Kapronasia Airwallex SFA SINGAPORE FINTECH ASSOCIATION



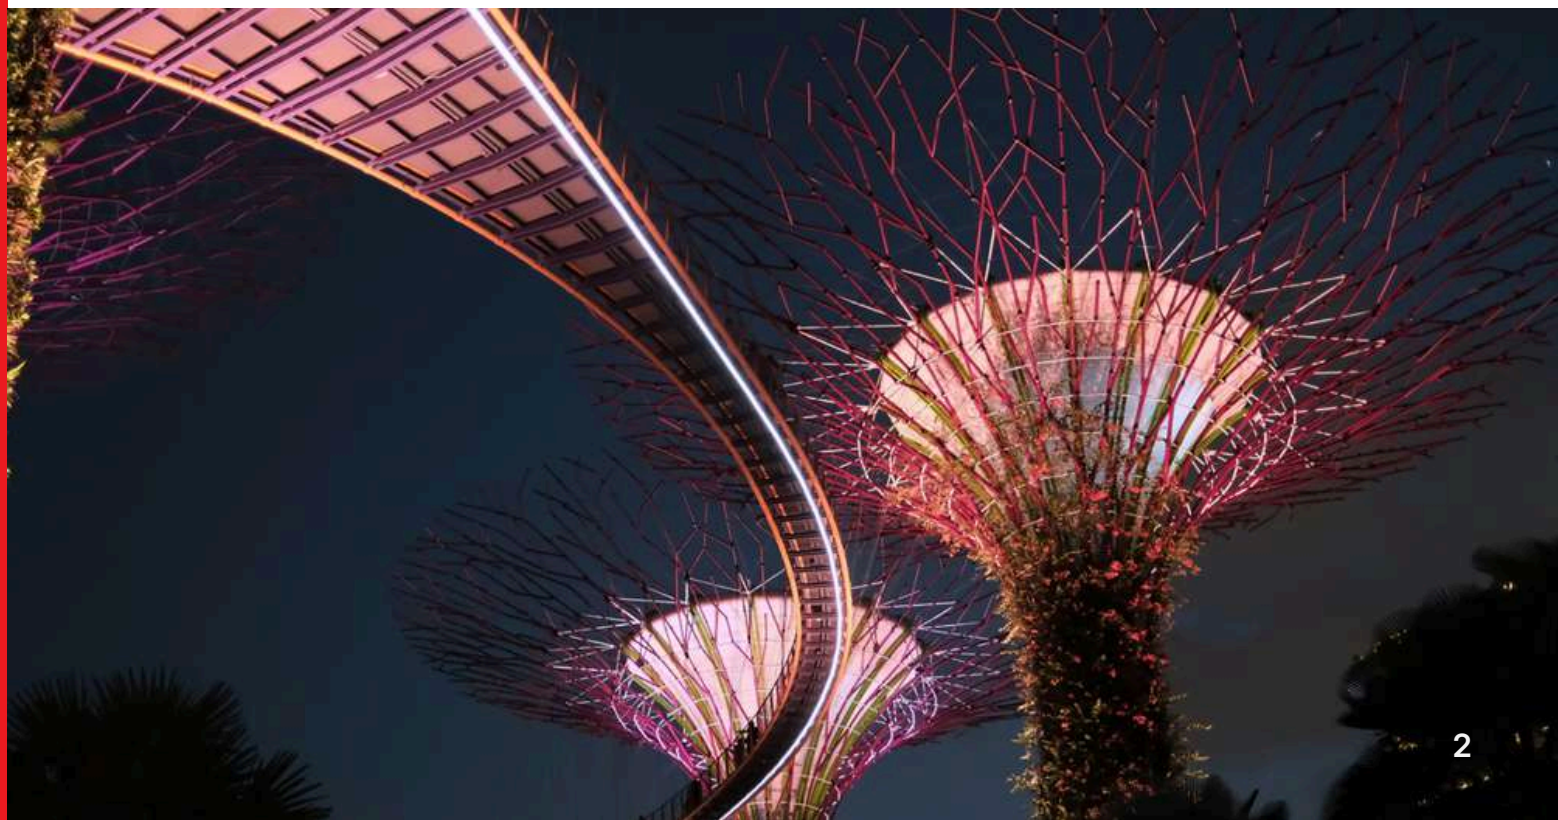
The AI Inflection Point: Rewiring Fintech

A research brief from Kapronasia in collaboration with Airwallex and SFA

November 2026

Contents

Foreword 3
Executive summary 4
AI inflection point and adoption maturity 6
Operational impact and high-value use cases 11
Agentic AI and the future of finance 14
The trust equation: Governance, Risk, and Internal Friction 17
Our framework 19
Conclusion 22
Acknowledgments 23



Foreword

AI is moving rapidly from experimentation into the way financial services are designed and delivered. For fintechs, this shift presents both an opportunity and a challenge: the question is no longer whether to adopt AI, but where it can deliver meaningful value, how deeply it can be embedded into the business, and how much work it can ultimately take on.

Singapore's fintech ecosystem is well placed to navigate this transition. Its digital-first businesses, strong technology capabilities, and close engagement with regulators provide a strong environment for experimenting with new AI applications. At the same time, financial services place a high bar on reliability, accountability, and customer trust, particularly as AI moves from assisting people towards taking actions on their behalf.

This research examines where Singapore's fintech sector stands in this transition. Drawing on a survey of 100 industry leaders and interviews with senior practitioners, it looks beyond headline adoption rates to assess how deeply AI is being embedded, where it is delivering value, and what is holding firms back from scaling further.

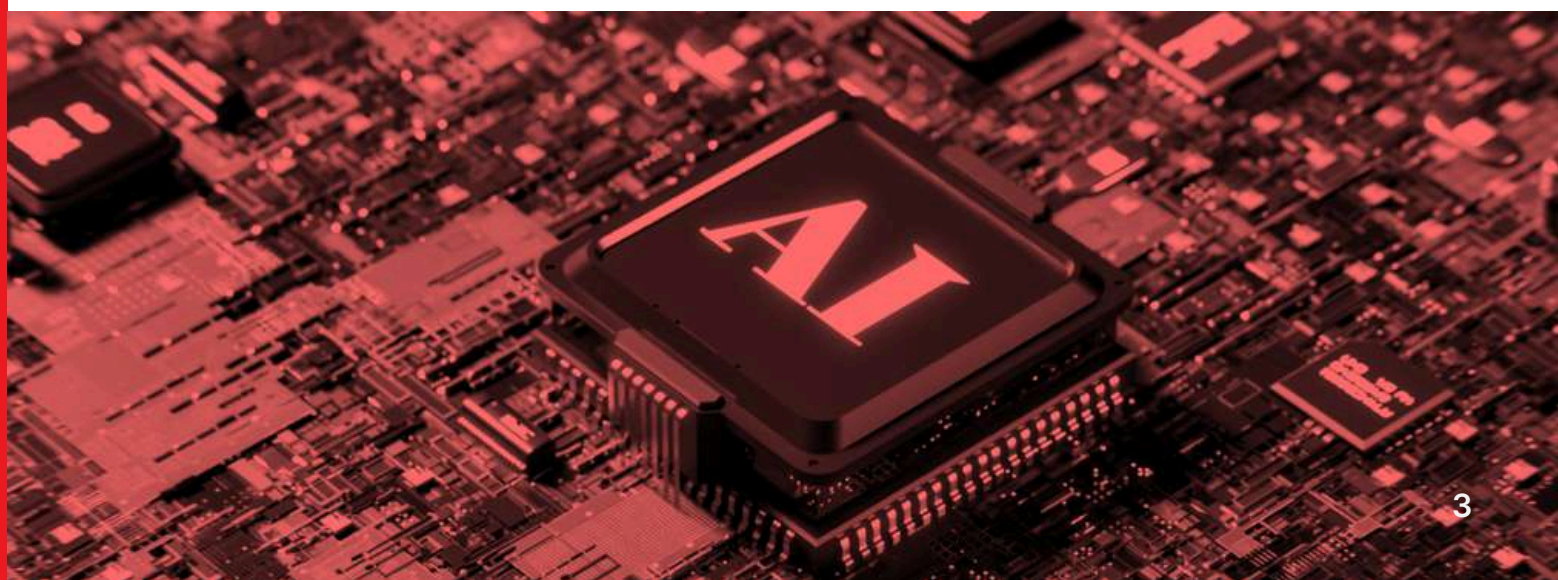
We hope this research provides a practical perspective for fintech leaders navigating that transition and helps inform the decisions needed to move from AI experimentation towards meaningful, scalable, and responsible adoption.



Arnold Chan
General Manager, Asia-Pacific
Airwallex



Tawishi Singh
Vice President
Singapore FinTech Association



Executive summary

Across financial services today, AI adoption has moved from a competitive advantage to an industry norm. Nearly all Singapore financial institutions and global financial institutions have either deployed or are piloting and experimenting with AI. The differentiator is no longer whether firms use AI, but how fast they are able to scale AI beyond simply pilots and into more advanced workflows and tasks.

Singapore's fintechs are at the forefront of this AI shift, with 86% already deploying AI, compared with 65% of Singapore financial institutions and 62% of global fintechs.[1,2,3] However, high adoption does not necessarily translate into greater AI maturity. While firms are already actively using AI, the percentage of firms that have moved beyond initial adoption to scale AI across the organization remains limited.

