, Mastercard

Gregory Van

Chief Executive Officer, Endowus

Heng Wang

Professor of Law, Yong Pung How School of Law, Singapore Management University

Maik Taro Wehmeyer

Chief Executive Officer, Taktile

Jonathan Welburn

Senior Researcher and Professor, RAND

James Whittington

Senior Vice-President, Financial Services, General Atlantic Service Company

Michael Wignall

Senior Director, Cloud, Solution Architecture, Microsoft

Oliver Winkenbach

General Manager, Data Applications, Guidewire Software

Edward Woodford

Co-Founder and Chief Executive Officer, Zero Hash

Haimera Workie

Vice-President and Head, Financial Innovation, Financial Industry Regulatory Authority (FINRA)

Liu Xi

Founder and Chief Executive Officer, FOMO Pay

Yu Yi

Senior Managing Director, Accenture Technology Lead, Greater China, Accenture Ali Business Group Lead, Accenture

Carol Yu

Founding Partner and Associate Dean, Shenzhen InnoX Academy

Markos Zachariadis

Professor and Chair in Financial Technology and Information Systems, University of Manchester

Fernando Zandona

Chief Executive Officer, Mambu Tech

Bryan Zheng Zhang

Executive Director and Co-Founder, Cambridge Centre for Alternative Finance, Cambridge Judge Business School, University of Cambridge

Sarah Zhang Jiachen

Founder and Chief Executive Officer, Guangzhishu Technology

Arthur Zou Jianglei

Chief Executive Officer, China, First Abu Dhabi Bank (FAB)

Thomas Zschach

Chief Innovation Officer, Global Head for Innovation and Architecture, SWIFT

Angela Zutavern

Partner and Managing Director, AlixPartners

Production

Laurence Denmark

Creative Director, Studio Miko

Sophie Ebbage

Designer, Studio Miko

Will Liley

Editor, Studio Miko

Cat Slaymaker

Designer, Studio Miko

World Economic Forum Public Engagement

Maxwell Hall

Creative Editorial Lead

Floris Landi

Design Lead

Gayle Markovitz

Head, Written and Audio Content

Sybille Penhirin

Head, Video and Design

Endnotes

1. Close, K., Wroblewski, S., Van Singel, J., Ahmad, S., et al. (2026). *Talent Reinventors: Delivering value with and for people in the age of AI*. Accenture. <https://www.accenture.com/us-en/insights/consulting/talent-reinventors-delivering-value-people-age-ai>.
2. Accenture. (2026). *Top Banking Trends for 2026. Unconstrained Banking: A new age of possibility*. <https://www.accenture.com/us-en/insights/banking/accenture-banking-trends-2026>.
3. Stanford University Human-Centered Artificial Intelligence. (2026). *Artificial Intelligence Index Report*. https://hai.stanford.edu/assets/files/ai_index_report_2026.pdf.
4. Executive Office of the President of the United States. (2025). *Ensuring a National Policy Framework for Artificial Intelligence*. <https://www.federalregister.gov/documents/2025/12/16/2025-23092/ensuring-a-national-policy-framework-for-artificial-intelligence>.
5. Cyber Risk Institute. (n.d.) *Financial Services AI Risk Management Framework*. <https://cyberriskinstitute.org/artificial-intelligence-risk-management/>.
6. The White House. (2026) *A National Policy Framework for Artificial Intelligence*. <https://www.whitehouse.gov/wp-content/uploads/2026/03/03.20.26-National-Policy-Framework-for-Artificial-Intelligence-Legislative-Recommendations.pdf>.
7. Financial Conduct Authority, United Kingdom (2026). *AI and the FCA: our approach*. <https://www.fca.org.uk/firms/innovation/ai-approach#revisions>.
8. European Union. (2024). *Regulation (EU) 2024/1689 of the European Parliament and of the Council*. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>.
9. Monetary Authority of Singapore (MAS). (2018). *Principles to Promote Fairness, Ethics, Accountability and Transparency (FEAT) in the Use of Artificial Intelligence and Data Analytics in Singapore's Financial Sector*. <https://www.mas.gov.sg/publications/monographs-or-information-paper/2018/feat>.
10. Monetary Authority of Singapore (MAS). (2026). <https://www.mas.gov.sg/schemes-and-initiatives/project-mindforge>.
11. Hong Kong Monetary Authority. (2019). *High-level Principles on Artificial Intelligence*. <https://www.hkma.gov.hk/media/eng/doc/key-information/guidelines-and-circular/2019/20191101e1.pdf>.
12. Financial Services Authority of Indonesia (OJK). (2025). *Tata Kelola Kecerdasan Artifisial Perbankan Indonesia*. <https://www.ojk.go.id/id/Publikasi/Roadmap-dan-Pedoman/Perbankan/Documents/Tata%20Kelola%20Kecerdasan%20Artifisial%20Perbankan%20Indonesia.pdf>.
13. Ministry of Communication and Digital Affairs of Indonesia. (2026). *Sekjen Ismail: Perpres AI Fondasi Tata Kelola Teknologi Masa Depan Indonesia*. <https://www.komdigi.go.id/berita/siaran-pers/detail/sekjen-ismail-perpres-ai-fondasi-tata-kelola-teknologi-masa-depan-indonesia>.
14. International Data Corporation (IDC). (2024) *IDC's Worldwide AI and Generative AI Spending – Industry Outlook*. <https://www.idc.com/resource-center/blog/idcs-worldwide-ai-and-generative-ai-spending-industry-outlook/#:~:text=IDC%20has%20recently%20unveiled%20the,significance%20in%20the%20AI%20landscape>.
15. Stanford University Human-Centered Artificial Intelligence. (2026). *Artificial Intelligence Index Report*. https://hai.stanford.edu/assets/files/ai_index_report_2026.pdf.
16. NVIDIA. (2026). *State of AI in Financial Services*. <https://resources.nvidia.com/en-us-2026-fsi-survey/finance-state-of-ai-report>.
17. Ibid.
18. Ibid.
19. World Economic Forum. (2026). *Conversation with Jamie Dimon, Chairman and CEO of JPMorgan Chase*. <https://www.weforum.org/meetings/world-economic-forum-annual-meeting-2026/sessions/conversation-with-jamie-dimon-chairman-and-ceo-of-jpmorgan-chase/>.
20. Case studies in this report were contributed by trusted members of the AI in Financial Services community. Submissions underwent editorial review; however, reported metrics and outcomes are self-reported and could not be independently verified. Results should be considered in the context of each organization's specific circumstances.
21. Thales. (2025). *Artificial Intelligence fuels rise of hard-to-detect bots that now make up more than half of global internet traffic, according to the 2025 Imperva Bad Bot Report*. <https://cpl.thalesgroup.com/about-us/newsroom/2025-imperva-bad-bot-report-ai-internet-traffic>.
22. Cambridge Centre for Alternative Finance, University of Cambridge. (2026). *The 2026 Global AI in Financial Services Report: Adoption, impact and risks*. <https://www.jbs.cam.ac.uk/wp-content/uploads/2026/04/ccaf-2026-04-28-global-ai-in-financial-services-report.pdf>.
23. Smith, K. and Abbott, M. (2024). *Banking in the age of generative AI*. Accenture. <https://www.accenture.com/us-en/insights/banking/generative-ai-banking>.
24. Shook, E. and Daugherty, P. (2024). *Work, workforce, workers: Reinvented in the age of generative AI*. Accenture. <https://www.accenture.com/us-en/insights/consulting/gen-ai-talent>.

