



State of AI in the Financial Services Industry

The untapped edge

March 2026

deloitte.com/us/state-of-ai





◆◆◆◆◆ Table of contents

Overview

About the report

03

06

Key findings

Worker access and pilot expansion

AI transformation impacts

AI fluency and work redesign

Sovereign AI

Agentic AI

Physical AI

AI preparedness

07

07

10

13

16

18

24

30

Tapping into AI's full potential

32

Close the gap between access and activation

32

Unlock human advantage

33

Build governance before scale

34

Address sovereign AI requirements

35

Build a "living" technology & data infrastructure

36

Pursue strategic reinvention

37

Acknowledgements

38

About the Deloitte AI Institute

40

Methodology

41





Overview

Organizations today stand at the untapped edge of AI's potential. Ongoing developments in agentic, physical, and sovereign AI present new challenges and opportunities. Momentum is building, yet the greatest gains still lie ahead as financial services organizations translate early progress into scalable impact. As AI expands beyond the digital core, success will ultimately hinge on a company's ability to move from ambition to activation—turning experiments and potential into real-world business value at scale.

Financial services organizations are advancing AI in an environment defined by margin pressure, heightened regulatory scrutiny, legacy technology, and rapidly evolving customer expectations. Leaders are prioritizing near-term value from use cases such as fraud detection, credit risk

assessment, customer service automation, and personalized marketing, while making longer-term investments in autonomous operations, AI-enabled compliance and risk management, and modernized core platforms. At the same time, data fragmentation, model risk, and talent constraints continue to limit the pace at which AI can be scaled across the enterprise.¹

Deloitte's latest *State of AI in the Enterprise* survey captured insights from more than 570 business and IT leaders in financial service organizations around the world with direct involvement in their companies' AI initiatives. This report explores the survey insights in detail, sharing critical actions for leaders to consider as they continue on their journey with AI.



Overview

Key findings

AI is moving from the pilot and experimentation phase to enterprise scaling as worker access to AI expands.

Surveyed financial services companies have doubled worker access to AI in just one year—growing from 30% to 62% of workers now equipped with sanctioned AI tools. While only 24% of respondents said their organization has moved 40% or more of their AI experiments into production to date, 53% expect to reach that level in the next three to six months.

AI transformation reveals productivity for most, business reimagination for a few.

AI is already boosting productivity and efficiency; just a subset are using it to rewrite the business. Today, 31% of financial services companies are starting to use AI to deeply transform their businesses, 28% are redesigning key processes around AI, and the remaining 41% are only using AI at a surface level with little or no change to underlying business processes. While each are capturing productivity and efficiency gains, just the first group are truly reimaging their businesses rather than optimizing what already exists.

Companies are focused on building AI fluency instead of redesigning work around AI.

Despite high expectations for automation, 84% of financial services companies have not redesigned jobs or the nature of work itself around AI capabilities. In line with other industries, financial services organizations cite insufficient worker skills as the biggest barrier to integrating AI, yet fewer than half are making significant adjustments to their talent strategies. Most are focused on educating employees, but far fewer are re-architecting roles, workflows, and career paths.

With sovereign AI taking hold, where technology is built matters as much as what it can do.

Sovereign AI is about more than technology ownership. It's about strategic independence. More than 3 in 4 financial services companies (80%) say the location of AI development is a key factor when choosing new technologies, signaling that geographic sovereignty is now as important as innovation. These results align with a broader, cross-industry shift toward recognizing sovereignty as a key part of decision-making.

AI agents are scaling faster than the guardrails.

Autonomous AI agents are racing into the enterprise, but oversight is lagging. Nearly 3 in 4 (71%) financial services companies plan to deploy agentic AI within two years. Yet, less than a quarter of respondents (23%) report having a mature model for governance of autonomous agents, raising the specter of unintended risks. Overall, the current level of at least moderate agentic AI adoption in the financial services industry (21%) is lower than the broader pattern across industries and is slightly below the cross-industry average (23%).

Physical AI is emerging in financial services—small today, scaling next.

Physical AI in financial services is in an early stage, but momentum is building: 32% of surveyed companies report using it today to some extent, and adoption is projected to reach 60% within two years. Even with this expected near-doubling, financial services remains the lowest-adopting industry—likely because it has fewer inherently physical processes than manufacturing, industrials, and health care.

71%

Nearly 3 in 4 financial services companies plan to deploy agentic AI within two years.

Overview

Leaders feel more strategically ready for AI than operationally ready in infrastructure and talent.

Despite the rapid evolution of AI beyond generative AI (GenAI) to agentic and physical AI, 47% of financial services companies believe their strategy is highly prepared for AI adoption and 32% say the same about risk and governance, both rising since last year. Perceptions of preparedness shifted down slightly for technical infrastructure and talent, revealing the persistent challenge of modernizing systems and skills at the speed of innovation.

47% of financial services companies believe their strategy is highly prepared for AI adoption.

About the Annual State of AI in the Enterprise Report

This annual report was fielded to 3,235 director-level to C-suite-level respondents across six industries and 24 countries between August and September 2025. Industries included: consumer; energy, resources, and industrials; financial services; life sciences and health care; technology, media, and telecom; and government and public services.* The survey data was augmented by additional insights from 15 interviews with global C-suite executives and AI and data science leaders at large organizations across a range of industries. For details on methodology, please see page 41. This annual report is part of an ongoing series by the Deloitte AI Institute™ to help leaders in business, technology and the public sector track the rapid pace of AI change and adoption.

[Learn more. →](#)

*Please note that government & public services survey data is not included in this analysis, given its smaller sample size and geographic focus on US organizations.