Increasing AI maturity matters because it impacts how much value firms can capture from their investments. Yet measuring that value remains challenging. The gap between expected and actual returns highlights the limits of traditional ROI metrics, which may not fully capture the value AI creates within specific functions and workflows. Firms therefore require a more nuanced understanding of how to measure the ROI of AI, beyond general revenue and cost metrics. Embedding AI into high-value workflows and linking deployment to measurable business outcomes can help firms better capture and demonstrate its impact.

As AI maturity advances, the opportunity extends beyond assistive use cases. Agentic AI systems can execute multi-step processes, use tools, and adapt to changing conditions with lower degree of human oversight required. At the same time, firms remain cautious in deploying AI into live systems and workflows, as greater autonomy raises the requirements for governance, data, talent, technology, and trust.

The opportunity is therefore to scale AI with the right balance of autonomy and control. This report proposes a practical framework to help firms move beyond siloed experimentation towards broader, AI-native deployment, with the foundations needed to deliver measurable value and trusted autonomy.

[1] Finastra (2026) Singapore banks surge ahead on AI deployment, modernization, and customer experience, according to Finastra Research. Available at: https://www.finastra.com/press-media/singapore-banks-surge-ahead-ai-deployment-modernization-and-customer-experience

[2] Kapronasia primary survey

[3] NVIDIA (2026). State of AI in Financial Services: 2026 Trends. Available at: https://www.nvidia.com/en-sg/industries/finance/ai-financial-services-report/



Background and methodology

To understand the current state of AI adoption in Singapore fintechs, Airwallex, The Singapore FinTech Association (SFA), and Kapronasia conducted a mixed-methods study to examine its state among Singapore fintechs, including where AI is delivering impact and what is constraining further adoption.

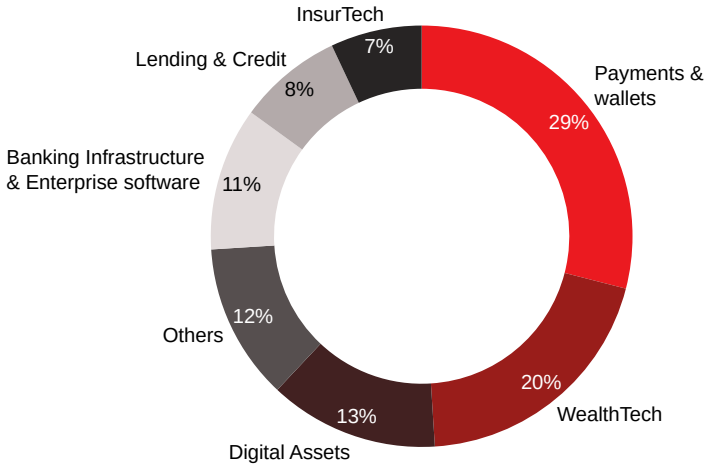
Research methodology

Our mixed-methods study consists of a quantitative survey and qualitative interviews with senior leaders in the industry. The goal was to assess the AI readiness of Singapore fintechs in terms of AI adoption, maturity, business impact, barriers to scale, and regulatory alignment across the ecosystem, with fieldwork conducted from July to September 2026.

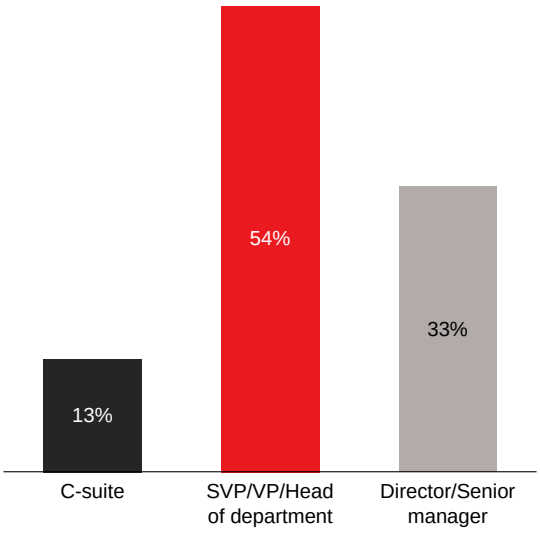
The quantitative survey features 100 respondents holding leadership positions across various functional areas in the fintech sector—including Payments & wallets, WealthTech, and Digital Assets—ensuring a balanced perspective on both strategic priorities and operational realities (Figure 1). In addition, qualitative interviews with 10 senior leaders were conducted to further understand the findings from the quantitative surveys.

Figure 1

Functional area of respondents, n=100



Seniority levels of survey respondents, n=100





AI inflection point and adoption maturity

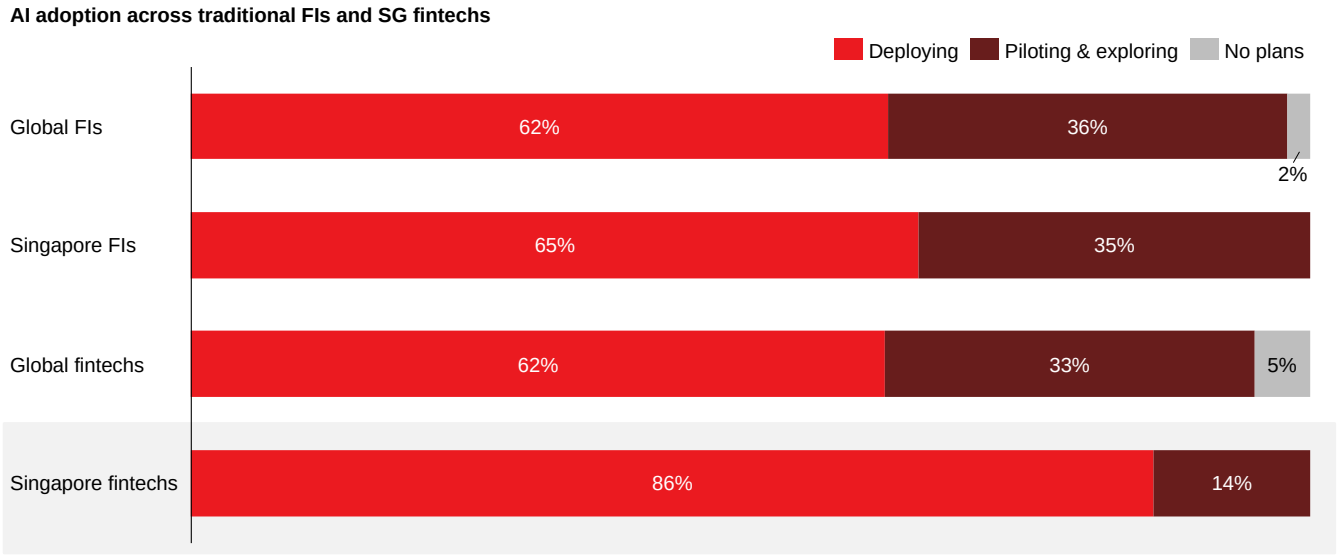
Singapore fintechs are ahead of both domestic and global peers in AI deployment, reflecting their digital-first operating models, streamlined infrastructure and greater organizational agility. With technology embedded deeply in how these firms operate and serve customers, many have been well positioned to adopt AI across their operations.

This momentum is reflected in adoption levels, with 86% of Singapore fintechs already deploying AI, compared with 65% of Singapore financial institutions (FIs) and 62% of both global FIs and fintechs (Figure 2). AI adoption is therefore no longer limited to experimentation, but is increasingly becoming part of day-to-day operations across the sector.

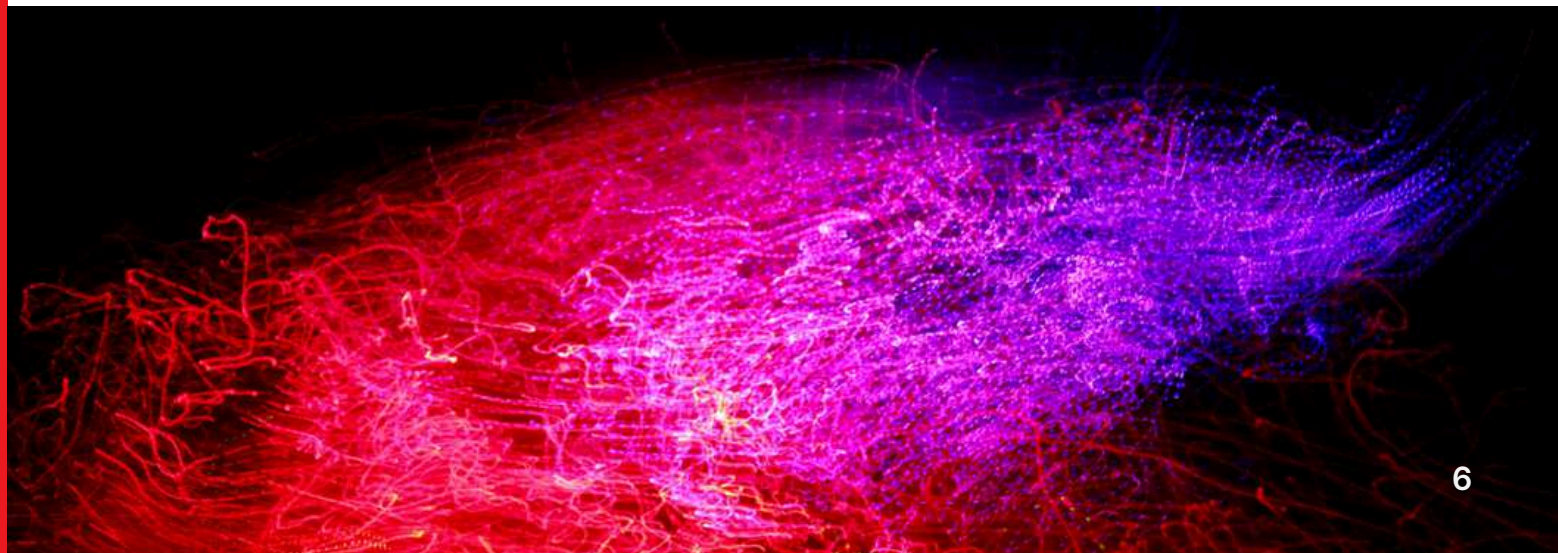
However, widespread deployment is only one measure of progress. As adoption becomes more widespread, the focus is shifting from simply deploying AI to building the organizational maturity needed to capture its full value. This means assessing not only whether firms are using AI, but how effectively they are translating adoption into meaningful and sustained organizational impact.

