SECURITIES AND FUTURES COMMISSION
證券及期貨事務監察委員會

Large Firm AI Readiness Survey

The adoption of artificial intelligence in Hong Kong's financial markets



Contents

Section	*Page*
EXECUTIVE SUMMARY	2
1. SURVEY BACKGROUND	3
1.1 Survey population and market footprint	3
1.2 Methodology and limitations	4
1.3 Use of AI in production of this report	4
2. KEY SURVEY INSIGHTS	5
2.1 Momentum: AI adoption is expected to accelerate	5
2.2 Reach: AI spans the operating model	6
2.3 Integration: connectivity changes the control perimeter	8
2.4 Autonomy: agentic AI is emerging, but maturity varies	10
2.5 Cyber resilience: frontier AI is a defensive opportunity as well as a risk driver	13
3. PRACTICAL BARRIERS AND CHALLENGES FOR AI ADOPTION	15
3.1 Barriers to broader adoption	15
3.2 Challenges in governance and risk management	17
4. IMPLICATIONS FOR THE SFC'S APPROACH TO AI	19
4.1 Investor protection: adapt safeguards as investor interactions change	19
4.2 Cybersecurity and resilience: keep pace with AI-enabled threats	19
4.3 AI governance and applications: keep humans at the centre as AI gains autonomy	20
4.4 Engagement and collaboration: build knowledge and capability together	20
5. CONCLUSION AND DIRECTION OF TRAVEL	21
GLOSSARY	22

Executive summary

Scope and key message. The Securities and Futures Commission (SFC) surveyed 20 large financial firms in July and August 2026, covering 78 affiliated licensed corporations (LCs), which collectively account for over 75% of reported securities transactions. The findings show artificial intelligence (AI) moving rapidly beyond standalone productivity tools towards wider deployment, deeper integration with internal systems and greater autonomy. As this transition gathers pace, the supervisory challenge also changes: if AI is to scale responsibly, new safety and resilience risks need to be understood, governance must extend beyond models and generic human oversight, and capability gaps need to be addressed.

Adoption is broadening and integration is deepening. Nearly all firms expect AI use to increase over the next 12 months. Current deployment spans research, client service, controls and operations: most firms report using AI for research and analytics, while a majority use AI to prepare client communications or documentation. Nearly all firms classify their AI use as partially integrated with internal systems. Connections extend from messaging and document repositories to customer data and trading-related systems, making data quality, confidentiality, access permissions and downstream workflows integral to the safe use of AI in business processes.

Agentic AI[1] is gaining traction, but maturity varies. More than half of the firms surveyed permit agentic AI in production or controlled pilots, and most expect adoption to increase. Applications range from information retrieval and drafting to software development, operational workflows and some trading-related activities, with varying degrees of autonomy and human involvement. The findings point to growing adoption, but not a uniform shift towards fully autonomous AI.

Cyber resilience is critical as frontier AI strengthens both attack and defence capabilities. Nearly all firms currently use or are actively evaluating frontier AI models for cyber defence, with reported benefits including faster analysis and broader detection coverage. At the same time, firms highlight false positives, the need to verify AI-generated findings and the specialist expertise required for effective use. Deployment ranges from early evaluation and pilots to production use, with some firms retaining human approval for remediation or response actions.

Broader adoption depends on overcoming practical implementation barriers and governance challenges. The main barriers are technology integration, moving pilots into resilient production, costs and uncertain business value, data constraints and specialist skills. Firms also face difficulty scaling AI inventories and lifecycle controls, assigning accountability, validating evolving models, maintaining meaningful human oversight, and managing third-party dependencies. Scaling AI therefore requires not only capable models and use cases, but also reliable infrastructure and data, adequate expertise, a viable business case and governance suited to increasingly connected and autonomous applications.

Implications for the SFC. The findings support a supervisory approach that is **safety-first, human-anchored and collaboration-led**. Under this approach, it is essential to address emerging safety, cybersecurity and resilience risks; preserve human accountability as greater authority is delegated to AI; and drive collaborative work among regulators, industry, technology providers and academia to close infrastructure, knowledge and capability gaps.

[1] Agentic AI refers to AI that can autonomously plan and take actions to achieve an objective with limited human direction.

1. Survey background

The SFC conducted a survey of LCs belonging to large financial firms in July–August 2026 to better understand how AI is being adopted across their Hong Kong businesses. Specifically, the survey:

- examined the use of AI in business activities and operational processes, the extent of its integration with internal systems, and the firms' plans for further adoption;
- studied how firms are adapting their governance and risk management arrangements as AI applications become more widely deployed and increasingly capable of performing tasks autonomously; and
- explored the challenges associated with responsible adoption, including data quality and security, model reliability, human oversight, third-party dependencies and the capacity of control functions to keep pace with technological change.

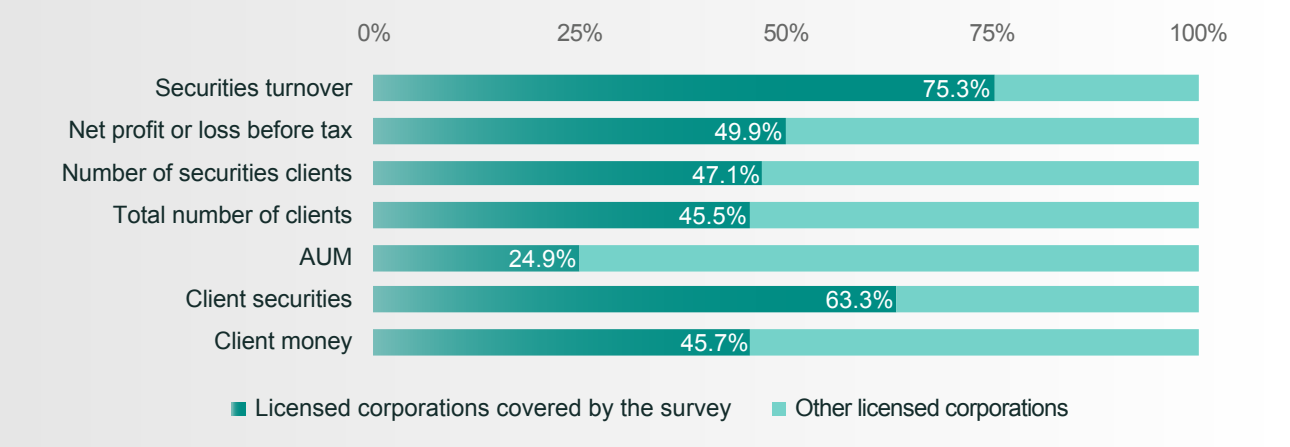
The responses inform the SFC's assessment of the opportunities and risks arising from AI adoption and the implications for its supervisory approach. This report presents the key insights and considers how the SFC can support responsible innovation in financial markets while maintaining appropriate standards of governance, accountability, operational resilience and investor protection. While focusing on large firms rather than the industry as a whole, the survey provides a useful basis for identifying emerging issues and priorities for further engagement and supervisory attention.

1.1 Survey population and market footprint

The survey covered 20 large firms, comprising 12 large global firms, five large Chinese Mainland firms and three large online retail brokers. This population was chosen to provide insights into AI adoption across large firms with different business models and operating structures, including internationally active financial groups and online retail brokerage businesses with a significant share of the Hong Kong market.

Taken together, the in-scope firms account for over 75% of the turnover of securities transactions reported by LCs for the six months ended 30 June 2026. Altogether, the survey covered 78 LCs affiliated to these 20 in-scope large firms and the responses in respect of individual LCs were consolidated at the firm level. Accordingly, a reported application or capability does not necessarily extend to every entity or business line within a firm. While the survey offers useful perspectives on developments among large firms, its findings should not be interpreted as representative of all SFC-licensed corporations.

Figure 1: Selected business indicators of the sampled firms (percentage of the total for all LCs)



Source: Information reported to the SFC by LCs under the financial returns for the six months ended 30 June 2026.