Key findings

AI is moving from the pilot and experimentation phase to enterprise scaling as worker access to AI expands

According to our latest survey, workforce access to AI in the financial services industry has doubled in just one year—growing from 30% to 62% of workers with sanctioned access to AI tools. Yet, among those workers with AI access, only 22% use it in their daily workflow more than


60% of the time. This suggests that while access is widening, enterprise AI remains underutilized, and its productivity and innovation potential are still largely untapped.

The scale acceleration is beginning

Moving from pilot to production is arguably the most important step in capturing AI value—yet this is where many companies stall. While enterprises are experimenting with AI at an accelerating pace, many struggle to scale these experiments into solutions that deliver measurable business impact.

Today, 24% of financial services industry respondents said their organization has moved 40% or more of their AI experiments into production to date. However, 53% expect to reach that level in the next three to six months, demonstrating the pathway to value is clear and achievable. These early scalers highlight an accelerated shift from pilots to enterprise-scale. The pace of progress for financial services companies aligns most closely with the consumer and energy, resources, and industrials industries (figure 1).

Figure 1: Organizations deploying 40%+ of AI experiments by industry (percent)



Question: In your estimation, what percentage of your AI experiments (e.g., pilots, test cases, etc.) have been deployed to date into your organization (moved into production)?

Key findings

The proof-of-concept trap

Why do so many pilots fail to reach production? The answer lies in a fundamental mismatch between pilot and production requirements. A pilot typically can run with a small team in a few months using cleansed data in an isolated environment. However, production deployment typically requires infrastructure investment, integration with existing systems, security reviews, compliance checks, monitoring systems, and ongoing maintenance—each of which demand significantly more resources and coordination.

In financial services, this gap is often magnified by the need to embed AI into highly regulated, data-intensive, and tightly integrated environments such as core banking platforms, trading systems, and claims and underwriting engines. Use cases such as fraud detection, credit decisioning, pricing, and customer service automation must

integrate with legacy core platforms, operate in near real time, and meet strict requirements for model risk management, data privacy, explainability, and regulatory compliance—often across multiple jurisdictions and lines of business.

Production deployment also surfaces challenges that pilots frequently obscure in financial services settings. Models that perform well on historical or curated datasets can degrade when exposed to live transaction flows, shifting customer behavior, adversarial activity, or changing market conditions. Data lineage gaps, inconsistent definitions across business units, and evolving regulatory expectations can undermine performance and slow approvals. Timelines that appear manageable in pilots can extend significantly once requirements for model validation, auditability, cybersecurity controls, and business sign-off are introduced. In production,



Key findings

failures are no longer theoretical—they can result in regulatory action, customer harm, financial losses, or reputational damage.

Organizations that experiment with AI often see positive results in controlled conditions but cannot consistently predict which use cases will yield the highest return on investment. This lack of clear value realization creates a vicious cycle where companies continue funding new pilots—which are relatively low cost and lower risk—rather than facing the harder work of scaling up existing successes.

The challenge of moving beyond isolated pilots to real enterprise impact is echoed by the people leading AI efforts on the ground across industries. According to one AI leader: “If there is no coherent AI strategy in organizations, you are likely to see pilot fatigue. You’re chasing the next shiny object, pressured to do something with AI without a real plan. I’ve seen many instances where people embark on pilots, but when asked how they’ll scale up if successful, they often don’t have an answer. Without a clear roadmap, executing a hundred pilots just leads to poor results and failed value creation.”



“If there is no coherent AI strategy in organizations, you are likely to see pilot fatigue.”

Key findings

AI transformation reveals productivity for most, business reimagination for a few

AI’s real-world business impact in the financial services industry is rising fast, with 85% of organizations increasing their AI investments and 76% of leaders expressing greater confidence in the technology. While 20% of leaders report that AI is having a transformative impact on their companies, many organizations are only at the edge of large-scale AI-driven transformation.


Beyond productivity and efficiency

AI is already delivering widespread gains in efficiency and productivity for financial services businesses; however, other types of benefits are taking longer to achieve. In particular, revenue growth largely remains an aspiration, with 75% of

financial services organizations hoping to grow revenue through their AI initiatives in the future compared to just 18% that are already doing so (figure 2). Financial services companies closely match industry averages in both the level of benefits achieved and their future aspirations, especially in the areas of productivity, innovation, and revenue.

These numbers suggest AI is on the verge of breaking out and delivering a wide range of benefits that go far beyond efficiency and productivity improvements. Ultimately, however, success with AI isn’t just about boosting efficiency or even growing revenue. It’s about achieving strategic differentiation and a lasting competitive edge in the marketplace.

Figure 2: AI benefits achieving today vs. hope to achieve in financial services (percent)




Question: With regards to benefits from your AI efforts: Which benefits are you achieving today? Which benefits do you hope to achieve?
N=573

The financial services industry’s approach to AI transformation differs from other industries.

Organizations are redefining how they work, but not all are diving to the same depth

Among the surveyed companies in the financial services industry, just over one-third (31%) are starting to use AI to deeply transform their businesses—creating new products and services, reinventing core processes, or even fundamentally changing their business models. Another 28% are redesigning key processes around AI but keeping their business models intact. And 41% are using AI at a more surface level, with little or no change to existing processes (figure 3). While each are capturing productivity and efficiency gains, only the first group are truly reimagining their businesses rather than optimizing what already exists.

Figure 3: Current approach to transformation with AI in financial services




Question: Which best describes your organization's current approach to process transformation during AI implementation?
N=573

Key findings

The financial services industry’s approach to AI transformation differs from other industries. Financial services organizations’ use of AI for deep transformation and process redesign aligns most closely with the consumer and life sciences and health care industries, while its surface-level AI use is the highest among industries (41%) (figure 4).

One company we interviewed took a bold, strategic approach to AI transformation by embedding AI into its core offerings and products. “AI is much more than a technology,” said the company’s head of AI and automation. “We wanted to give it to everyone for everyday usage and make it pervasive everywhere, but we also wanted to disrupt the market.” Rather than focusing solely on internal improvements, the organization is reimagining its business by using AI to unlock new value for both itself and its clients, and to create new digital solutions and revenue streams.