25. World Economic Forum. (2025). *The Future of Jobs Report 2025*. <https://www.weforum.org/publications/the-future-of-jobs-report-2025/>.
26. Note: SR 26-2 explicitly excludes generative AI and agentic AI from its scope; the agencies have indicated a forthcoming Request for Information specifically addressing AI governance.
27. Cambridge Centre for Alternative Finance, University of Cambridge. (2026). *The 2026 Global AI in Financial Services Report: Adoption, impact and risks*. <https://www.jbs.cam.ac.uk/wp-content/uploads/2026/04/ccaf-2026-04-28-global-ai-in-financial-services-report.pdf>.
28. U.S. Department of Education. (2025). *U.S. Department of Education to begin federal student loan collections, other actions to help borrowers get back into repayment*. <https://www.ed.gov/about/news/press-release/us-department-of-education-begin-federal-student-loan-collections-other-actions-help-borrowers-get-back-repayment>; Board of Governors of the Federal Reserve System. (2026). Consumer credit – G.19. <https://www.federalreserve.gov/releases/g19/current/>.
29. Hanson, M. (2024). *143 student loan forgiveness programs*. Education Data Initiative. <https://educationdata.org/student-loan-forgiveness-programs>.
30. Granville, P., Nilaj, E., & Zhang, J. (2026). *Trump's student loan delinquency crisis, unmasked. The Century Foundation & Protect Borrowers*. <https://tcf.org/content/report/trumps-student-loan-delinquency-crisis-unmasked/>.
31. The Institute for College Access & Success. (2025). *On the edge of a "default cliff": New survey shows student loan borrowers are struggling to keep up*. <https://ticas.org/affordability-2/2025-student-debt-survey-blog/>.
32. Fidelity Investments. (2026). *Fidelity research reveals many borrowers delaying major life milestones due to student loan debt*. <https://newsroom.fidelity.com/pressreleases/fidelity-2026-state-of-student-debt/s/e9bdd85e-328d-43ec-8f65-bf5a97eaaabe>.
33. Mateega, S., Georgescu, C., & Tang, D. (2025). *FinanceQA: A benchmark for evaluating financial analysis capabilities of large language models*. arXiv. <https://arxiv.org/abs/2501.18062>; Farrington, R. (2025). *I asked ChatGPT about student loans. Here's what it got wrong*. The College Investor. <https://thecollegeinvestor.com/63419/chatgpt-on-student-loans/>.
34. Candidly. (2026). *Early findings: How Cait drives engagement*. <https://static.getcandidly.com/docs/Cait-Early-Findings-Report.pdf>.
35. Ibid.
36. Down, A. (2026). *Google warns quantum computers could hack encrypted systems by 2029*. The Guardian. <https://www.theguardian.com/technology/2026/mar/26/google-quantum-computers-crack-encryption-2029>.
37. NIST. (2026). *What Is Post-Quantum Cryptography?* <https://www.nist.gov/cybersecurity-and-privacy/what-post-quantum-cryptography>.



COMMITTED TO
IMPROVING THE STATE
OF THE WORLD

The World Economic Forum, committed to improving the state of the world, is the International Organization for Public-Private Cooperation.

The Forum engages the foremost political, business and other leaders of society to shape global, regional and industry agendas.

World Economic Forum
91–93 route de la Capite
CH-1223 Cologny/Geneva
Switzerland

Tel.: +41 (0) 22 869 1212
Fax: +41 (0) 22 786 2744
contact@weforum.org
www.weforum.org