1.2 Methodology and limitations

This report presents anonymised findings from the survey and considers their implications for the SFC's approach to AI. Reported practices are not an endorsement of particular technologies or evidence that controls are effective. The discussion does not replace applicable legal or regulatory requirements.

Unless otherwise stated, percentages are based on the 20 consolidated firm responses, with each firm accounting for five percentage points. Responses may encompass multiple affiliated entities; a reported capability may therefore reflect use within part of a group rather than across every Hong Kong entity or business line. Where respondents could select more than one category, the results may overlap and the percentages may not add up to 100%. The findings reflect respondents' own descriptions and classifications and should not be interpreted as independently verified assessments of their systems or the effectiveness of their controls.

Where applicable, the report distinguishes between applications in production, controlled pilots or tests, and development or consideration. Statements about expected growth reflect respondents' intentions rather than forecasts of deployment. Counts of adoption barriers and governance challenges are drawn from the thematic analysis of responses, with themes normalised across firms and each firm counted no more than once per theme. These categories may overlap and involve analytical judgement; their frequency should not be interpreted as a ranking of severity.

Comparisons should also take account of differences in how respondents may interpret terms such as AI, agentic AI and frontier AI. The illustrative cases encompass both language-model applications and other AI technologies and do not necessarily represent equivalent capabilities or levels of autonomy.

1.3 Use of AI in production of this report

Generative AI tools assisted with the initial thematic analysis of survey responses and preparation of draft contents. SFC staff reviewed and validated the substantive analysis, findings and conclusions, and remain responsible for the content of this report. The cover illustration was generated using AI.



2. Key survey insights

MOMENTUM	REACH	INTEGRATION	AUTONOMY	CYBER RESILIENCE
95% of firms surveyed expect AI use to grow	**100%** of firms surveyed already use AI in back-office or support functions	**90%** of firms surveyed already connect AI to internal systems	**55%** of firms surveyed already permit agentic AI in pilot or production	**90%** of firms surveyed already use or are evaluating frontier AI for cyber defence

2.1 Momentum: AI adoption is expected to accelerate

The survey asked respondents to indicate whether they expected their firm's use of AI to increase over the next 12 months and, if so, whether the increase would be significant or more modest. They were also asked to identify the five business functions or areas in which they anticipated the greatest growth in AI use.

Figure 2: Expected growth in AI adoption



Out of the 20 surveyed firms, 19 expect their use of AI to increase over the next 12 months: nine anticipate a significant increase and 10 a more modest increase. The remaining firm considers it too early to say. Taken together, the responses indicate plans to expand AI across revenue-generating activities, supporting operations and the control environment, rather than within a single part of the business:

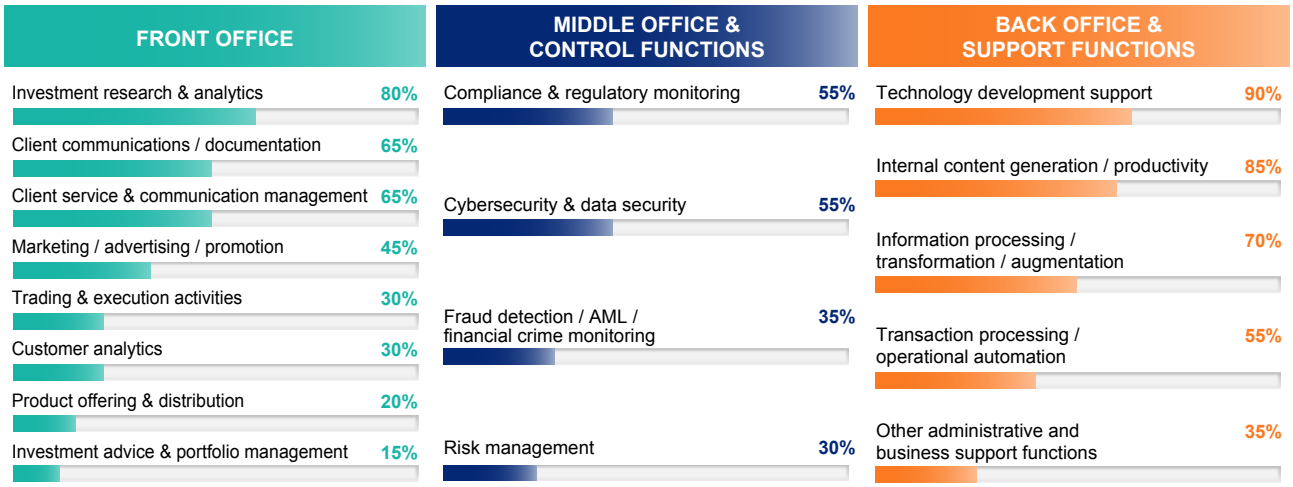
- Operational efficiency, workflow automation and internal productivity are the most widely identified areas of anticipated growth, selected by 17 firms, suggesting that improving internal processes remains a common priority;
- 14 firms anticipate growth in investment, markets, research and analytics; 13 identify client service, sales and marketing-related activities, pointing to further expansion into business and client-related functions;
- 12 firms expect growth in risk management, legal, compliance and cybersecurity, and
- 11 firms expect growth in technology development and code modernisation.

Survey finding. The responses point to strong momentum for further AI adoption across both front-and back-office activities. Internal productivity and workflow automation remain the most widely identified areas of expected growth, but firms' ambitions also encompass investment and research, client-related services, control functions and technology development. This breadth suggests that AI is becoming more integral to firms' operating models, rather than remaining standalone productivity tools. However, it remains to be seen how quickly planned applications will enter production or whether they will deliver their intended benefits.

2.2 Reach: AI spans the operating model

The survey asked respondents to identify the business functions or processes in which AI systems were currently deployed within their Hong Kong operations or in support of Hong Kong-regulated business. The question covered front-office activities, middle-office and control functions, and back-office and support functions, allowing respondents to select multiple categories.

Figure 3: Functions or processes where AI is currently deployed in Hong Kong operations or in support of Hong Kong-regulated business (percentage of firms surveyed deploying AI in each area)



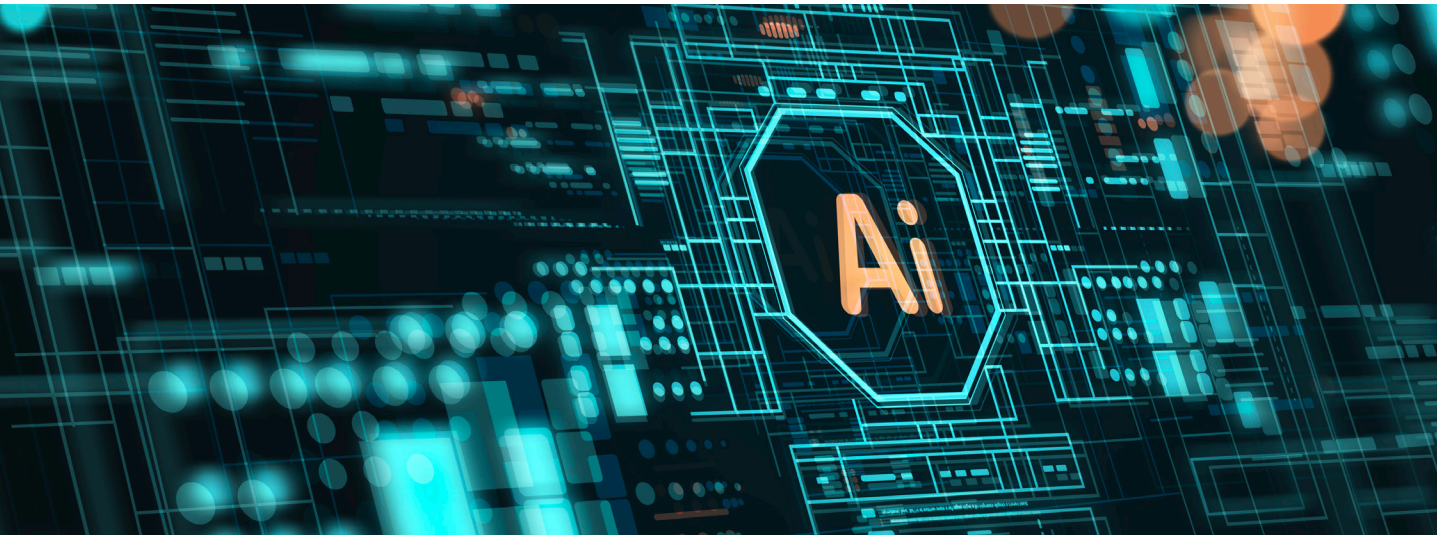
In the front office, Figure 3 shows that AI is most widely deployed in investment research and analytics (80%), followed by the preparation of client communications or documentation and client service and communication management (65% each). Adoption is less widespread in marketing, advertising and promotion (45%), trading and execution and customer analytics (30% each), product offering and distribution (20%), and investment advice and portfolio management (15%). This pattern indicates a broader reported footprint in research, information preparation and service support than in investment advice, portfolio management or execution.