Figure 4: Current approach to transformation with AI by industry* (percent)



Question: Which best describes your organization's current approach to process transformation during AI implementation?
* Figures may not sum to 100% due to rounding.

Key findings

Companies are focused on building AI fluency instead of redesigning work around AI

Within a year, just under a third of surveyed financial services companies (31%) expect at least 10% of their jobs to be fully automated. Over the next three years, the vast majority (81%) expect at least 10% of their jobs to be fully automated.

These changes require careful thinking about career pathways. Leaders in the qualitative interviews expressed concerns about potential disruption to professional development pipelines as a result of automation. Entry-level jobs involving data entry, reconciliation, and first-level customer support at their companies are being prioritized for automation, but these jobs are often the starting point for longer careers. Organizations will likely need to develop alternate pathways for professional advancement, ensuring that employees have expertise that includes foundational processes.



31%

of surveyed financial services companies expect at least 10% of their jobs to be fully automated within a year.

Key findings

Most companies have yet to redesign jobs around AI

Despite high expectations for automation, 84% of financial services companies have not redesigned jobs around AI capabilities.

AI doesn't just augment existing processes. It often requires fundamentally rethinking operating models and how work gets done.² A credit or fraud operations manager who has traditionally relied on rules, manual reviews, and analyst judgment to approve transactions or flag risk must now work alongside AI systems that generate real-time risk scores, alerts, and recommended actions. This raises questions: *When should they override the AI? How do they document and justify those decisions to regulators and auditors? What happens to their professional judgment, accountability, and career trajectory as decision-making becomes increasingly automated?*

Entry-level and task-aligned roles could be most affected, as automation may replace common, time-consuming tasks. However, as front-line jobs become more automated, supervisor and managerial roles will likely shift toward orchestration of human-AI teams. This is prompting many financial services organizations to explore flatter structures: 49% have considered pod-based or non-hierarchical models since fewer roles require supervision of large teams; however, only 16% have moved to such models to a great or maximum extent.

Talent strategies are falling short

Worker sentiment about AI is mixed but cautiously positive. While 12% of non-technical workers in the financial services industry are highly enthusiastic about AI and are proactively seeking to use it—and 56% are at least open to exploring it—skepticism remains: 23% prefer not to use AI but will do so if required, and 3% actively distrust and avoid it.

Key findings


According to the leaders surveyed, insufficient worker skills are the biggest barrier to integrating AI into existing workflows. Yet, fewer than half of financial services companies are making significant adjustments to their talent strategies, with most (59%) simply focusing on educating employees to raise AI fluency (figure 5). While most are focused on educating employees, far fewer are re-architecting roles, workflows, and career paths. The financial services industry’s survey results in this area largely align with those

of the overall survey pool, skewing a bit higher on talent investments (e.g., AI fluency and upskilling).

A director of AI and innovation at a major logistics organization emphasized the importance of workforce upskilling and supporting business teams as they adapt to new AI technologies: “We are reskilling our people on the business side—investing a lot to ensure they adopt the new AI tools so they can deliver bigger, better, and smarter.” The company’s vision for workforce

transformation isn’t just limited to basic upskilling. Rather, the focus is on making sure employees can move from traditional roles into more strategic positions—supported by AI tools. “For example, in the future we would like to see AI enable today’s pricing analysts to become pricing strategists.”

Figure 5: Talent strategy adjustments due to AI in financial services (percent)



Question: How is your organization adjusting its talent strategies because of the adoption of AI tools/capabilities?
N=573

Key findings

With sovereign AI taking hold, *where* technology is built matters as much as *what* it can do

Sovereign AI is when a country—and the companies operating within it—design, train, and deploy AI under their own laws, on infrastructure they control, using locally governed data. The goal is to reduce dependence on foreign vendors for critical AI capabilities.

A new boardroom issue

More than 8 in 10 financial services companies (86%) view sovereign AI as at least moderately important to their strategic planning, and over half (57%) rate it as very important (figure 6). Financial services is in line with other industries on “moderately important” but is above the cross-industry average on “very important” (57% versus 43%).

Similarly, 65% of financial services companies express at least moderate concern about reliance on foreign-owned AI technologies and infrastructure, with 26% very concerned (figure 7).

Figure 6: Importance of data residency constraints and compute considerations in strategic planning in financial services*
(percent)




Figure 7: Level of concern over reliance on foreign-owned AI in financial services
(percent)



Key findings

Sovereign AI in practice

The rise of sovereign AI has immediate practical implications. Financial services companies working across borders must navigate complex requirements that vary by country, forcing them to build customized solutions for different markets. Eight in 10 financial services companies (80%) now factor an AI solution's country of origin into their vendor selection decisions, and more than half (53%) now build their AI stacks primarily with local vendors. This signals that geographic sovereignty is now as important as innovation.

As the former vice president of observability at a major global organization noted, "I've been working with a lot of international companies lately that are adamant we use an in-country infrastructure. We're taking the approach with some of those customers to do a distilled small language model for them that meets the import/export control rules because we build it in their

country. With state-run companies in particular, there is skepticism when you're using something from outside the country."

Sovereign AI pressures vary significantly by geography and industry. Only 11% of companies in the Americas rely on foreign-sourced solutions for the majority of their AI stack, compared to 32% of Europe/Middle East/Africa (EMEA) companies. In the US, state-level deployment considerations are increasingly shaping the landscape.

Ultimately, sovereign AI isn't just about technology ownership. It's about strategic independence. By building on infrastructure within its own control (fueled by its own data, models, talent, and ecosystem), a company has the ability to innovate securely and responsibly.



80%

of surveyed financial services companies now factor an AI solution's country of origin into their vendor selection decisions.