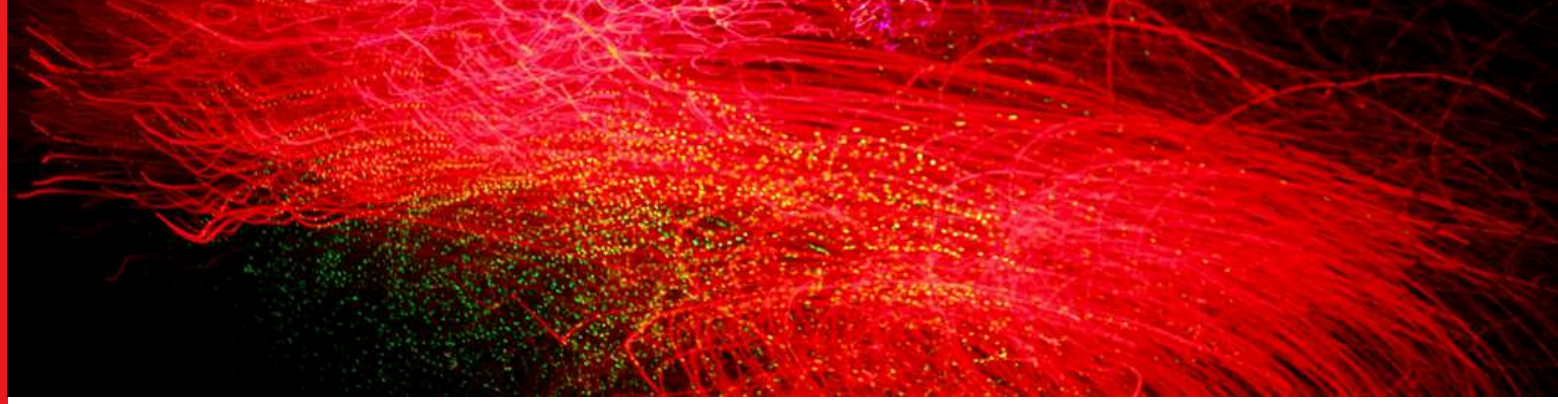
Figure 2

Singapore fintechs lead traditional Financial Institutions and Global fintechs in AI adoption, but adoption alone is no longer a differentiator



Source: Finastra, NVIDIA, Kapronasia primary survey data





To assess how effectively Singapore fintechs are translating AI adoption into organizational impact, we examine AI maturity across five dimensions:

- Workforce adoption of AI: How much of workforce is using AI in day-to-day activities?
- AI adoption breadth: How many functions have adopted AI?
- AI integration depth: What % of AI used are standalone tools vs embedded in core functions?
- AI autonomy: To what degree are these AI deployed chatbots vs copilots vs agents?
- Infrastructure and data readiness: How ready are the key foundational levers for scaling AI?

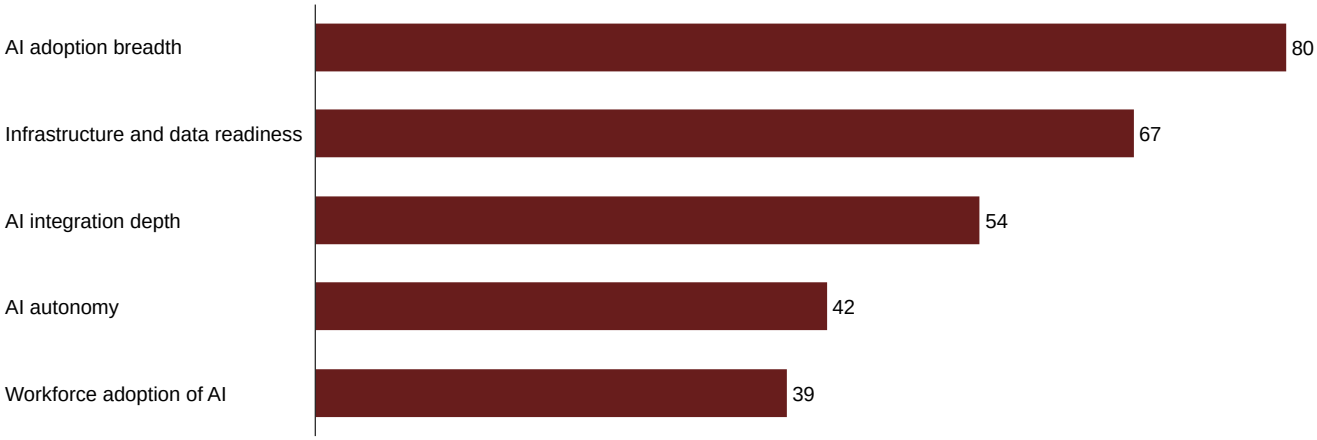
The results reveal uneven progress: fintechs score highest in AI adoption breadth (80) and infrastructure and data readiness (67), but lower in AI integration depth (54), autonomy (42) and workforce adoption (39) (Figure 3). This suggests that while firms have established strong foundations for AI, embedding it into core workflows and driving consistent workforce use remain key areas for development.

The next challenge is to close this gap between readiness and implementation. Sustaining progress will require fintechs to strengthen the processes, governance and workforce capabilities needed to embed AI into everyday work and realize its long-term value. The following sections examine each dimension in detail.

Figure 3

Singapore fintechs have built the foundations for AI, but not yet embedded it

Singapore fintech maturity score (0 – 100) – five dimensions, n=100



Note: Scores are index values from 0 to 100 calculated using Kapronasia's AI maturity scoring methodology, not the share of firms (n=100). Each dimension is scaled so that the highest response option scores 100. Workforce adoption: average position on the share-of-workforce-using-AI scale. Adoption breadth: average share of 12 business areas where AI is deployed. Integration depth: average level of integration (standalone to end-to-end) across areas with AI deployed. Autonomy: average level of autonomy (assistant to multi-agent) across areas with AI deployed. Infrastructure and data readiness: average of self-assessed infrastructure and data readiness. "Don't know" responses count as no AI for breadth and are excluded from the depth and autonomy scores.

Source: Kapronasia primary survey data



> *At Carta, AI adoption is driven from the top, with our CEO setting the tone and AI leaders embedded across departments to drive adoption and share practical use cases. Our CEO actively participated in AI training and encouraged employees to do the same, setting the tone and having the rest of us follow.*

Vanessa Fushi Chin[4]
Head of Business Development APAC & Middle East

carta

Workforce adoption of AI: How much of workforce is using AI in day-to-day activities?

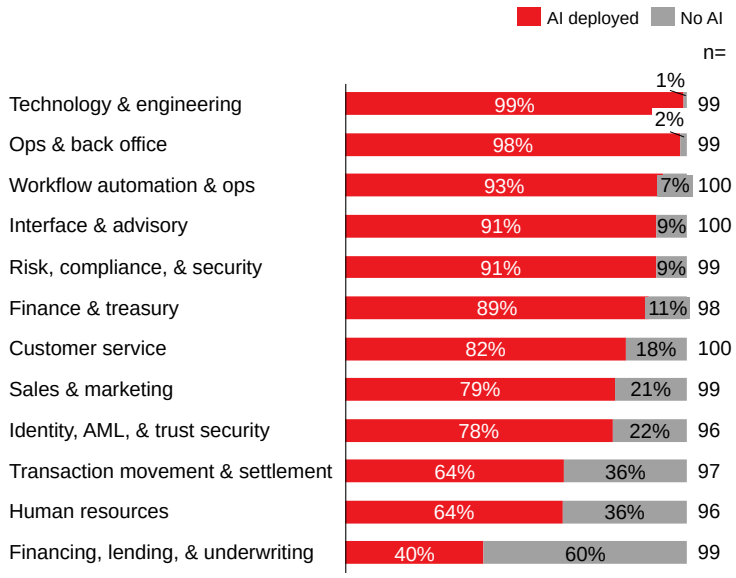
AI adoption in the workforce has yet to become an enterprise-wide norm across most fintechs. While AI is widely available across most functions, employee adoption has not kept pace. Nearly 6 in 10 respondents report that AI is used by half or less of their workforce, indicating that broad access to AI has yet to translate into widespread use in daily work (Figure 4).