2. Key insights from the survey

Within middle-office and control functions, 55% of firms reported AI use in compliance and regulatory monitoring, and the same proportion in cybersecurity and data security. Reported adoption is lower in fraud detection, anti-money laundering and financial crime monitoring (35%) and risk management (30%). These findings show that AI is being incorporated into the control environment as well as the activities it oversees, offering potential to support information analysis, monitoring and the identification of risks. The variation across functions suggests an uneven deployment footprint rather than a uniform transformation of controls. The survey did not seek to measure the effectiveness of AI-enabled controls in mitigating the relevant risks – a topic for future studies.

In back-office and support functions, AI is most widely deployed in technology development support (90%), followed by internal content generation and productivity (85%) and information processing, transformation and augmentation (70%). Reported adoption is lower in transaction processing and operational automation (55%). These responses indicate a stronger reported footprint in technology development, content preparation and information handling than in transaction processing and operational automation. The pattern is consistent with AI supporting staff productivity and established internal workflows, although the reported deployment figures do not establish the extent of automation or the effectiveness of individual applications.

Survey finding. The responses show that AI already spans front-, middle- and back-office activities, with particularly widespread use in technology support, content generation and research. Deployment in client-related services, control functions and operations further illustrates its reach across firms' operating models. However, adoption remains uneven: reported use is more common in information analysis, preparation and support than in investment advice, portfolio management or execution. This pattern suggests that AI's current foothold is strongest in augmenting staff capabilities and supporting established workflows, while adoption remains more selective in activities directly involving investment decisions and transactions. AI can therefore influence client information, services and outcomes both directly and through staff-assisted workflows. Client-related deployment does not, however, necessarily imply autonomous interaction or decision-making.



2.3 Integration: connectivity changes the control perimeter

The survey asked respondents to indicate whether AI systems currently in use are connected to internal data sources or applications and to describe their level of integration. A related question asked them to identify the sources or systems connected to AI, including email and messaging platforms, document repositories, customer-data systems, trading and portfolio systems, and compliance-monitoring systems. Respondents could select multiple sources or systems.

Figure 4: AI integration with internal systems and connected data sources (percentage of surveyed firms connecting each source/system to AI)

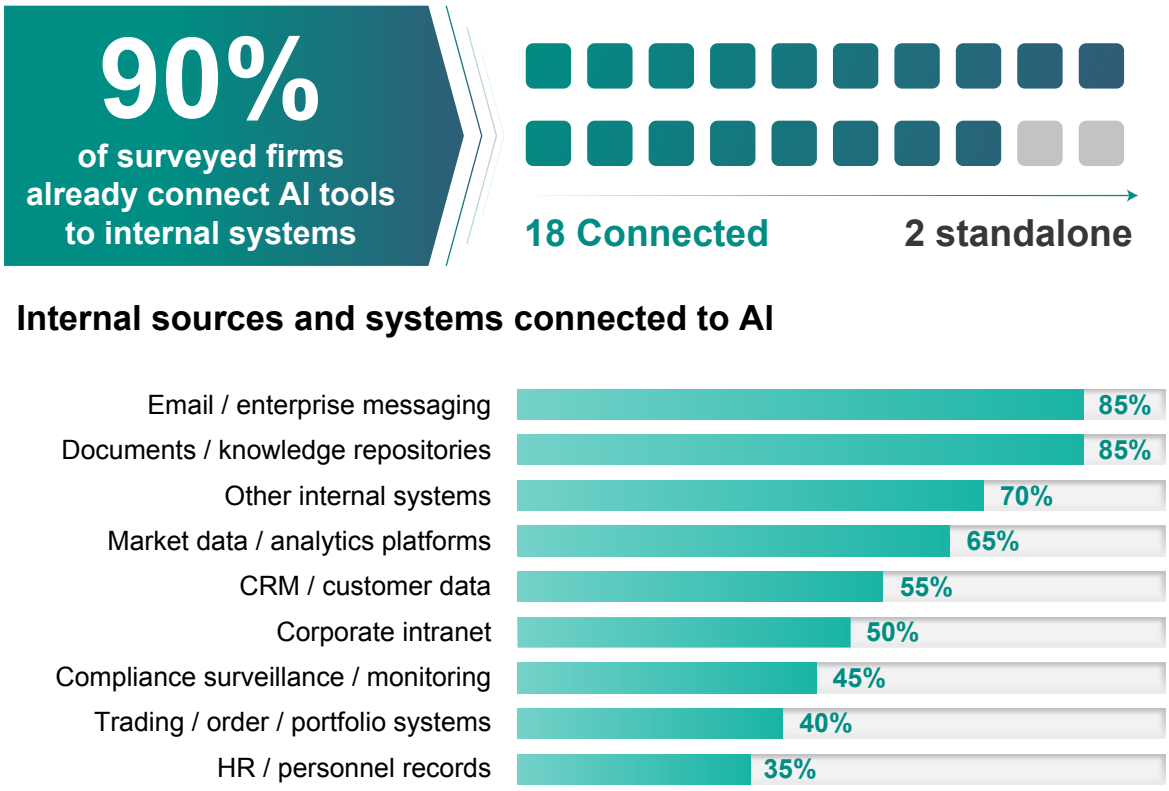


Figure 4 indicates that connected AI applications are already widespread among the surveyed firms: 18 firms (90%) classified their systems as partially integrated with internal systems, while two classified them as standalone. The separate breakdown of connected sources shows the greatest prevalence in email or enterprise messaging and document or knowledge repositories, each selected by 17 firms (85%). This pattern is consistent with AI supporting access to firms' internal information and accumulated knowledge, rather than relying exclusively on standalone prompts or external material. Market data or analytics platforms were selected by 13 firms (65%), while other internal systems were selected by 14 (70%); the breadth of the latter category limits conclusions about the specific connections involved. Connectivity also extends to customer relationship management or customer-data systems (11 firms; 55%), intranets (10; 50%), compliance monitoring (nine; 45%), trading, order or portfolio systems (eight; 40%) and personnel records (seven; 35%).

2. Key insights from the survey

The central insight is that integration changes the scope of AI controls: reliable model behaviour is a necessary focus, but insufficient when an application can draw on internal communications, client information or operational systems. The high prevalence of connections to messaging and knowledge repositories makes source quality, confidentiality and appropriate information access particularly relevant. Connections to customer and personnel records bring additional sensitivity, while links to compliance and trading-related systems warrant an understanding of how outputs enter monitoring, decision-making or transaction workflows. These are implications of the reported connectivity, not findings that unauthorised access or control failures have occurred.

Assessment of the specific risks of a given AI application should distinguish its ability to retrieve information from its authority to modify records, initiate an action or approve a decision. As agentic capabilities expand, that distinction becomes increasingly important to defining authorised activities and the points at which human review or approval operates. Oversight also needs to consider the full chain of dependencies between the model, data sources, interfaces and downstream processes, including whether outdated information, access changes or disruption in a connected service could affect the application's operation.

For firms and the SFC, the findings support an end-to-end view of workflows, encompassing objectives, operating boundaries, data flows, permissions and dependencies. These workflows should have clear ownership, with effective monitoring and intervention proportionate to the sensitivity of the information and the consequences of permitted actions.