Key findings

AI agents are scaling faster than the guardrails

After years of non-AI chatbots that answered basic questions, companies are now deploying sophisticated AI agents that can set goals, reason through multi-step tasks, use tools and APIs, and coordinate work with people or other agents. This shift transforms AI from a source of information and insights into a system that could perform in a different capacity.


Agentic AI will surge

In last year’s State of AI survey, 26% of respondents in the financial services industry said their organizations were already exploring autonomous agent development to a large or very large extent. Those early exploration efforts are now starting to translate into real-world use and are poised to drive the sharp rise in adoption expected in the near term. Today, 21% of financial

services companies are using agentic AI at least moderately. However, within the next two years agentic AI is expected to become nearly ubiquitous in the industry, with nearly 3 in 4 financial services companies (71%) using it at least moderately, 24% using it extensively, and 5% fully integrating it as a core component of their operations (figure 8).

The current level of at least moderate agentic AI adoption in the financial services industry (21%) is lower than all other industries except energy, resources, and industrials. Over the next two years, adoption rates are expected to grow significantly with shared momentum across industries—aside from technology, media and telecom, which projects a notably higher rate.

Figure 8: Extent of agentic AI usage in financial services (percent)



Key findings

AI agents unlock a vast array of use cases

While agentic AI is expected to have the highest impact in customer support, use cases for supply chain management, R&D, knowledge management, and cybersecurity are also seen as having particularly high potential.³ The enterprises we interviewed across industries are already deploying AI agents across multiple functions:



A financial services company is building agentic workflows to automatically capture meeting actions from video conferences, draft communications to remind participants of their next steps, and track follow-through.



An air carrier is using AI agents to help customers complete the most common transactions, such as rebooking a flight or rerouting bags, freeing up time for human agents to address more complex matters.



A manufacturer is using AI agents to support new product development initiatives, leveraging AI to find the optimal balance between competing objectives such as cost and time-to-market.

91% of financial services companies expect to customize agents to fit the unique needs of their business.

Key findings

AI agents don't eliminate the value of humans

Adoption of agentic AI may increase the need for uniquely human strengths, such as adaptivity and judgement, in the near term. "We thought we were going to automate jobs," said the former VP of observability at a major global organization. "The truth is, you're not. You're going to give existing workers force multipliers where they can be more effective. Maybe someday these things will start to become headless where they just feed off a dashboard metric and you can pull back staff to wait on an alert that wakes

somebody up or flashes red on the screen if something really bad happens. But initially it is going to be more work for those people. They're not going to be cooling their heels; they're going to be watching these agents, making sure the volume metrics are right, making sure the qualitative metrics are right, and being there to interact with them if they hit a human-in-the-loop gate and need to interact with a human for accountability purposes."

"We thought we were going to automate jobs. The truth is, you're not. You're going to give existing workers force multipliers where they can be more effective."



Key findings

Managing the risks of agentic AI

Almost one-quarter (23%) of financial services companies surveyed report currently having a mature model for governance of autonomous agents. Given the technology's rapid adoption trajectory, this could be a significant limitation. As agentic AI scales from pilots to production deployments, establishing robust governance should be essential to capturing value while managing risk. For many financial services companies, the risks are amplified because agents could influence payments and funds movement, trading and portfolio actions, credit and underwriting decisions, customer communications and disclosures, and fraud/dispute actions—areas where errors can quickly drive customer harm, compliance breaches, model-risk findings, or material financial loss.

Governing agentic AI requires new approaches beyond traditional AI oversight. Unlike conventional AI systems that provide recommendations for humans to act upon, agents take actions directly—making purchases, sending communications, or modifying systems. Organizations need to establish clear boundaries for agent autonomy, defining which decisions agents can make independently versus which require human approval. Real-time monitoring systems that track agent behavior and flag anomalies are essential, as are audit trails that capture the full chain of agent actions to help ensure accountability and enable continuous improvement.



23%

of financial services companies report currently having a mature model for governance of autonomous agents.

Key findings

In financial services, this may include explicit guardrails around customer-impacting and regulated actions—such as initiating or approving payments and transfers, executing trades or rebalancing portfolios, changing credit decisions or limits, modifying customer communications and disclosures, and making privileged access or configuration changes in core banking systems—to help prevent unintended financial, model-risk, or compliance consequences.

Companies seeing the most success are taking a measured approach—starting with lower-risk use cases, building governance capabilities, and scaling deliberately. This includes cross-functional governance structures that bring together IT, legal, compliance, and business unit leaders to set policies, monitor performance, and manage escalations. Rushing to deploy agents widely before establishing these governance foundations can expose organizations to significant risks.



Companies seeing the most success are taking a measured approach—starting with lower-risk use cases, building governance capabilities, and scaling deliberately.

Key findings

A closer look at governance

AI governance as the catalyst for growth

As AI moves from experimentation to deployment, governance is the difference between scaling successfully and stalling out. In the AI era, governance is more than guardrails—it’s the catalyst for responsible growth.


The AI risks that financial services companies are most worried about all relate to governance. Data privacy and security tops the list at 80%, followed by legal, intellectual property, and regulatory compliance (55%), governance capabilities and oversight (48%), and model quality, consistency, and explainability (43%) (figure 9).

Insights gathered during qualitative interviews with AI leaders reveal that governance is a concern. In some organizations, leaders are learning AI models have been deployed

into production without formal oversight or monitoring processes. In particular, one AI leader discovered that there wasn’t a clear inventory of all the AI tools and models currently active since development efforts occurred without systematic tracking or centralized visibility over what is running in production.

Organizations that are serious about capturing AI value should treat governance as a strategic capability, not an afterthought. Those that build strong governance frameworks now will be positioned to scale AI quickly and safely. On the other hand, those that treat governance as a checkbox exercise may find themselves unable to move AI from pilot to production, held back by the very risks they failed to address.