Strong leadership, combined with bottom-up adoption is crucial to driving widespread workforce adoption of AI. StashAway, for example, mandates AI usage across the company, while Carta and Airwallex encourage non-technical employees to build their own AI-powered applications. Technology deployment alone, however, does not create workforce adoption. When leadership sets a clear baseline and teams are empowered to experiment, AI is more likely to become an enterprise-wide operating standard.

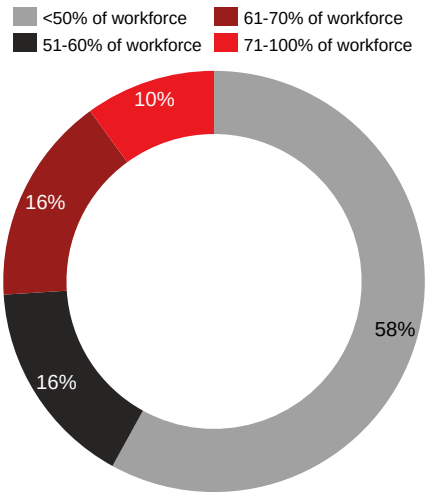
Figure 4

Widespread availability of AI has not translated into widespread workforce adoption

% of AI deployed across functions



How widely is AI used in employees' daily work?, n=100



Source: Kapronasia primary survey data

[4] Note: Vanessa Fushi Chin has since departed Carta after her contributions to this whitepaper

AI adoption breadth: How many functions have adopted AI?

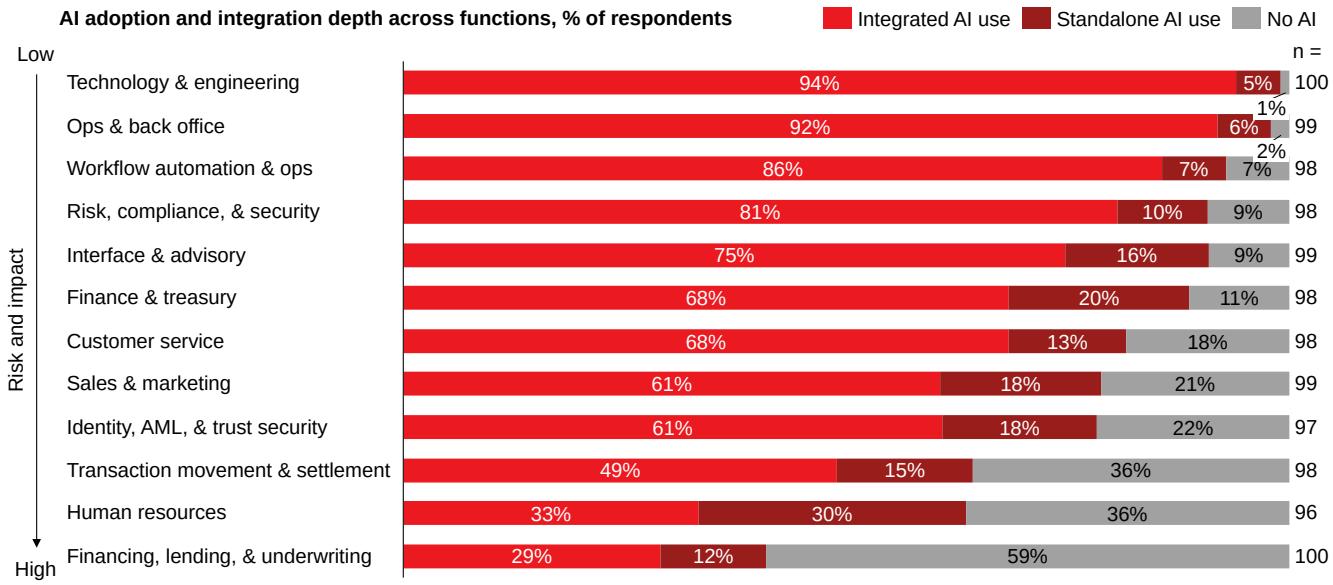
AI adoption is highly concentrated in Technology & Engineering (99%) and Operations & Back office (98%) but remains limited in higher-risk functions. By contrast, 59% reported no AI adoption in Financing, Lending & Underwriting (Figure 5). This suggests that firms are more cautious about extending AI into functions where outcomes carry greater risk and impact, particularly where AI can influence financial decisions or customer outcomes. This caution extends beyond whether AI is deployed to how deeply it is integrated into these functions.

AI integration depth: What % of AI tools are standalone tools vs embedded in core functions?

The same gap becomes more pronounced when examining integration depth. Technology & Engineering leads at 94%, while financing, lending & underwriting falls to just 29% (Figure 5). More importantly, the gap between integrated and standalone use widens as risk and impact of each function increase. This points to a clear boundary around deeper AI integration, as firms remain more cautious about allowing AI to take on greater responsibility in higher-risk functions.

Figure 5

AI adoption and integration depth decrease as functions' risk and impact increase



Note: Data has been re-based and normalized to exclude "Don't know," "Unsure," and "Not applicable" answers. Percentages represent the proportion of respondents with definitive selections across each functional area. Percentages may not sum to 100% due to rounding.
Source: Kapronasia primary survey data



AI autonomy: To what extent are these AI deployed chatbots vs copilots vs agents?

Despite widespread adoption, most firms remain at the AI assistant or copilot stage. Even the functions with the highest AI adoption, Operations & Back Office and Technology & Engineering, have relatively low agentic adoption at 28% and 26%, respectively (Figure 6). This reflects deliberate caution rather than a technical barrier. Agentic tools are concentrated in structured, high-volume functions, with adoption becoming more selective where greater autonomy, system access and the consequences of errors increase risk, or where human relationships remain important.

Infrastructure and data readiness: How ready are the key foundational levers for scaling AI?

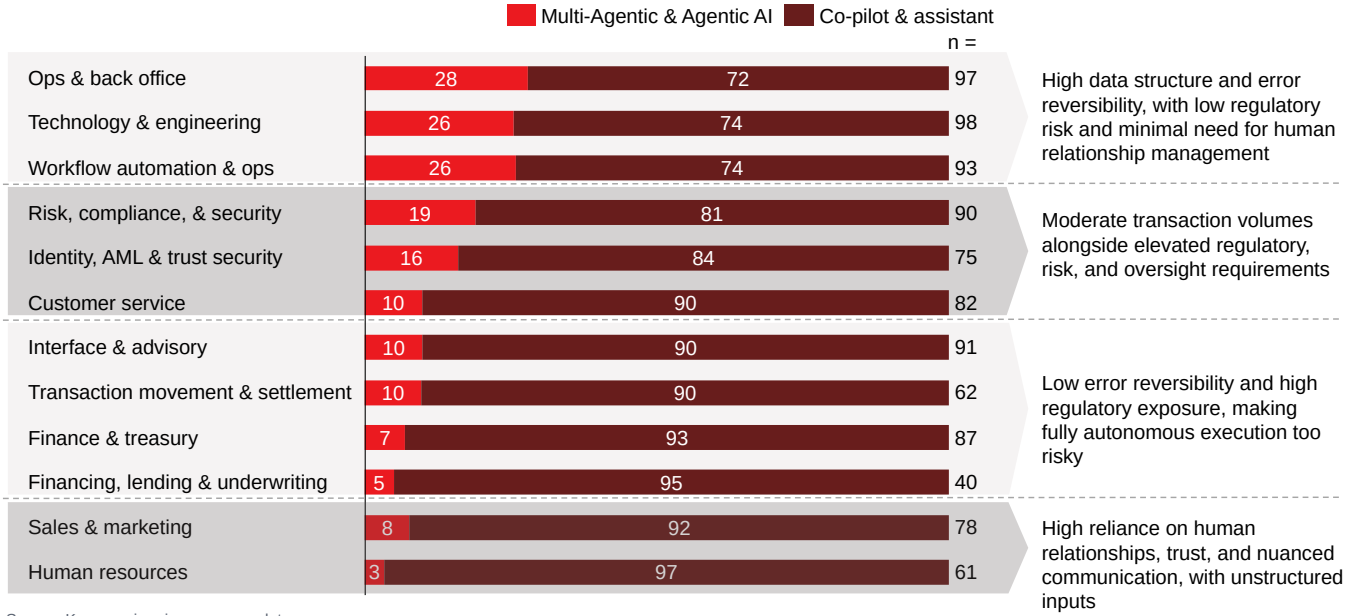
Technical readiness ranks the second highest across the five dimensions, indicating that it is no longer a primary barrier to scaling AI among fintechs. Fintechs report strong technical readiness, with 87% citing infrastructure and 85% data readiness.[5] By comparison, only 23% of organizations globally reported having infrastructure ready to scale AI.[6] Fintechs therefore stand out as a sector where the core technical foundations for AI are already largely in place. As a result, any remaining ceiling on AI deployment is increasingly an operational choice rather than a system constraint.

Fintechs, unlike many other industries, have competitive advantages owing to strong technical readiness. Yet in a sector that demands high trust and risk management, these demands can limit how fast and far firms are able to tap into this edge. Heightened regulatory scrutiny, the sensitivity of financial data, and the consequences of AI-driven decisions require fintechs to scale with greater operational restraint, maintaining human oversight even where the underlying technology is capable of more. The more important question to ask is how can firms balance efficiency gains from AI with the necessary controls.