Survey finding. The responses show that AI is becoming embedded in firms' internal information environments rather than operating solely as a standalone tool. Connections are most widespread in messaging and knowledge repositories, consistent with a strong focus on retrieving and using institutional information. Links to customer data, compliance platforms and trading-related systems show that integration also reaches business and operational workflows. Taken together, this pattern points to a growing role for AI as an interface between staff, information and existing systems, although connectivity is distinct from authority to change records or execute transactions.

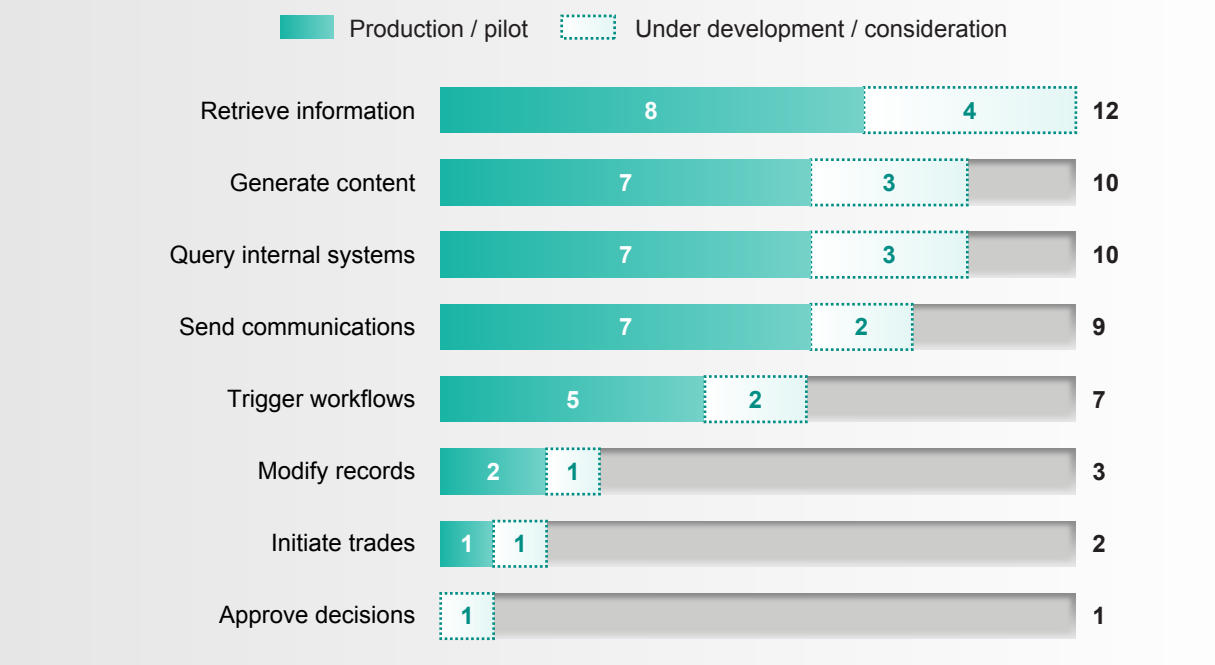


2.4 Autonomy: agentic AI is emerging, but maturity varies

The survey asked respondents to indicate whether their firm permitted autonomous or agentic AI systems and whether these were in live production, controlled pilots or tests, or development or consideration. They were also asked to identify the actions such systems could perform autonomously, ranging from retrieving information and generating content to modifying records, initiating transactions and approving decisions without human intervention. A further question asked whether they expected to adopt more agentic AI over the next 12 months and the anticipated extent of any increase.

Six firms permit agentic AI in live production, while five permit it in controlled pilots or tests. Together, these 11 firms account for 55% of total. The remaining nine firms are developing or considering agentic AI use cases.

Figure 5: Autonomous actions by AI agents reported by firms (number of firms)



Counts in the above chart represent firms selecting each action, not numbers of agents used or transactions impacted. Multiple selections were permitted.

The survey asked respondents to provide details of each active AI use case, including its purpose, primary business function, deployment status and model source; whether it interacted with clients or related to regulated activities; its role in decision-making; and the human oversight, intervention arrangements and risk assessment applied.

These responses provided the basis for the following eight illustrative AI examples which were all flagged by survey respondents either as agentic AI or as including partially or fully automated decision making. They encompass different technologies and levels of autonomy, ranging from decision-support tools to more autonomous workflows.

2. Key insights from the survey

1



Application

Research search and question answering

Reported functionality and oversight

An AI application searches published analyst research and generates responses to natural-language questions. The application operates without direct client interaction and is reported as fully agentic, with continuous testing to assess whether outputs remain within established quality and accuracy thresholds.

Production

2



Application

Cybersecurity operations

Reported functionality and oversight

An application combines AI technologies to support cybersecurity operations, with decisions remaining subject to human oversight. It does not interact directly with clients.

Production

3



Application

Enhanced due-diligence assistance

Reported functionality and oversight

An AI application retrieves and summarises information obtained through enhanced due-diligence work and uses reviewed source material to draft consolidated investigation reports. It serves as a decision-support tool, with human oversight retained.

Production

4

Application

Software-development support



Reported functionality and oversight

An AI agent performs multiple steps within the software-development workflow, including researching technical issues, writing and debugging code, testing software and analysing data. The application is classified as decision support, with human-in-the-loop controls applying specifically to operations that write or execute code.

Production

5

Application

Algorithmic execution and hedging



Reported functionality and oversight

An application uses deep reinforcement learning to generate intraday execution and hedging decisions, subject to electronic-trading governance, risk limits and monitoring. Trading users retain oversight of hedging activities, and the application does not interact directly with clients.

Production

6



Application

Request-for-quotation processing

Reported functionality and oversight

An AI application restructures requests for quotation into the format required for market-intelligence and sales workflows, with users reviewing the application's responses. There is no direct client interaction.

Pilot

7



Application

Margin-call and collateral processing

Reported functionality and oversight

An AI application extracts information to support partially automated margin-call and collateral-processing and booking workflows. Users review the extracted information and approve booking. Exceptions are handled manually, and validation controls apply.

Production

8



Application

Client onboarding and know-your-client document processing

Reported functionality and oversight

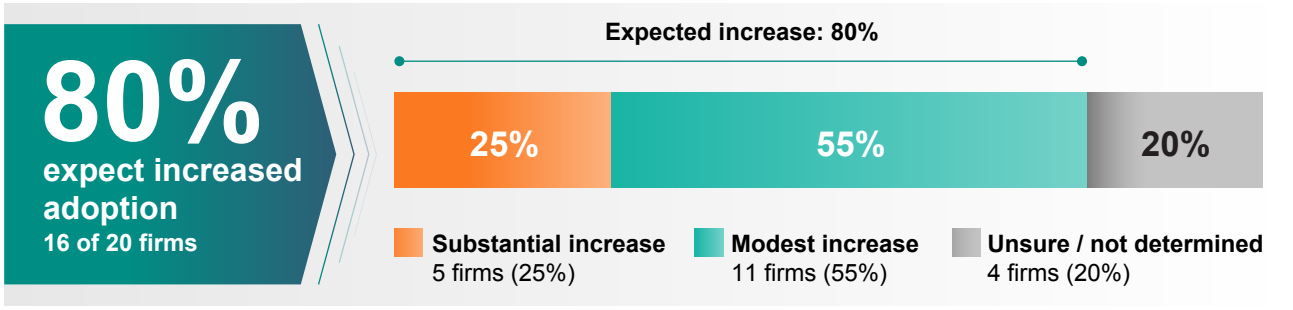
An AI application extracts information from regulated client-onboarding and know-your-client documentation containing client-identifying data. Human validation of the extracted data is required, supported by guardrails, access controls, observability and data-lineage controls.

Production

These examples illustrate different ways of combining agentic AI with human involvement, rather than a single required oversight model. Human oversight need not entail manual approval of every output or action; the relevant considerations include who defines the system's objectives and authority, how activity is monitored, and when and how people can intervene, proportionate to the application's potential impact.