Figure 9: AI risks most concerning in financial services (percent)



Question: Which of the following risks related to AI tools/applications is your organization most concerned about?
N=573


Key findings

Physical AI is emerging in financial services—small today, scaling next

Physical AI is the class of AI systems that perceive the real world, make decisions, and drive physical actions through machines or control systems. It sits at the intersection of AI and machine learning, sensors, controls, and robotics.

Physical AI integration in financial services is still early-stage, but momentum is building. Today, 32% of surveyed companies report at least limited use of physical AI, and among these, 8% are leveraging it to a moderate or greater extent. Over the next two years, adoption is expected to accelerate, with 60% of financial services companies using physical AI in some capacity—10% extensively and 1% fully integrating it as a core element of their operations (figure 10).

Figure 10: Extent of physical AI usage in financial services (percent)



Question: To what extent is your organization currently utilizing physical AI (e.g., robotics, automated machinery) in its operations?
N=573


Key findings

Financial services organizations report the lowest level of physical AI adoption—defined as moderate or greater use—among industries and are expected to remain the lowest in two years (figure 11). This low adoption is not surprising given that financial service companies have fewer inherently physical processes than industries such as manufacturing, industrials, and health care. However, physical AI can still be relevant in environments like offices, branches, ATMs, and data centers (e.g., security, safety, and facilities monitoring). It may also be directly relevant in niche use cases such as telematics-based monitoring and assessment of real-time driver behavior in auto insurance.

Even though financial services firms’ use of physical AI is expected to more than double in two years, its growth is notably slower than software-based agentic AI, which is on track to jump from 21% to 71% in the same timeframe.

The difference reflects the inherent challenges of physical deployment: higher costs and capital requirements, longer development cycles, stricter safety regulations, and the need for specialized hardware and maintenance.

Figure 11: Moderate to greater physical AI usage across industries: Now and in 2 years (percent)



Question: To what extent is your organization currently utilizing physical AI (e.g., robotics, automated machinery) in its operations?

Key findings

Controlled environments are leading the way

Physical AI applications span a wide range of industrial and commercial settings. For example, one company we interviewed is automating package sorting and routing while granting warehouse robots more autonomy to decide where and how to store items to maximize floorspace. Other common use cases include collaborative robots (cobots) on assembly lines, inspection drones with automated response capabilities, robotic picking arms, and autonomous forklifts. Adoption is especially advanced in manufacturing, logistics, and defense, where robotics, autonomous vehicles, and drones are already reshaping operations.

In the financial services industry today, physical AI adoption is largely limited to the back office where controlled conditions and minimal public interaction reduce operational and regulatory risks. Early use cases include cash handling operations—where robotic arms sort and bundle banknotes in secure cash centers—and autonomous security systems monitoring ATMs and facilities.



Physical AI adoption in financial services is largely limited to the back office where controlled conditions reduce operational and regulatory risk.


Key findings

Types of physical AI with the greatest expected impact

As physical AI gains broader adoption, certain types are expected to have a bigger long-term impact than others in the financial services industry—most notably, intelligent security systems and/or smart monitoring (41%), followed by digital twins (17%) and IoT-driven retail (14%) (figure 12).

These use cases may rise to the top because they align with financial services’ most critical operational priorities: physical and cybersecurity, regulatory compliance, risk management, and operational efficiency in controlled environments. Intelligent security systems help protect branches, ATMs, cash centers, and data facilities while enabling real-time threat detection and response. Digital twins allow institutions to model risk scenarios, simulate operational changes, test new systems, and optimize resource allocation

Figure 12: Types of physical AI expected to have the greatest impact in financial services (percent)



Question: Which area of physical AI do you believe will have the greatest impact on your industry?
N=573

Key findings

before deployment in live environments. IoT-driven retail enhances branch operations and customer experiences through connected devices and environmental monitoring, while autonomous logistics and collaborative robotics support back-office functions like cash handling and document processing where reliability and security can be tightly controlled.

To help ensure public acceptance and trust of these solutions, financial services companies should focus on making them secure, interoperable, and resilient against disruptions and cyberthreats.⁴ Unlike software AI that operates in digital environments, physical AI

systems may interact with people, equipment, and infrastructure in ways that can pose safety risks if they malfunction or are compromised. This raises the stakes for testing, certification, and ongoing monitoring.

Financial services companies must also contend with complex regulatory environments that vary by industry and geography. Physical AI systems often require approval from safety regulators, compliance with industry-specific standards, and adherence to liability frameworks that don't always account for autonomous systems. These requirements add time and cost to deployment but are essential for responsible adoption.

Financial services companies must also contend with complex regulatory environments that vary by industry and geography.

Key findings

Accounting for the full cost of physical AI

In our survey, cost was cited most often as a key barrier to physical AI deployment. When evaluating business cases for physical AI, decision makers in financial services organizations should account for total cost of ownership and not just initial equipment costs. The full picture includes facility retrofits to accommodate new equipment, sensors and robots themselves, integration with existing systems and workflows, maintenance and spare parts, and potential downtime during implementation and troubleshooting.

These costs can significantly exceed the initial investment in AI models and software. A warehouse automation project might require hundreds of thousands of dollars in AI development but millions of dollars in physical infrastructure, robotic systems, and facility modifications. Companies that underestimate these costs risk project delays or abandonment partway through implementation.


When evaluating business cases for physical AI, decision makers in financial services organizations should account for total cost of ownership and not just initial equipment costs.

Key findings

Leaders feel more strategically ready for AI than operationally ready in infrastructure and talent

Despite the rapid evolution of AI beyond GenAI to agentic and physical AI, 47% of financial services companies believe their strategy is highly prepared for AI adoption and 32% say the same about risk and governance, both rising since last year's report (+4 and +7 percentage points, respectively) (figure 13). These areas have likely advanced more quickly as they depend primarily on executive decision-making and policy development.