Figure 6

AI remains largely assistive, with greater autonomy concentrated in operational and technology functions

Degree of AI autonomy across functions, % of respondents



Source: Kapronasia primary survey data

[5] Kapronasia primary survey

[6] Turner, A. (2026) 'Enterprises' data infra is biggest barrier to leveraging agentic AI', Mobile Europe, 19 May. Available at: https://www.mobileeurope.co.uk/enterprises-data-infra-is-biggest-barrier-to-leveraging-agentic-ai/



I see the value of AI in terms of outcomes rather than traditional ROI metrics. AI is the foundation for our core business. From a product perspective, we measure this in terms of velocity—what's the speed at which we are able to build and ship things. With AI, we have been able to accelerate our product suite into one of the most extensive in the market.

Varun Mittal
Managing Director Country Head
Keenai

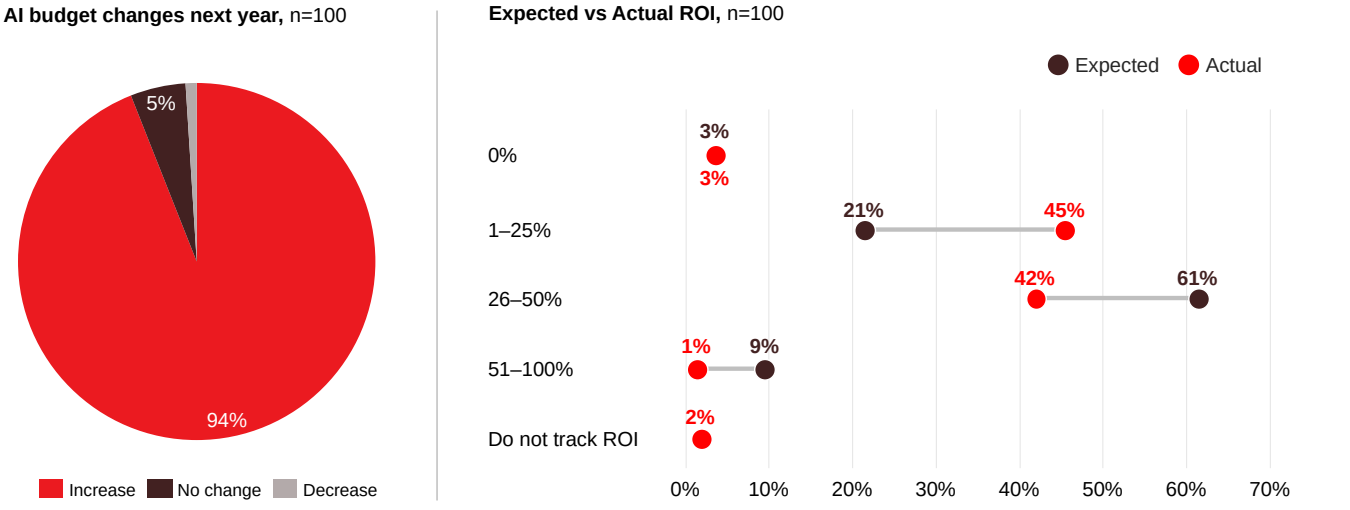
Operational impact & high-value use cases

With AI already widely adopted, fintechs are increasing investments to move beyond initial deployments and capture greater value from AI. Nearly all firms (94%) expect their AI budgets to increase next year (Figure 7). These continued investments suggest that firms see further opportunities to deepen AI's role across the business and expand its impact beyond existing use cases.

However, measuring the returns on AI investments remains challenging, as AI impacts can be difficult to measure using traditional ROI models. This is reflected in the gap between expected and actual returns. Most respondents (61%) expected returns of 26–50%, but only 42% reported actual returns in that range. At the same time, 45% reported returns of 1–25%, compared with just 21% who had expected returns in that range. This gap suggests that expected returns can differ considerably from actual returns, while also pointing to a broader measurement challenge: traditional ROI can capture cost savings and revenue gains but may miss benefits that are harder to quantify financially.

Figure 7

Fintechs continue to invest in AI, but traditional ROI measures struggle to capture its full value



Source: Kapronasia primary survey data

When asked where AI is creating value today, fintechs see the greatest benefits in cost reduction and efficiency, rather than in revenue growth. Operating cost reduction was cited as the primary value driver by 41% of firms, while net-new revenue generation ranks last at just 4% (Figure 8). This may reflect the fact that cost savings and efficiency gains are easier to link directly to AI deployments, while revenue gains are influenced by multiple factors beyond AI. As a result, the value AI creates can be harder to quantify through traditional ROI measures, particularly when its contribution to revenue is difficult to isolate.

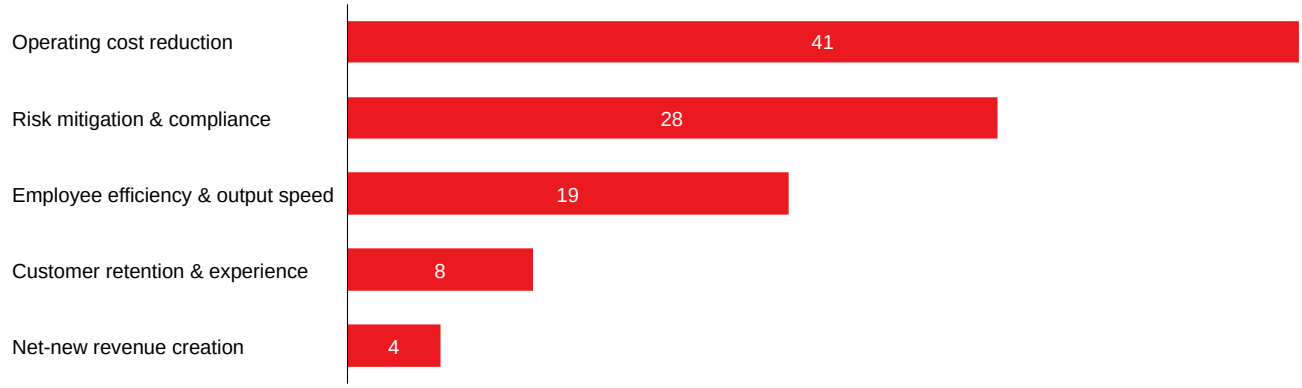
Leading fintechs are therefore shifting towards function-specific measures that link AI deployment to the outcomes that matter for each part of the business. Airwallex, for example, measures sales through deals closed and operations through SLA throughput. This approach makes AI's contribution more visible and helps firms identify where further deployment could create value.

The following use cases illustrate this approach in practice.

Figure 8

AI value remains concentrated in efficiency and risk, with revenue generation still largely untapped

Where AI is most valuable to fintechs today, n=100



Source: Kapronasia primary survey data



How Endowus' WhatsApp Bot freed up 30% of advisors' time to do what AI can't[7]

At Endowus, advisors and support staff have been spending a significant share of their time on routine queries, time that could otherwise go toward the deeper conversations and relationship-building to support trusted advice.

To address this, they deployed their first generative AI-based WhatsApp support bot in 2023, early relative to the rest of the industry, giving them a substantial body of data on what AI can and can't do in this space. Client support is now Endowus' largest area of AI adoption.

Today, about 30% of WhatsApp and in-app queries are handled by AI, freeing advisors and support staff to focus on the human parts of the work AI cannot replace. Client satisfaction with the support process has improved, though this has been driven more by faster resolution times than by AI outperforming human responses.

How Cleanverse scaled outreach & lead generation using autonomous agents[8]

As a lean startup scaling globally, Cleanverse avoided expanding sales overhead by adopting an "AI-first, human-assisted" strategy. Autonomous AI agents handle front-end prospect discovery and engagement, allowing the team to focus on high-value sales qualification and relationship building.

Cleanverse extends this approach to transactions through its AI Payments Platform—a safety layer enabling agents to discover, negotiate, and place orders on a user's behalf. The platform incorporates verified identities, trust-based payment controls, escrow settlement, signature requirements, and automated dispute rules.

Together, these capabilities demonstrate how Cleanverse is redesigning operating models around intelligent delegation, explicit human control, and risk-appropriate protection.

Airwallex's Agentic Business Accounts, a hub for AI-driven financial operations[9]

As AI agents become more capable, businesses are moving beyond using AI for insights towards allowing agents to take action across financial workflows. Airwallex's Agentic Business Accounts are designed for this shift, bringing accounts, balances and activity across multiple entities into one view, while allowing businesses to manage their finances with their preferred AI agents.

Within permissions, rules and approvals set by the business, AI agents can automate tasks such as anticipating funding needs, rebalancing liquidity, putting idle funds to work and timing currency conversions to manage FX exposure.

By combining AI agents with the infrastructure and governance needed to move money, Airwallex is building towards a model where AI can execute financial operations on behalf of finance teams, while businesses retain control over how and when those actions take place.