Looking ahead, five firms expect a substantial increase in agentic AI adoption and 11 a modest increase; four are unsure or have not determined their plans. Taken together, 16 of the 20 firms anticipate further adoption over the next 12 months, indicating that interest extends beyond firms already permitting live production use. The survey results suggest that incremental expansion is more widely anticipated than a substantial step-up in deployment. This outlook sits alongside differing starting points: some firms already permit agentic applications in production, others are testing them in controlled environments, and the remainder are still developing or considering their approach. The responses therefore point to continued expansion at different speeds, rather than a uniform move towards greater autonomy.

Figure 6: Direction of travel for use of agentic AI

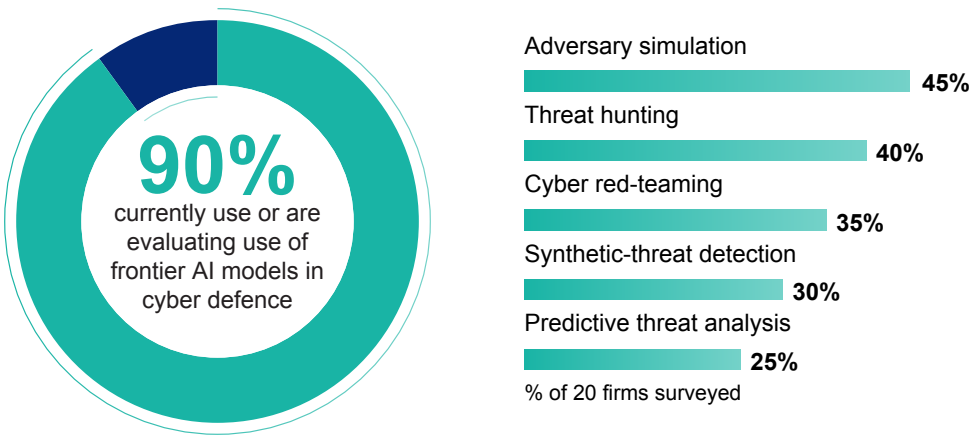


Survey finding. The responses point to growing adoption of agentic AI, but the levels of maturity and autonomy differ. Some firms already permit live production use, while others remain in controlled testing or development, and most expect adoption to increase. Reported applications extend beyond retrieval and drafting to multi-step software development, trading and operational workflows. The examples show that automation and human involvement coexist: systems may gather information or perform intermediate tasks autonomously while users retain review or approval over consequential actions. Agentic AI is therefore taking shape through a range of application-specific arrangements rather than a uniform shift towards fully autonomous decision-making.

2.5 Cyber resilience: frontier AI is a defensive opportunity as well as a risk driver

The survey asked respondents to indicate whether their firm was using, or actively evaluating, frontier AI models to enhance cyber defence against AI-enabled or other advanced cyber threats. They were asked to identify the relevant applications and describe their maturity, the types of models used, and observed benefits, effectiveness or limitations. Where such models were neither used nor being considered, respondents were asked to explain the reasons.

Figure 7: Cyber defence applications where frontier AI is being used or under active evaluation



A total of 18 firms (90%) reported using or actively evaluating frontier AI for cyber defence, indicating widespread interest but differing maturity. Adversary simulation was selected by nine firms (45%), followed by AI-assisted threat hunting by eight (40%), cyberattack simulation or defensive red-teaming by seven (35%), detection of synthetic threats by six (30%) and predictive threat analysis by five (25%).

Respondents described a range of deployment stages, from production capabilities to controlled pilots and early evaluation. Production examples included AI-enabled threat detection, code scanning and security analysis, sometimes alongside proofs of concept or planned extensions. Pilot applications included assistance with penetration testing, vulnerability discovery, threat hunting and security-operations investigations. Other firms were evaluating vendor tools or conducting contained proofs of concept to compare models before deciding on longer-term deployment. Some responses combined established AI-enabled security tools with newer frontier-model initiatives, illustrating that firms are building on different technological starting points.

Reported benefits centred on the speed and breadth of analysis and the ability to direct specialist attention towards more significant findings. One respondent described faster and broader testing, validation of legacy scan findings and improved analyst productivity. Another reported earlier vulnerability discovery, wider detection coverage, faster triage and better prioritisation across code scanning and security-detection activities. A further respondent stated that AI-assisted source-code analysis could identify vulnerabilities missed by regular penetration testing. Other firms described expected improvements rather than demonstrated outcomes, while several provided no specific observations on effectiveness.

The responses given in the survey also highlighted the effort required to turn model outputs into reliable defensive action. One respondent cited false positives and the need to validate findings during operation, retaining human approval before remediation or response actions. Another emphasised the skills and supporting tooling required for effective use and described a cautious approach to adoption. The respondent reporting vulnerabilities missed by conventional testing also noted that the work required substantial effort and a dedicated expert task force. These examples suggest that faster analysis can coexist with significant demands for verification and specialist judgement.

Survey finding. The survey responses point to widespread interest in frontier AI as a means of strengthening cyber defence, but differing levels of deployment maturity. Applications span testing defences against simulated attacks, identifying threats and discovering vulnerabilities, with reported benefits in speed, coverage and prioritisation. Some firms have moved into production, while others remain in pilots or active evaluation. Respondents' accounts also show that AI-assisted analysis continues to rely on validation and specialist judgement, with some retaining human approval over remediation or response actions. This pattern suggests that frontier AI is developing as a complement to cybersecurity expertise and established defensive processes, rather than a uniform shift towards autonomous cyber defence.

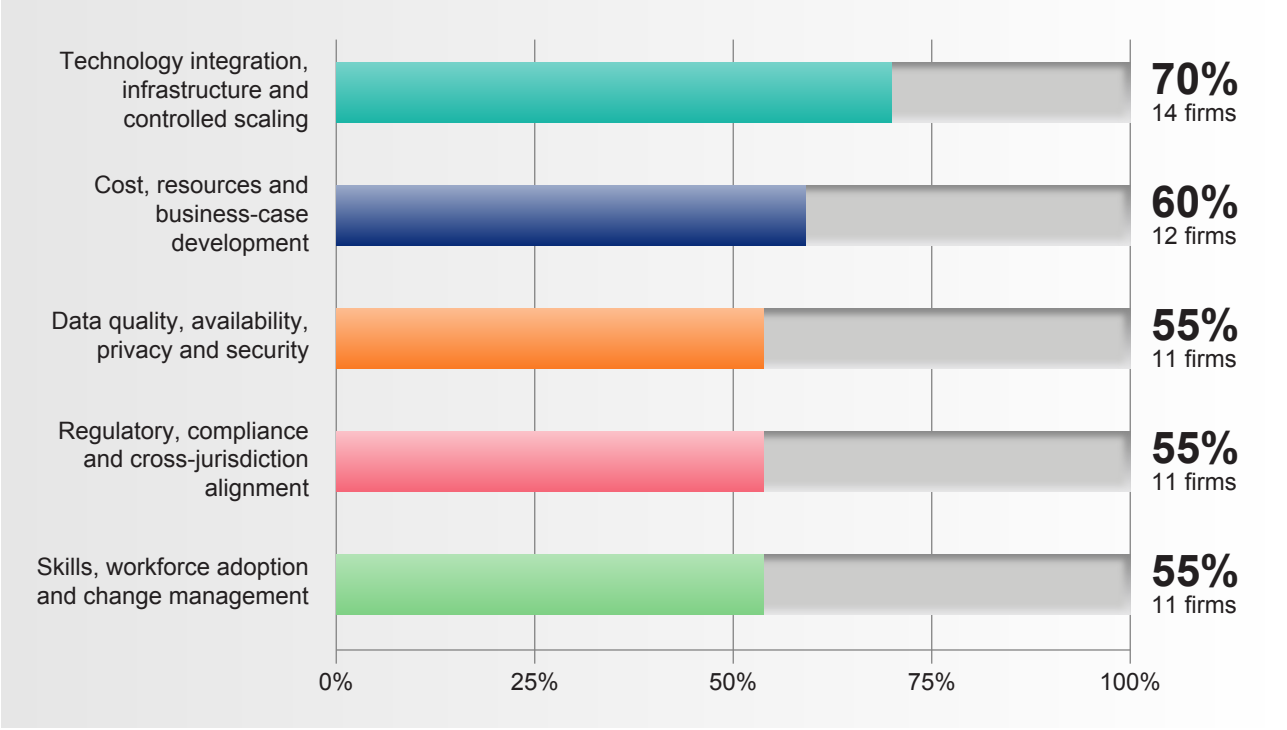
3. Practical barriers and challenges for AI adoption

The survey explored two related but distinct constraints: barriers to broader AI adoption and challenges in scaling governance and risk management. The following subsections present these separately. Themes may overlap, and their frequency does not indicate severity or establish control failures.