Figure 13: Level of preparedness for AI adoption in financial services (percent)



Question: For each, rate your organization's level of preparedness with respect to broadly adopting AI tools/applications.
N=573
● Percentage point change vs last year

Key findings

Meanwhile, perceptions of high preparedness have shifted down compared with last year for technical infrastructure (38%) and talent (19%), revealing the persistent challenge of modernizing systems and skills at the speed of innovation. In fact, most financial services respondents believe that resolving the key challenges for their organization's priority AI initiatives will take more than a year—far too long in today's fast-moving environment.

As a head of AI strategy we interviewed explained: "Many organizations prepared for an AI future by building infrastructure and governance for traditional AI models. With LLMs, those efforts were upended. Suddenly, there was a new capability unlike previous AI. Now, traditional AI use cases—training models from scratch, custom interfaces—have diminished. Nearly 80–90% of new use cases are generative AI. So yes, companies are prepared, but for a different future. GenAI needs a new set of capabilities."

"Nearly 80–90% of new use cases are generative AI. So yes, companies are prepared, but for a different future. GenAI needs a new set of capabilities."

Tapping into AI's full potential

The research is telling: AI's transformational potential is real but capturing it requires far more than just technology investments. Financial services companies should treat AI as foundational. The most successful won't be those with the most AI projects or the biggest budgets, but those who build AI into the foundation of how they operate, compete, and grow.

Here are six key focus areas to help your business capture AI's untapped edge:



Close the gap between access and activation

Most organizations have deployed AI tools, but far fewer have achieved meaningful usage. The gap between availability and adoption is now the primary barrier to value. Successful companies focus on activation, not just access.

High-performing implementations start with empowered employees who experiment, share early wins, and become internal champions. Top-down directives alone rarely drive meaningful change. Grassroots adoption supported by senior sponsorship creates momentum and helps ensure solutions align with real workflows.

Activation requires early attention to practical constraints: system integration, data permissions, and operational reliability. For organizations applying AI not only to digital processes but also to physical systems—such as robotics, IoT devices, or machinery—early planning for these operational realities is especially critical. Organizations that design for deployment from the outset, rather than treating scale as an afterthought, see far higher adoption. Hands-on, role-specific training and visible executive advocacy materially shift employee behavior. Leaders that treat pilots as stepping stones to production, not isolated experiments, are likely to achieve faster and more durable impact.



Unlock human advantage by redesigning work around AI

AI is reshaping work at every level. While most organizations currently focus on personal productivity, leaders are rebuilding processes, roles, and career paths around expanded AI capabilities.

The most successful organizations reimagine jobs to seamlessly combine human strengths and AI capabilities, ensuring both aspects are used to their fullest potential. New roles—AI operations managers, human-AI interaction specialists, quality stewards, and others—signal a deeper shift: AI is now a structural component of how work is organized. Advanced organizations streamline workflows that AI can execute end-to-end, while humans focus on judgment, exception handling, and strategic oversight. The goal isn't

to replace humans or merely assist them, but to create complementary working relationships between humans and AI, in which the combined output exceeds what either could achieve alone.

Organizational structures are beginning to flatten as AI absorbs routine execution tasks. Some companies are merging technology and people-leadership functions to ensure that systems and workforce design evolve together. The pace varies by industry, but the direction is consistent: Roles, skills, and career paths should be rebuilt, not simply adjusted. Organizations should take an AI-native approach and redesign work holistically rather than layering AI onto legacy processes.



Build governance before you scale and make it everyone's role

Governance is no longer a compliance exercise; it's the mechanism that enables rapid, confident scaling. Enterprises where senior leadership actively shapes AI governance achieve significantly greater business value than those delegating the work to technical teams alone. True governance makes oversight everyone's role, embedding it into performance rubrics so that as AI handles more tasks, humans take on active oversight. This shared responsibility empowers employees to help identify challenges and guide safe, trusted AI use.

Effective governance integrates with existing risk and oversight structures, not parallel "shadow" functions. It focuses on identifying high-risk applications, enforcing responsible design practices, and ensuring independent validation where appropriate. Leading organizations proactively monitor evolving legal requirements and build systems that can demonstrate safety, fairness, and compliance.

Autonomous systems heighten needs for data and cybersecurity governance.⁵ Organizations need to define where humans should remain in control, how automated decisions and data use are audited, and which records of system behavior should be retained. Cross-functional teams—technology, legal, compliance, and business—establish governance frameworks early so that scale does not outpace control. At the same time, governance should be calibrated to balance risk management with innovation, ensuring that oversight enables experimentation rather than constraining it. The objective is not to add bureaucracy but to create clear, adaptive guardrails that allow responsible progress at speed.



Address sovereign AI requirements with focus and discipline

As national governments accelerate efforts to establish sovereign AI capabilities, enterprises will likely navigate increasingly complex expectations around data control, model transparency, compliance, and localization. Sovereign AI is no longer limited to the public sector; it is reshaping requirements for every organization handling sensitive data or operating across jurisdictions. At the same time, compute strategy becomes a core component, requiring careful evaluation of both data residency and processing locations (e.g., in cloud, on-prem, hybrid, or edge environments) to remain aligned with evolving regulations and performance needs.

Leading companies take a focused approach: assessing which data and workloads must remain within national or regional boundaries, determining where local model hosting is mandatory, and clarifying how transparency, auditability, and documentation standards differ across markets. They establish clear policies for data residency, model retraining, and cross-border flows, supported by infrastructure capable of meeting multiple regulatory regimes simultaneously.