[7] Interview with Deepak Sarda, Chief Technology Officer, Endowus

[8] Interview with Ceridwen Choo, Co-founder & CEO, Cleanverse

[9] Interview with Timothy Wong, Global VP of Data and AI, Airwallex

> *The shift to agentic AI is bigger than automating individual tasks. It changes how you think about running an organization. The question becomes: what should a human do, what should an AI agent do, and where should they work together? At Airwallex, we're increasingly designing our operating model around that question rather than simply adding AI to the workflows we already have.*



Timothy Wong
Global Vice President, Data and AI

Airwallex

Agentic AI and the future of finance

The shift to agentic AI is already underway. Implementations like Airwallex's Airi[10] and Agentic Business Accounts demonstrate how fintechs are preparing for a future where AI extends beyond supporting individual actions - across both consumer commerce and business financial operations. Agents can complete processes end-to-end, making decisions, using tools, handling multiple steps, and adapting based on feedback. In this model, people move from directing every action to supervising end outcomes.

This shift represents a bigger change than simply adopting more advanced AI tools. Rather than assist staff with existing workflows, agentic AI promises to transform how firms organise and run their businesses. Most firms today work through established systems and workflows, with AI layered on top to assist employees. As agents become more capable, that model could reverse: agents could execute core processes while people supervise their work, manage exceptions, and step in when human judgement is needed. The shift is therefore not just about what AI can do, but about how work gets done.

Despite 42% of Singapore fintechs claiming readiness to implement agentic AI, a mere 11% have actively deployed it in live environments.[11] This significant gap highlights the hesitation companies experience when transitioning from sandbox testing to live operations, especially when autonomous execution could impact sensitive or high-risk functions. At Endowus and StashAway, for example, both MAS-regulated entities, core financial advice cannot simply be delegated to autonomous agents. Rather than introduce autonomy into production, these firms use secure sandbox environments to test agentic tools for robustness and safety before considering live deployment.

Even where agentic AI is deployed, firms continue to keep humans firmly in control. Oversight of deployed agents across fintechs still relies largely on manual reviews and controls. Nearly all senior leaders interviewed expect human oversight to remain a permanent feature of financial services, particularly where agents are involved in critical financial decisions. This is aligned with our survey findings, where 74% of fintechs cite human checkpoints as a requirement for agentic AI (Figure 9). This suggests that agentic AI is currently being introduced within human-controlled boundaries, rather than as fully autonomous systems.

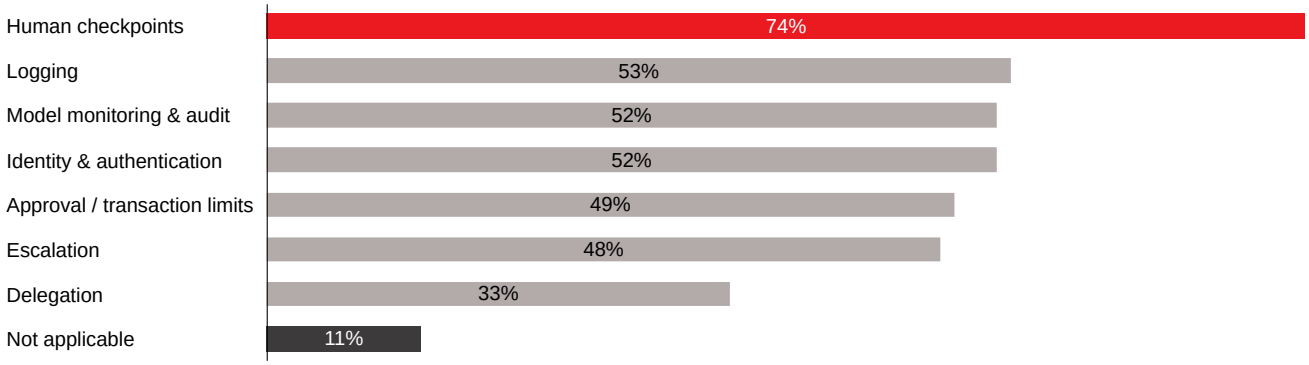
[10] The Airwallex Editorial Team (2026). Introducing Airi: Increase your conversion with a global customer network. Available at: https://www.airwallex.com/global/blog/introducing-airi

[11] Kapronasia primary survey, n=100

Figure 9

Agentic AI is advancing, but human oversight remains the guardrail to autonomous execution

Controls required for agentic AI deployment, n=100



Source: Kapronasia primary survey data

However, this view of human oversight is not universal. Conversations with younger startups, firms with 50 employees or fewer, reveal a different perspective. These firms recognize that accountability remains essential, but do not necessarily see human involvement as the only way to achieve it.[12] With an AI-native mindset, they see full autonomy as an eventual outcome rather than a fundamental limitation, with trust and guardrails determining when it becomes possible. This points to a broader shift in oversight as agentic AI matures, from requiring human involvement to relying on the right controls.



"I would not rule out any specific business function from agentic AI. Almost every decision I can think of—from finance and payments to legal review and audit—could eventually be handled by autonomous agents. I do not see technical capability as the constraint. The key barriers are trust and guardrails. If we are able to solve the trust problem and put the right guardrails around it, many of these decisions can be done agentically and autonomously.

Pramodh Rai
Co-founder & CEO
Cyber Sierra

[12] Kapronasia expert interviews

As this shift takes hold, the implications extend beyond how agents are governed to how work itself is designed and executed. At StashAway, this means moving towards an AI-native way of working, with agentic systems built into core business processes. As AI becomes embedded in everyday workflows, it moves from being a tool employees use to becoming part of how the organization operates. This marks a broader shift from using AI to augment existing work to redesigning work around AI.

This shift is also expected to happen relatively quickly, even as firms remain uncertain about how to get there. Six in ten surveyed respondents believe Agentic AI will become mainstream within the next two years.[13]

However, few leaders see a clear roadmap for getting there.[14] Instead, many are taking a learn-as-they-go approach, experimenting with agents today while working towards a longer-term vision of increasingly autonomous, agent-to-agent processes.

The direction of travel is becoming clearer even if the path is not. While agentic AI has the potential to revolutionize fintech operations, the rate of transformation hinges on factors beyond technology alone. Ultimately, regulatory frameworks, robust governance guardrails, and public trust in AI-driven financial decisions will govern how rapidly firms transition from supportive AI tools to fully autonomous agents operating on their behalf.

> *"To scale AI, you can't just bolt AI onto yesterday's processes. Making an existing task faster is automation, not workflow redesign. We are stepping back and rebuilding core business functions around agentic systems—designing workflows where humans and AI agents work together from the outset, an operating model that is purpose-built for multi-agent autonomy rather than constrained by legacy process bottlenecks.*



Uthra Parameswaran
Chief Compliance and Risk Officer

StashAway

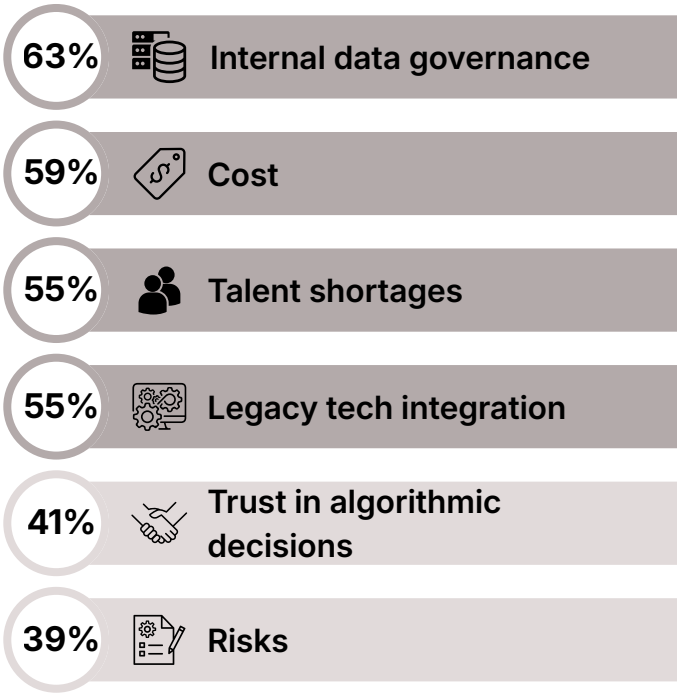
[13] Kapronasia primary survey, n=100

[14] Expert interview with senior leaders

The trust equation: Governance, Risk, and Internal Friction

The technology may be ready for agentic AI, but most firms are still working out how to deploy it with confidence at scale. Internal data governance, cost, talent, and legacy infrastructure are the four biggest barriers cited by fintechs. The challenge is therefore no longer just technical. Firms need to strengthen how they govern agents, measure their value, build the necessary skills, and integrate them into existing systems and ways of working.