3.1 Barriers to broader adoption

The survey asked respondents to identify and explain the five main barriers or challenges currently limiting broader AI adoption within their firms. The open-ended question sought their own assessments of the factors constraining adoption, rather than requiring them to select from predetermined categories. The responses were subsequently grouped into the themes presented below.

Figure 8: Barriers to broader AI adoption



Technology integration, infrastructure and controlled scaling. 14 firms (70%) identified difficulties in integrating AI with existing technology and scaling applications beyond pilots. Legacy applications, fragmented data pipelines and compatibility constraints can require substantial architectural changes before AI tools can be embedded in operational workflows. Respondents highlighted the effort involved in establishing secure data connectivity, identity and access controls, audit logging, monitoring and operational support. Computing capacity, model-response times and dependencies on external interfaces can further complicate implementation. These responses indicate that demonstrating an application's feasibility is only one stage of adoption; moving into resilient production also requires a supporting technology and control environment capable of sustaining its use.

Cost, resources and business-case development. 12 firms (60%) cited constraints relating to investment, resources or the demonstration of business value. Expenditure extends beyond acquiring a model or software licence to specialist staff, computing infrastructure, cloud services, model usage, testing, governance and ongoing support. Some respondents noted rising usage costs and the need for greater cost transparency and controls over consumption. At the same time, benefits such as improved productivity, customer experience and risk management can be difficult to measure consistently or attribute directly to AI. These challenges complicate prioritisation and funding decisions, particularly where the full costs of operating and controlling an application emerge as deployment expands.

Data quality, availability, privacy and security. 11 firms (55%) reported data-related barriers to broader adoption. Fragmented sources, inconsistent definitions, unstructured information and restricted access can limit the availability of reliable data for AI applications. Respondents also highlighted confidentiality obligations, concerns about exposing client or proprietary information to external models, and cross-border data restrictions. Preparing information for use therefore involves not only improving its completeness and accuracy, but also establishing its provenance, permitted uses and appropriate access protections. The findings underline that firms need to develop data access and safeguards together. Making more information available does not, by itself, make an AI application reliable or appropriate for a regulated activity.

Regulatory, compliance and cross-jurisdiction alignment. 11 firms (55%) identified challenges in applying evolving regulatory and compliance requirements to AI deployment. Respondents sought practical clarity on matters including model governance, data protection, cloud services and automated decisions. Internationally active groups also reported difficulties in reconciling local requirements, data-sovereignty restrictions and differences in technology availability with group-wide policies and deployment approaches. These factors can lead to additional reviews, region-specific implementations and repeated testing or approval. The responses support continued regulatory dialogue and practical engagement, while recognising that a coherent group framework must still accommodate applicable Hong Kong requirements.

Skills, workforce adoption and change management. 11 firms (55%) cited constraints involving specialist expertise, staff readiness and changes to established working practices. Respondents reported shortages of personnel combining AI engineering and data skills with financial-sector knowledge and an understanding of risk and compliance. Adoption can also be constrained by uneven AI literacy, concerns about output accuracy and reliability, and the effort required to incorporate approved tools into existing workflows. Training and clear communication are therefore important not only for encouraging effective use, but also for helping staff understand the technology's limitations, their responsibilities and the role of human review. These findings suggest that sustainable adoption depends on organisational readiness and changes to working practices as well as technical capability.

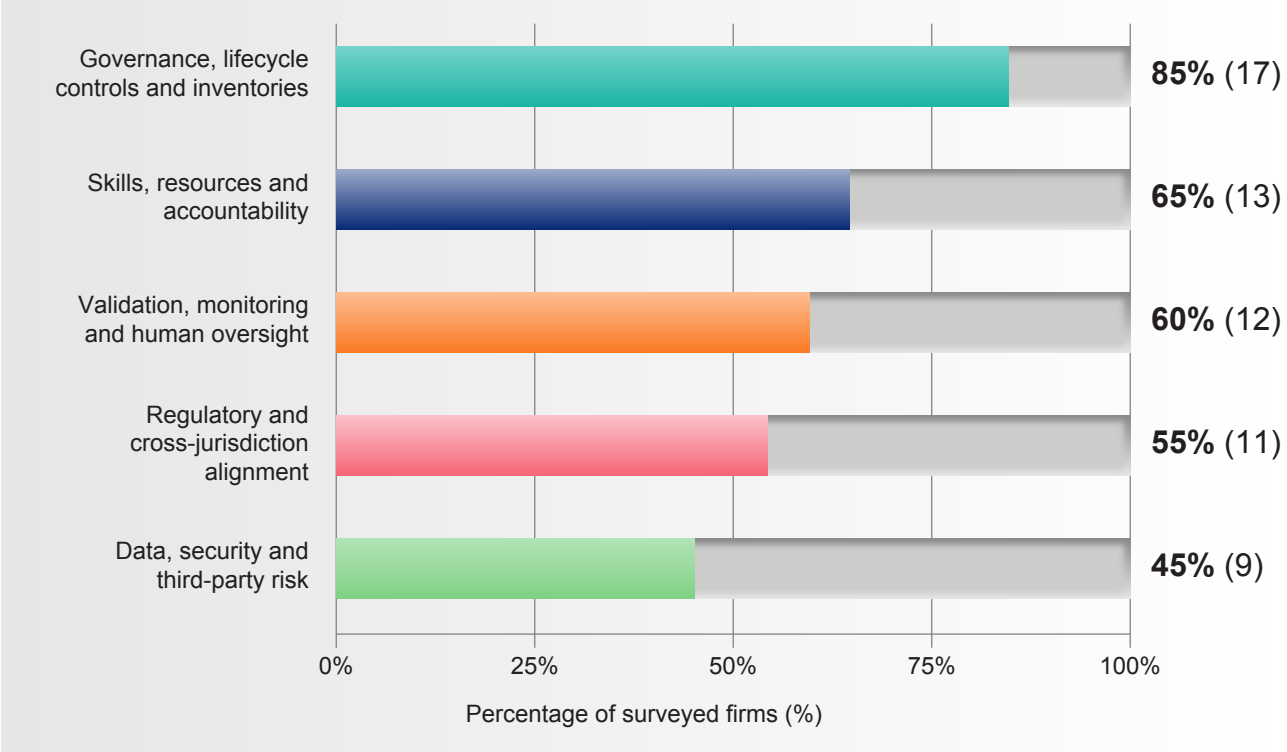
Survey finding. The responses show that ambitions for broader AI adoption coexist with practical difficulties in embedding applications into everyday operations. Technology integration and the transition from pilots to sustained production are the most widely reported barriers, alongside costs and uncertainty over business value, data constraints, evolving requirements across jurisdictions, and workforce readiness. Taken together, these themes suggest that firms' ability to expand AI use depends as much on their supporting infrastructure, information and organisational capabilities as on the models themselves. The major challenge is moving from technically feasible applications to reliable, economically viable tools that staff can use effectively within established business processes.

3. Practical barriers and challenges for AI adoption

3.2 Challenges in governance and risk management

The survey asked respondents to identify the five main challenges in scaling up or fully implementing their firm's AI governance and risk management framework. The open-ended question focused on the practical difficulties of extending governance and controls as AI use develops. Responses were subsequently grouped into the themes presented below, with each firm counted no more than once per theme.