Enterprises that ignore sovereign AI constraints will face escalating operational disruption, higher compliance risk, and restricted access to key markets. Those that proactively engage build strategic advantage: They can reduce regulatory uncertainty, enhance customer trust, and position themselves as preferred partners in industries where sovereignty concerns dominate. Sovereign AI readiness is now a core element of enterprise resilience and global competitiveness, not a specialized compliance task.

Tapping into AI's full potential



Build a “living” technology and data infrastructure for tomorrow’s AI

Legacy data and infrastructure architectures cannot power real-time, autonomous AI. As AI capabilities extend beyond software into devices, machinery, and edge locations, organizations need to evaluate if their technology foundations are ready to support potential physical AI deployments. Modernization should create a living AI backbone: an organization-wide, real-time system that adapts dynamically to business and regulatory change, elevating infrastructure from IT initiative to strategic capability.

Leaders are enabling modular, cloud-native platforms that securely connect, govern, and integrate all data types, fostering rapid experimentation and seamless scaling. They break down silos with domain-owned data products and embed privacy, sovereignty, and

security-by-design, while enforcing enterprise standards for quality, interoperability, and lineage. This balanced approach delivers decentralized innovation supported by centralized control.

A unified, trusted data strategy is indispensable. Poor or fragmented data compounds risk and undermines every AI initiative. Forward-thinking organizations converge operational, experiential, and external data flows and invest in evolving platforms that anticipate the needs of emerging AI.⁶ Infrastructure determines enterprise velocity; those that modernize early will likely accelerate while others remain constrained.





Pursue strategic reinvention, not incremental efficiency

A widening performance divide separates companies treating AI as core to strategy from those viewing it as a cost-saving tool. Leading organizations invest heavily in using AI to reshape operations and create new revenue streams, resisting the pressure to chase every trending technology in favor of initiatives that genuinely advance strategic goals and deliver real value.

These organizations pursue growth across multiple horizons: strengthening current operations, expanding into adjacent markets, and building entirely new businesses enabled by AI. They rethink their organizations from the ground up and imagine how to build without legacy constraints, rather than digitizing old processes. This extends to reimagining business models and adapting to emerging trends like sovereign AI. This

intentional reinvention is one of the strongest predictors of achieving outsized returns.

Autonomous AI systems are accelerating this shift. In knowledge-intensive industries, they can absorb substantial routine work, enabling people to focus on higher-order activities. High performers are reorganizing around systems that perceive context, make decisions, and act independently, balancing bold transformation with operational continuity. They move at a pace suited to their organization's readiness, making thoughtful trade-offs and fostering informed decision-making grounded in evidence rather than hype. The strategic opportunity is discovering new sources of value that competitors cannot easily replicate.

With developments in agentic, physical, and sovereign AI rapidly expanding the boundaries of what's possible, financial services organizations today are at the edge of tapping into AI's full potential. Whether it's figuring out how to capitalize on the latest cutting-edge innovations, making the leap from pilots to large-scale deployment, or using AI to create an enduring competitive advantage, enterprises around the world are on the edge of transforming themselves with AI. The challenge now is activation: bridging the gap from tool access to meaningful adoption, moving beyond experimentation to operationalizing AI at scale, embedding AI into core business processes—transforming technology potential into enterprise value.

Acknowledgements

Authors

Business leadership



Jim Rowan
US Head of AI
Deloitte Consulting LLP
jimrowan@deloitte.com



Beena Ammanath
Executive Director
Global Deloitte AI Institute
Deloitte LLP
bammanath@deloitte.com



Nitin Mittal
Principal
Global AI Leader
Deloitte Consulting LLP
nmittal@deloitte.com



Costi Perricos
Global GenAI
Business Leader
Deloitte LLP
cperricos@deloitte.co.uk

Continue the conversation

Industry leadership



Neal Baumann
Global Financial Services
Industry Leader
Deloitte Consulting LLP
nealbaumann@deloitte.com



Ryan Lockard
US Financial Services
AI Leader
Deloitte Consulting LLP
rlockard@deloitte.com

Acknowledgements

Special thanks

Technology and research team



Vivek Kulkarni
Managing Director
US AI Transformation Leader
Deloitte Services LP
vivkulkarni@deloitte.com



Lisa Hohener
Senior Manager
Deloitte AI Institute
Deloitte Consulting LLP
lhohener@deloitte.com



Caroline Ritter
Manager
US AI Transformation
Deloitte Services LP
carritter@deloitte.com

Special thanks

Global AI subject matter expert network

The authors would like to thank the Deloitte subject matter specialists who contributed to the development of the survey and report:

Javier Echaniz Alonso, Bjoern Bringmann, Lou DiLorenzo, Kyle Forrest, Marcus Goetz, Cliff Goss, Scott Holcomb, Jas Jaaj, Saurabh Kumar, Anjani Kumar Sripada, Patrick Laurent, Chris Lewin, Parth Patwari, Antonio Russo, Prakul Sharma, Rohit Tandon, Ed Van Buren, Greg Vert, Blaine Woodcock.

Additional contributors

The authors would also like to thank our team of professionals who brought this report and campaign to life, including:

Stephanie Anderson, Andrew Ashenfelter, Hali Austin, Hannah Caslin, Vanessa Carney, Abigail Claiborne, Victoria Estacio, Behzad Eftekhari, Robert Fabbro, Alex Frole, Bharathi Gabbita, Karen Hogger, Michael Holt, Lisa Iliff, Anish Indukur, Bavapriyan K, Omar Khan, Ganapathy Subramaniam Kumar, Amy Lando, Fiona Maguire, Cullen Marriott, Meredith Mazzotta, Pamela Zamora Miranda, Hye Ra Moon, Melissa Neumann, Inal Olmez, Kendall Phillips, Jose Porras, Priyadarshini, Arockia Pulavar, Preetham Raghu, Emily Rosenberg, Kate Schmidt, Brenna Sniderman, Lesley Stephen, Tina Trabucco, Olivia Tytla, Vaibhav Vasireddy, Ivana Vucenovic, Emilie Kuo Yang.