Top barriers to scaling AI[15]



Internal data governance

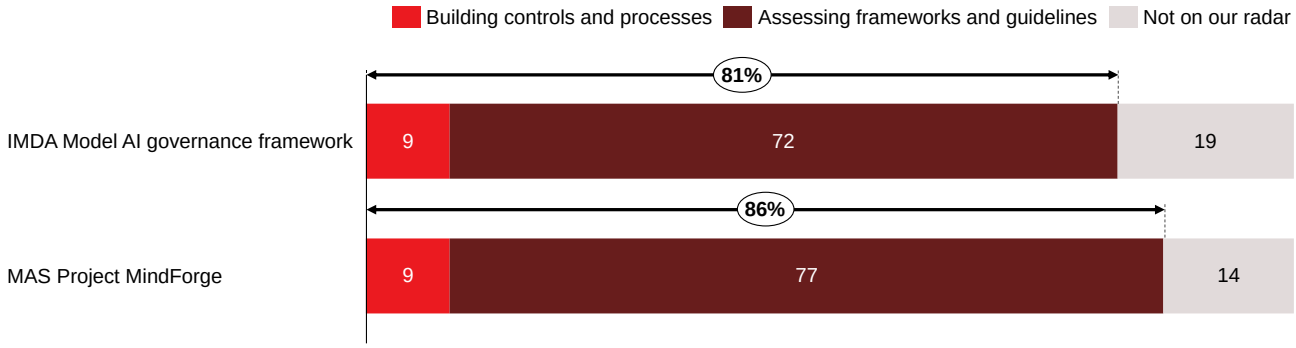
Internal data governance is the biggest constraint to scaling AI, with 63% of fintech experts citing internal data governance as their top barrier. Agentic AI raises the stakes because agents can access data and take actions across multiple processes. Firms therefore need clear boundaries around what agents can access, what they can do and when human intervention is required.

In Singapore, regulators are helping firms navigate this challenge rather than holding them back. This is reflected in strong engagement with their guidance: 81% of firms are engaging with IMDA's Model AI Governance Framework, while 86% are engaging with MAS's Project MindForge (Figure 10). Nevertheless, high engagement has yet to yield widespread adoption; merely 9% of firms actively implement controls mapped to the IMDA framework, and an identical 9% have established SOPs and guardrails under MindForge. This stems

Figure 10

SG Fintechs demonstrate high regulatory awareness across key frameworks, but operational readiness remains low

Firm readiness and implementation of AI governance frameworks, n=100



Source: Kapronasia primary survey data

[15] Kapronasia primary survey, n=100

in part from the swift velocity of technological and regulatory changes, and organizations face significant friction when translating high-level guidelines into actionable, operational controls.

Cost barrier

Cost is emerging as a concern as firms scale AI beyond experimentation. In our interviews, firms described rising token usage as a growing concern as they adopt more powerful models. The challenge is managing costs as AI usage grows, rather than simply the cost of individual models. Coinbase, for example, developed an automated routing framework that matches tasks to different AI models based on their complexity.

The hybrid-skills gap

The talent challenge is not simply a shortage of AI expertise. Firms need people who possess both the business knowledge and the AI competency to leverage AI. Firms are responding by hiring for an AI mindset and developing internal change champions and redefining existing roles. But as agents move into core processes, these skills will need to sit across the organization rather than within technical teams.

Legacy tech integration

Legacy infrastructure is another barrier to scaling agentic AI. Agents need access to connected data and systems to execute complex workflows, but data silos and older technology can limit that access. As a result, firms often keep deployments to narrow, low-risk tasks, even when the underlying AI is capable of doing more.

While cost and legacy technology present tangible hurdles to scaling AI, the more fundamental challenge is creating the conditions for greater autonomy. Internal data governance and workforce capabilities are critical to this shift: leaders need confidence that agents can access only the data and systems they are authorised to use, operate within clear boundaries, and be effectively overseen by employees with the right skills. Without these foundations, even capable AI systems are likely to remain confined to narrower, more controlled use cases. Scaling agentic AI therefore requires more than advancing the technology itself; firms also need the governance, skills, infrastructure and economics to support it in practice. The remainder of this paper outlines a practical framework to help firms build these foundations and move from AI experimentation towards scaled, operational use.

> *As internal AI adoption scaled across Coinbase, frontier models were the natural default. To optimize performance and resource allocation, we experimented with smart defaults using open-weight models alongside an automated routing framework that evaluates task complexity and model pricing. The system automatically dispatches routine or execution-focused tasks to cost-effective models, while reserving frontier tools for high-complexity planning and reasoning. This approach has enabled us to decouple our AI expenditure from company-wide token growth.*



Pete Patanapanlert
Product Lead

coinbase

Our framework

Scaling AI involves looking beyond piloting and testing AI in silos. In particular, there are six core areas that firms will have to consider across the entire implementation lifecycle: vision and strategy; implementation and outcomes measurement; governance and risk; people and talent; data and infrastructure; and operating models and processes (Figure 11).

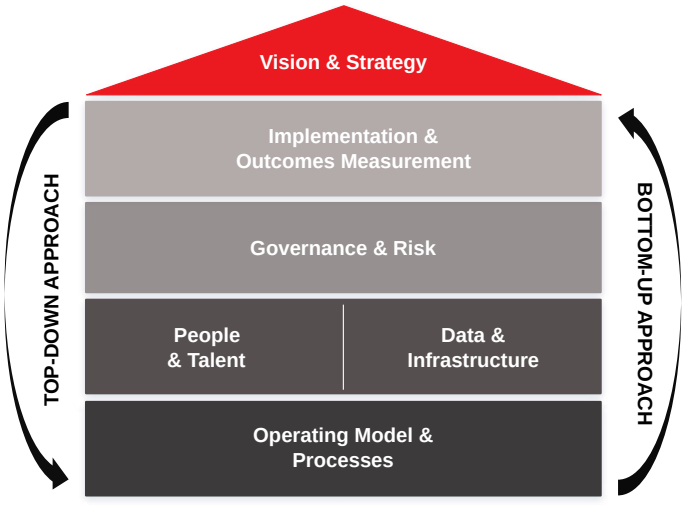
There are 3 approaches firms can take to tackle these 6 areas, depending on their organizational priorities and bias for speed:

- **Top-down approach:** Firms start with enterprise vision and strategy, establishing governance boundaries and organizational alignment before redesigning core operational processes and workflows.
- **Bottom-up approach:** Firms begin by reimagining specific operating models and workflow processes around AI capabilities, iteratively building the supporting data, talent, and governance infrastructure as they scale.
- **Hybrid approach:** Firms establish enterprise priorities and risk controls from the top down, while simultaneously allowing agile teams to prototype, redesign, and prove high-impact processes from the bottom up.

Although each model provides unique advantages, leading organizations increasingly emphasize bottom-up execution or adopt a hybrid framework in established institutions to balance enterprise-level governance with operational speed. This approach enables firms to gain maximum agility and accelerate momentum, operating much like a high-velocity startup.

Figure 11

Fintechs that implement AI with a bottom-up or hybrid approach stand to gain speed and leverage



Approach	Benefit	Suited for
Top-down	• Allows firms to balance enterprise priorities, **before altering** core operations and structure	• Larger enterprises that have **multiple stakeholders** involved • Segments where **stakes are high**
Bottom-up	• Enables firms to **move quickly** and design operating models with **less redundancy**	• **Younger** and **challenger** AI-native firms • **Leaner teams** & **fewer legacy constraints** let them design the operating model around AI
Hybrid	• **Balances** both **enterprise priorities** and **controls**, with selected areas able to **iterate fast** and see returns	• **Established enterprises** that still seek to have speed and competitive edge whilst staying cautious and intentional

Source: Kapronasia analysis

Matching autonomy to risk and repeatability

To balance efficiency gains from AI with the appropriate degree of risk and governance considerations, firms should think beyond simply maximizing autonomy or applying the highest degree of oversight or controls as a blanket approach.

There are two dimensions that guide the degree to which autonomy and controls can be applied to each specific use case:

- **Business risk:** the magnitude of consequences of an unreliable output or action.
- **Task repeatability:** this determines how reliable and consistent AI can become, based on the amount of established patterns and rules it can learn from.

This creates a basis for determining the appropriate level of autonomy and matching controls (Figure 12).

1. **High-risk, less repeatable tasks** mean that AI cannot be trusted for reliability, and failure is not tolerated. This requires the highest degree of human oversight, where humans should review and approve everything. There should also be a high degree of technical controls, to further mitigate risks.

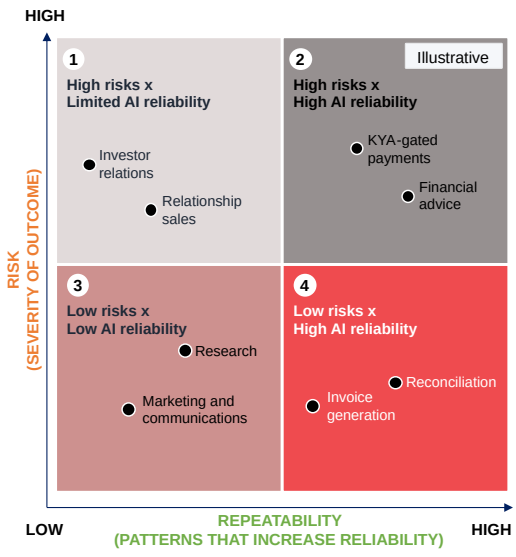
2. **High-risk, highly repeatable tasks** mean there is a high degree of confidence that the output produced by AI is reliable and accurate, but failure remains intolerable. Here, AI can help perform tasks more efficiently, whereby humans do not have to oversee everything (but remain as the final decision maker at material decision points). There should also be a high degree of technical controls, to further mitigate risks.

3. **Low-risk, less repeatable tasks** mean whilst the impact of failure is limited, the capability of AI to perform a task reliably and accurately is equally limited. Such tasks may benefit from AI assistance while retaining a limited degree of human verification to ensure the quality of the final output.