Figure 9: Share of 20 surveyed firms reporting each theme as a challenge



Scaling governance, lifecycle controls and inventories. 17 firms (85%) identified challenges in extending governance arrangements across expanding portfolios of AI models, applications and use cases. Respondents highlighted fragmented inventories, rapidly changing vendor capabilities and the increasing availability of AI features embedded in third-party products. These developments make it difficult to maintain a complete and current view of AI use and apply consistent controls throughout development, approval, deployment, monitoring and retirement. Firms also described the need to calibrate review processes to risk so that resources remain focused on material applications without creating unnecessary bottlenecks for lower-risk uses. The challenge is therefore not simply to establish governance policies, but to translate them into scalable processes that remain effective as adoption grows.

Skills, resources, accountability and organisational readiness. 13 firms (65%) reported challenges in providing sufficient expertise, tools and capacity to support AI governance. Growing demand for applications places pressure on the teams responsible for risk assessment, validation, approvals and monitoring, particularly where specialist technical knowledge must be combined with financial-sector and regulatory expertise. Respondents also highlighted the importance of clear responsibilities across business, technology, data, risk, compliance and legal functions. Distributed governance arrangements can create coordination difficulties where end-to-end ownership, escalation routes or decision-making authority are unclear. Effective implementation therefore depends on both adequate resources and an operating model that assigns accountability clearly and supports timely, informed decisions throughout the AI lifecycle.

Model validation, monitoring, explainability and human oversight. 12 firms (60%) identified difficulties in testing and overseeing AI applications whose outputs can vary and whose underlying models or data may change. Respondents reported challenges in validating generative outputs, detecting model drift and anomalies, maintaining traceability and scaling independent review across growing volumes of activity. Limited transparency into third-party models can further constrain understanding of their behaviour and limitations. Human oversight also needs to remain meaningful without relying on manual review processes that become impractical as usage expands. These responses underline the need to connect pre-deployment assessment with ongoing monitoring, appropriate audit trails and clearly defined human-review arrangements, rather than treating approval as a one-off assurance exercise.

Regulatory change and cross-jurisdiction alignment. 11 firms (55%) cited challenges in keeping governance frameworks aligned with evolving regulatory expectations and different requirements across jurisdictions. Respondents reported the effort required to update policies, control standards, documentation and training while maintaining consistency across entities and business lines. For internationally active groups, this involves reconciling group-wide frameworks with local requirements and deployment constraints, rather than assuming that a single implementation can be applied uniformly. Rapid technological change adds to this burden by altering the capabilities and risks that governance arrangements must address. The findings support continued regulatory engagement and coordinated internal change management, while recognising that enterprise-wide consistency must accommodate applicable local obligations.

Data, privacy, security and third-party or concentration risk. Nine firms (45%) identified challenges in governing data and external dependencies as AI adoption expands. Reliance on third-party models, cloud infrastructure and data providers can limit visibility into training data, model changes and supporting controls, while creating confidentiality, contractual and operational-resilience concerns. Respondents highlighted issues involving data-use rights, provider due diligence, audit rights and the ability to secure appropriate contractual protections. Dependence on a limited number of providers also raises concentration and access risks, reinforcing the importance of credible continuity, exit and transition arrangements. These challenges extend beyond initial vendor selection: firms need ongoing oversight that reflects changes in providers, models and the applications that rely on them.

Survey finding. The responses show that firms' governance challenges focus on keeping oversight aligned with expanding AI use and rapidly changing capabilities. Scaling lifecycle controls and maintaining reliable inventories are the most widely reported difficulties, alongside pressures on expertise, resources and accountability, and challenges in validation, monitoring and human oversight. Evolving requirements across jurisdictions and reliance on data and external providers add further complexity. Taken together, these themes point to a shift from establishing governance frameworks to making them work consistently across diverse applications and throughout their lifecycle. AI capabilities and ambitions are advancing rapidly, but responsible scaling depends on infrastructure, expertise and governance capabilities that take longer to build. As applications become more connected and autonomous, these challenges increasingly concern governance of the complete workflow: maintaining visibility over what AI is authorised to access and do, assigning clear ownership, and ensuring that monitoring and intervention remain effective as use expands.

4. Implications for the SFC's approach to AI

The survey points to a shift in the supervisory challenge as AI becomes more widely deployed, more deeply integrated with firms' internal systems and increasingly capable of autonomous action. At the same time, firms reported challenges in moving from pilots to resilient production, scaling governance and controls, building specialist expertise and managing growing technology and third-party dependencies.

These findings support a supervisory direction that is **safety-first, human-anchored and collaboration-led**:

- *"Safety-first"* encompasses investor protection, cybersecurity and resilience;
- *"Human-anchored"* means keeping clear human accountability and oversight at the centre of AI governance and applications as greater authority is delegated to AI;
- *"Collaboration-led"* means building knowledge, capability and supervisory innovation with industry, technology providers, academia and other authorities.

These principles underpin the four pillars of the SFC's approach below.

4.1 Investor protection: adapt safeguards as investor interactions change

Section 2.2 shows that AI is already used across client communications and client service, with more selective deployment in marketing, product distribution, investment advice and portfolio management. Section 2.4 further illustrates how AI is entering workflows such as client onboarding and know-your-client processes, with varying degrees of automation and human involvement. Together, these findings suggest that AI may increasingly shape how investors interact with firms across different stages of the investor journey.

As these interactions become more automated and personalised, the SFC should consider how existing investor protection safeguards apply to AI-enabled delivery. Relevant questions include when AI-generated interactions amount to solicitation or active marketing, when information or assistance crosses into investment advice, and how requirements relating to account opening and operation, suitability assessment and client communications should operate. Greater personalisation may also warrant closer attention to the identification and treatment of vulnerable investors.

The underlying regulatory objectives remain unchanged, but their application may need to evolve as AI changes how financial services are communicated and delivered. Licensed firms should remain responsible for compliance and investor outcomes regardless of whether an interaction is delivered by people, AI or a combination of both.

Supervisory implication: The SFC should examine how AI is changing the investor journey and assess whether existing conduct and investor protection frameworks remain effective as interactions become more automated and personalised. Initial areas of focus could include solicitation and active marketing, investment advice, account opening and operation, and the treatment of vulnerable investors.

4.2 Cybersecurity and resilience: keep pace with AI-enabled threats

Section 2.5 shows that frontier AI is increasingly relevant to cyber defence. Out of the 20 surveyed firms, 18 use or actively evaluate frontier AI for activities including threat hunting, adversary simulation, red-teaming and vulnerability detection. Firms reported potential benefits in speed, coverage and prioritisation, but also emphasised the need for verification, specialist expertise and, in some cases, human approval over remediation or response actions.

Frontier AI can strengthen defensive capabilities, but it can also increase the speed and sophistication of cyber threats. The SFC's June 2026 circular[2] highlighted the potential for AI to accelerate vulnerability discovery and exploitation, coordinate attacks across interconnected systems and lower barriers to phishing, social engineering and deepfake impersonation. As firms connect AI more deeply to internal data, systems and workflows, effective cyber defences, access controls, monitoring, containment and recovery become increasingly important to responsible adoption. The findings also reinforce the importance of cybersecurity and operational resilience as AI adoption accelerates.

Supervisory implication: The SFC should treat cybersecurity and operational resilience as fundamental to responsible AI adoption, while continuing to monitor how AI changes both attack and defence capabilities. Supervisory work should focus on whether firms' defences can keep pace with emerging threats and whether AI-enabled cyber tools are supported by appropriate expertise, verification, human oversight and effective response arrangements. On operational resilience, the SFC will assess how firms identify and manage dependencies on AI systems and third-party providers, evaluate their ability to maintain critical operations during disruptions, and establish adequate recovery arrangements.

4.3 AI governance and applications: keep humans at the centre as AI gains autonomy

The survey shows AI becoming more integrated with firms' information and systems and, in some applications, gaining greater autonomy. More than half of surveyed firms permit agentic AI in production or controlled pilots, but the authority delegated to these applications and the role of human involvement varies significantly.