About the Deloitte AI Institute

The Deloitte AI Institute™ helps organizations connect all the different dimensions of the robust, highly dynamic, and rapidly evolving Artificial Intelligence ecosystem. The AI Institute leads conversations on applied AI innovation across industries, with cutting-edge insights, to promote human-machine collaboration in the “Age of With™.”

The Deloitte AI Institute aims to promote the dialogue and development of AI, stimulate innovation, and examine challenges to AI implementation and ways to address them. The AI Institute collaborates with an ecosystem composed of academic research groups, start-ups, entrepreneurs, innovators, mature AI product leaders, and AI visionaries to explore key areas of artificial intelligence including risks, policies, ethics, the future of work and talent, and applied AI use cases. Combined with Deloitte’s deep knowledge and experience in artificial intelligence applications, the Institute helps make sense of this complex ecosystem, and as a result, delivers impactful perspectives to help organizations succeed by making informed AI decisions.

No matter what stage of the AI journey you are in: whether you are a board member or a C-Suite leader driving strategy for your organization—or a hands-on data scientist bringing an AI strategy to life—the Deloitte AI Institute can help you learn more about how enterprises across the world are leveraging AI for a competitive advantage. Visit us at the Deloitte AI Institute for a full body of our work, subscribe to our podcasts and newsletter, and join us at our meet-ups and live events.

Let’s explore the future of AI together. →

Methodology

To obtain a global view of how AI is being adopted by organizations on the leading edge of AI, Deloitte surveyed 3,235 leaders between August and September 2025. Respondents were senior leaders in their organizations and included board and C-suite members, and those at the president, vice president, and director levels. The survey sample was split equally between IT and line-of-business leaders. Twenty-four countries were represented: United States (n=1,200), Canada (n=175), Brazil (n=115), Mexico (n=100), United Kingdom (n=220), Germany (n=170), France (n=150), Netherlands (n=50), Italy (n=75), Spain (n=100), Austria (n=50), Portugal (n=15), Poland (n=40), Czechia (n=30), Saudi Arabia (n=70), United Arab Emirates (n=130), Egypt (n=45), Kuwait (n=10), Qatar (n=10), Oman (n=5), Japan (n=100), India (n=200), Singapore (n=75), Australia (n=100). All participating organizations have one or more working implementations of AI being used daily. Plus, they have pilots in place to explore AI or have one or more working implementations of being used daily. Respondents were required to meet one of the following criteria with respect to their organization's AI and data science strategy, investments, implementation approach, and value measurement: influence decision-making, are part of a team that makes decisions, are the final decision-maker, or manage or oversee AI technology implementations. All statistics noted in this report and its graphics are derived from Deloitte's annual survey, conducted between August and September 2025; *The State of AI in the Enterprise* report series. N (Total leader survey responses) = 3,235. The survey data was supplemented with case studies and qualitative findings derived from 15 interviews with executives and AI and data science leaders at large organizations across a range of industries and countries.

Endnotes

1. Michelle Gauchat, Val Srinivas, 2026 Banking and Capital Markets Outlook, October 30, 2025, <https://www.deloitte.com/us/en/insights/industry/financial-services/financial-services-industry-outlooks/banking-industry-outlook.html>, accessed January 13, 2025.
2. Sue Cantrell, David Mallon, Aniket Bandekar and Simona Spelman, "The path to achieving value from AI: Scaling your human edge," Deloitte Insights2Action", October 2025, <https://action.deloitte.com/insight/4740/the-path-to-achieving-value-from-ai-scaling-your-human-edge>, accessed December 12, 2025.
3. Nitin Mittal, Beena Ammanath and Jim Rowan, "The AI Dossier," The Deloitte AI Institute, <https://www.deloitte.com/content/dam/assets-zone3/us/en/docs/services/consulting/2025/ai-institute-ai-use-cases-10-06-2025.pdf>, accessed December 12, 2025.
4. Mark Osis, Raquel Buscaino, and Caroline Brown, "Robotics & Physical AI: Intelligence in motion," Deloitte, November 2025, <https://www.deloitte.com/content/dam/assets-zone3/us/en/docs/about/2025/robotics-and-physical-ai-tech-futures-report.pdf>, accessed December 12, 2025.
5. Sunny Aziz, Adnan Amjad, Naresh Persaud, Mark Nicholson, and Ed Burns, "Tech Trends 2026: The AI dilemma: Securing and leveraging AI for cyber defense," Deloitte, December 2025, <https://www.deloitte.com/us/en/insights/topics/technology-management/tech-trends/2026/using-ai-in-cybersecurity.html>, accessed December 12, 2025.
6. Vivek Kulkarni, Scott Holcomb, Prakul Sharma, Greg Vert and Caroline Ritter, "Agentic enterprise 2028: A blueprint for cost savings, job creation, and faster growth through agentic AI," Deloitte, September 2025, p. 14, <https://www.deloitte.com/us/en/what-we-do/capabilities/applied-artificial-intelligence/articles/agentic-ai-enterprise-2028.html>, accessed December 12, 2025.



As used in this document, “Deloitte” means Deloitte Consulting LLP, a subsidiary of Deloitte LLP. Please see www.deloitte.com/us/about for a detailed description of our legal structure. Certain services may not be available to attest clients under the rules and regulations of public accounting.

This publication contains general information only and Deloitte is not, by means of this publication, rendering accounting, business, financial, investment, legal, tax, or other professional advice or services. This publication is not a substitute for such professional advice or services, nor should it be used as a basis for any decision or action that may affect your business. Before making any decision or taking any action that may affect your business, you should consult a qualified professional advisor. Deloitte shall not be responsible for any loss sustained by any person who relies on this publication.

Copyright © 2026 Deloitte Development LLC. All rights reserved.