4. **Low-risk, highly repeatable tasks** mean the impact of failure is limited, and there is a high degree of confidence that the output produced by AI is reliable and accurate. This is where AI has the highest potential to create reliable efficiency gains. To tap into this, firms can consider a human-out-of-the-loop approach with lighter controls and periodic monitoring

Figure 12

Balancing efficiency with human oversight and technical controls can be mapped to the degree of risk, or how repeatable a task performed by AI is



CATEGORY	DEGREE OF CONTROLS (NOT EXHAUSTIVE)*		RATIONALE
	HUMAN	TECHNICAL**	
1 High risks x Low AI reliability	• [Human-in-the-loop] Human reviews and approves **everything**	• Bounded credentials • Logging • Model monitoring & Audit • Tight limits • Escalation • No Delegations	• Agent isn't reliable yet, and stakes are high — controls should assume it will get it wrong and mitigate risk
2 High risks x High AI reliability	• [Human-in-the-loop] Human approves at **significant steps** and can intervene any time	• Bounded credentials • Logging • Model monitoring & Audit • Tight limits • Escalation • Narrow Delegation	• Agent is reliable, but a rare miss is costly — controls should catch the tail, not slow the median case
3 Low risks x Low AI reliability	• [Human-in-the-loop] Human approves at critical steps or steps in at failures	• Standard authentication • Logging • Looser limits • Periodic audits only	• Low stakes, but the agent still isn't reliable — controls stay light on consequence, not light on verification
4 Low risks x High AI reliability	• [Human-out-of-the-loop] No human approvals needed, but might audit	• Standard authentication • Logging • Looser limits • Periodic audits only	• Pattern is well-learned and errors are low-stakes — audits can be done for quality checks

HIGH ↓ LOW

*Firms should refer to IMDA & MAS guidelines on types of controls recommended, and supplement with own internal policies
**IMDA recommends deterministic, system-level controls over prompt-layer instructions (IMDA MODEL AI GOVERNANCE FRAMEWORK FOR AG ENTIC AI SECTION 2.3.1).

This framework of matching autonomy and controls to the risk and reliability of each use case, allows firms to balance efficiency with the necessary degree of oversight and controls.

Building an AI-native organization

Central to the success of any organization is people and talent. Building an AI-native organization requires a mix of top-down and bottom-up initiatives to ensure organizational alignment, close gaps in capability and knowledge, and cultivate a culture of experimentation and sharing (Figure 13).

Top-down: Create alignment

Leadership should set clear AI priorities, establish expectations for responsible use and demonstrate willingness to learn and experiment. This creates the direction and organizational support needed for adoption.

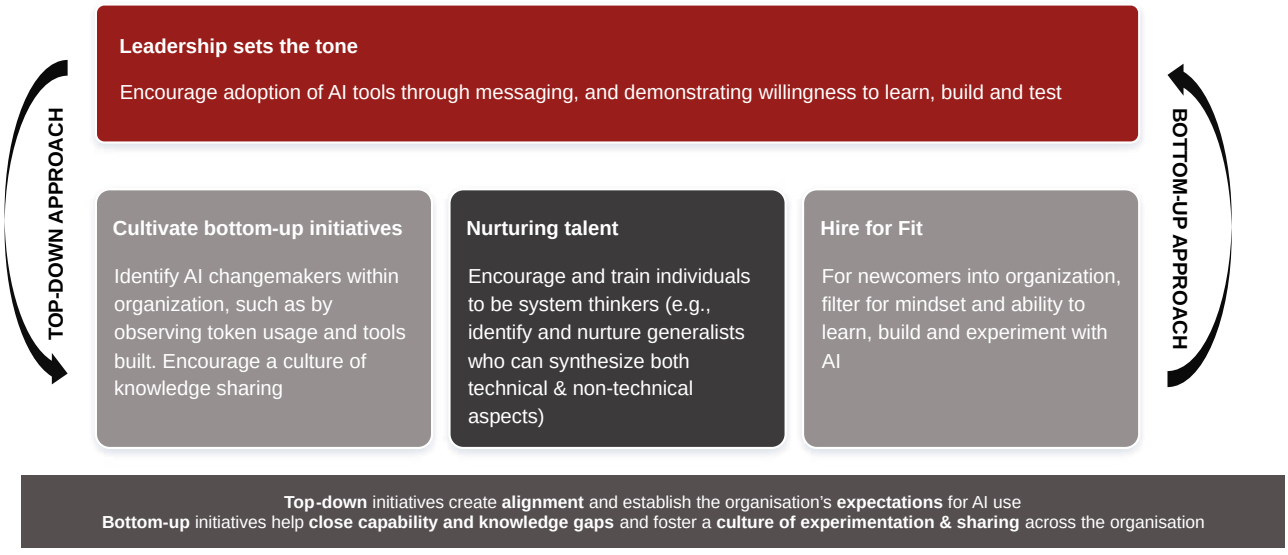
Bottom-up: Build capability

Firms should identify employees already experimenting with AI, create channels to share successful use cases and develop systems thinkers who can connect AI capabilities with business processes, controls and customer outcomes. Recruitment should also assess candidates' ability to learn and work across technical and non-technical domains.

Both are required. Top-down direction establishes where and how AI should be used. Bottom-up capability enables employees to turn that direction into practical adoption and, ultimately, AI-native ways of working.

Figure 13

Cultivating an AI-native organization requires both a mix of top-down and bottom-up initiatives



Conclusion

Singapore's fintech sector has moved beyond the question of whether to adopt AI. The primary hurdle now lies in determining the extent to which AI can be integrated into operational frameworks and how much autonomy can be safely entrusted to these systems. Today, organizational capabilities, specifically deploying AI with strong oversight and governance, represent a far more critical factor than technical readiness alone.

To navigate this transition, organizations must move away from isolated AI applications and systematically realign their workflows, operational models, and decision-making processes. While strategies will vary across institutions, firms should aim to establish the requisite governance structures, technical skills, and operational environments necessary to scale AI responsibly without sacrificing accountability.

The next phase of AI maturity will therefore not be measured by deployment alone, but by how effectively firms turn AI capability into sustained business value. For Singapore's fintechs, the opportunity is to build on their strong adoption and technical foundations and move from AI that assists people towards AI that can responsibly act on their behalf.

Ultimately, the advantage will belong not simply to firms that adopt AI first, but to those that can determine where AI should act, where people should remain in control, and how the two can work together at scale.



Acknowledgments (interviewee profiles)



Timothy Wong
Global VP of Data & AI

Airwallex



Deepak Sarda
Chief Technology Officer

Endowus



Vanessa Fushi Chin
Head of Business Development (APAC, ME)

carta



Kathy Qu
Head of Issuance Platform & Settlement Platform, On-Chain

HASHKEY GROUP



Ceridwen Choo
Co-founder & CEO

Cleanverse



Varun Mittal
Managing Director Country Head

Keenai



Pete Patanapanlert
Product Lead

coinbase



Jeremy Foo
Singapore Country Manager

StashAway



Pramodh Rai
Co-founder, CEO

Cyber Sierra



Uthra Parameswaran
Chief Compliance and Risk Officer

StashAway

Acknowledgments

Airwallex

Arnold Chan
General Manager, Asia-Pacific

Benjamin Goh
Associate Director, Corporate Affairs

Timothy Wong
Global Vice President, Data and AI

Singapore FinTech Association

Bryan Tay
General Manager

Samantha Soong
Program Lead, Subcommittees and Working Groups

Swati Sodhani
Program Lead, Subcommittees and Working Groups

Centre for Finance, Technology, and Entrepreneurship

Huy Nguyen
Co-founder

About us

Airwallex is an AI native financial operating system for modern businesses. We are building the future of global banking for a borderless, real-time, intelligent economy.

We serve more than 676,000 businesses worldwide, either directly or through our platform customers. From startups to public enterprises, our customers use Airwallex to manage their global banking and financial operations, or to build and monetise their own financial products.

Founded in Melbourne in 2015, Airwallex holds 85+ licenses across North America, Europe, the Middle East, and Asia-Pacific, forming one of the most comprehensive financial infrastructures in the world. This regulated backbone powers Airwallex products at global scale, including payment acceptance, billing, global accounts, corporate cards, and spend management.

The company is co-headquartered in San Francisco and Singapore with over 2,600 employees across 27 offices. Learn more at www.airwallex.com.

SINGAPORE FINTECH ASSOCIATION

The Singapore FinTech Association (SFA) is a cross-industry non-profit organisation. Its purpose is to support the development of the FinTech industry in Singapore and facilitate collaboration among the participants and stakeholders. The SFA is a member-based organization with over 700+ members. It represents the full range of stakeholders in the FinTech industry, from early stage innovative companies to large financial players and service providers. To further its purpose, the SFA also partners with institutions and associations from Singapore and globally to cooperate on initiatives relating to the FinTech industry. Well-connected globally, the SFA has signed over 70 international Memorandum of Understanding (MoU) to lay the network for its members and ecosystem.

For more information, please visit www.singaporefintech.org

Kapronasia

Kapronasia, an Atlas Technologies Group Company, is a leading consulting and market research firm specializing in fintech, banking, payments, and capital markets. Our services aim to equip clients across the region with the necessary insights to capitalize on their most valuable opportunities and maintain a competitive edge in the market.

For more information, please visit sg.atlstech.com