The SFC will guide firms towards human-anchored, end-to-end workflow governance. This means defining objectives and boundaries, controlling access and authorised actions, monitoring activity, enabling effective intervention and maintaining clear accountability. Human-anchored governance should not require manual approval of every output or action; controls should be proportionate to an application's authority and potential impact and should avoid unnecessarily constraining beneficial uses.

Supervisory implication: The SFC should develop practical, principles-based guidance for applications involving greater autonomy or consequential actions. The focus should be on preserving human accountability and oversight across the workflow and ensuring that controls remain proportionate to what AI is authorised to access, decide and do, rather than prescribing particular technologies or imposing blanket human-approval requirements.

4.4 Engagement and collaboration: build knowledge and capability together

Collaboration-led supervision should have three practical dimensions:

- **Knowledge:** maintain a current understanding of fast-moving technologies, risks and controls
- **Capability:** address common gaps in expertise and infrastructure that impede responsible adoption
- **Supervisory innovation:** develop practical applications that improve regulatory reporting, monitoring and risk identification

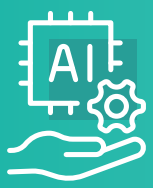
The SFC should build on existing engagement through more continuous knowledge building and practical collaboration with industry, technology providers, academia and other authorities. This can support shared understanding of areas such as agentic AI, cyber resilience and end-to-end workflow governance. Collaboration can address common capability gaps, but it should not dilute individual firms' responsibility for the risks, compliance obligations and outcomes of their own AI deployments.

Initiatives such as CrypTech and the planned Business Risk Management Questionnaire (BRMQ) revamp illustrate a broader opportunity: use collaboration not only to discuss AI, but also to redesign how regulatory information is generated, exchanged and analysed. Better data, common standards and technology-enabled processes can reduce reporting frictions while improving the timeliness and effectiveness of supervision.

Supervisory implication: The SFC should pursue continuous knowledge creation, targeted capability building and practical supervisory experimentation, while preserving each firm's responsibility for its own AI risks and outcomes.

[2] Circular to licensed corporations, SFC-licensed virtual asset service providers and associated entities: Enhanced cybersecurity measures to address evolving risks arising from artificial intelligence-enabled cyberattacks

5. Conclusion and direction of travel



AI adoption among large firms is expected to continue growing, while becoming broader, more integrated and autonomous. The survey shows strong momentum across business, operations and control functions, although deployment remains uneven and applications differ materially in maturity, authority and human involvement.



The constraints are increasingly shifting from model capability alone to safety, infrastructure, governance and expertise. Moving from pilots to resilient production requires reliable data and systems, a viable business case, specialist skills and controls that can scale across connected and increasingly autonomous workflows.



The SFC's response will therefore be safety-first, human-anchored and collaboration-led. Safety-first supervision will focus on investor protection, cybersecurity and resilience. Human-anchored governance will preserve clear responsibility across end-to-end workflows as greater authority is delegated to AI. Collaboration-led supervision will build knowledge and capability and develop new approaches to regulatory information, monitoring and risk identification, without diluting individual firms' responsibility for their deployment.



The objective is not AI adoption for its own sake. It is to create the conditions in which useful AI can scale safely, with clear human responsibility and the capabilities needed to sustain trust.

Glossary

Term	Meaning in this report
Agentic AI	AI systems capable of autonomously determining and executing steps to achieve an objective, using result to decide the next step, including accessing information or tools and taking actions.
AI agent	An AI-enabled system or component that interprets an instruction or objective and iteratively selects and takes actions, potentially using data, tools or other systems. An agent may operate alone or within a multi-agent system.
AI inventory	A maintained register of AI models, systems and applications that supports oversight of their purposes, responsible owners, deployment status and material dependencies. For agentic systems, the inventory also supports visibility over material agents and their authorised activities.
AI lifecycle	The stages through which an AI model or application passes, from development or acquisition and assessment through deployment, operation, monitoring and change to decommissioning. Governance and accountability extend throughout these stages.
Anti-money laundering (AML)	The policies, procedures and controls used to prevent and detect the laundering of criminal proceeds, including risk assessment, customer due diligence, ongoing monitoring, suspicious transaction reporting and record-keeping.
Artificial intelligence (AI)	A family of technologies enabling machine-based systems to infer from data or other inputs how to perform tasks, including problem-solving, recommendations, prediction, decision-making and content generation.
Assets under management (AUM)	The value of assets managed for clients under asset-management mandates.
Autonomy	The degree to which the system selects its approach, makes intermediate decisions and acts without contemporaneous human direction or approval. Level of autonomy should be assessed separately from automation: an automated rules-based process may have little decision-making autonomy.
Customer relationship management (CRM)	Processes and supporting systems for recording and managing customer information, interactions and service relationships. In this report, CRM refers principally to internal customer-data systems connected to AI applications.
Cyber threat hunting	A proactive, evidence-led search across networks, systems and security data for malicious activity or compromise that existing detection controls may not have identified. It is distinct from investigating only automatically generated alerts.
Data-lineage	A traceable record of data's origins, movement, processing and transformations from source to subsequent use, supporting an understanding of how an output was produced.
Deep reinforcement learning	A machine-learning approach combining deep neural networks with reinforcement learning, in which an agent learns a policy for selecting actions through interaction with an environment and feedback expressed as rewards or penalties.
Explainability	The extent to which the operation of an AI system and the basis for its outputs can be made understandable to relevant users and oversight functions. Explanations should reflect the application and audience and acknowledge technical limitations.

Glossary

Frontier AI	Advanced AI models at or near the leading edge of available capabilities, potentially applicable across a broad range of tasks.
Generative AI	AI that generates content, including text, images, audio or computer code, in response to inputs, using patterns learned from training data.
Guardrails	Technical or procedural safeguards that constrain an AI system's permitted inputs, data access, outputs or actions within defined boundaries. They may be incorporated into its design or applied during operation.
Human-in-the-loop	An oversight arrangement in which a person participates directly at specified stages of an AI-assisted process, such as validating an output or approving a consequential action before it is executed. It differs from supervision in which a person only monitors an otherwise automated process.
Know your client (KYC)	Processes for establishing a client's identity and obtaining information needed to understand the client and the intended relationship.
Licensed corporation (LC)	A corporation licensed by the SFC under Part V of the Securities and Futures Ordinance to carry on one or more regulated activities, within the scope and conditions of its licence.
Model drift	A change over time in a model's performance or behaviour that may cause it to cease operating as originally intended, including when input data or operating conditions change.
Model validation	A structured assessment and testing process to determine whether a model is fit for its intended use, covering relevant inputs, outputs and associated systems, its performance and limitations, and material risks. Validation is undertaken before deployment and revisited following significant changes, alongside ongoing performance monitoring.
Observability	The ability to assess a system's operational state and behaviour using recorded evidence, including logs, metrics and traces of inputs, outputs, tool use and actions. Observability supports oversight and incident investigation but does not itself establish model explainability.
Operational resilience	The ability to deliver critical operations through disruption, supported by the capacity to identify and protect against threats and failures, respond and adapt, and recover and learn from disruptive events.
Pilot	A controlled trial or limited deployment used to evaluate an application's operation, risks and user experience before deciding whether to expand its use.
Production	Use of an application in a live operational environment to perform its intended business function, rather than solely for feasibility assessment or controlled testing.
Red-teaming	Authorised adversarial testing conducted within an agreed scope to identify weaknesses and assess preventive, detection and response capabilities. In cybersecurity, it simulates relevant attacker behaviour; AI red-teaming may instead test model or application safeguards.

Request for quotation (RFQ)	A request to one or more dealers or counterparties for a price quotation for a specified financial instrument or transaction. An RFQ initiates price discovery or negotiation and is not itself an executed trade.
Synthetic threat	For this report, artificially generated or manipulated content or activity used to deceive, impersonate or facilitate a cyber threat. The term is broader than deepfakes and does not imply that all synthetic content is malicious.
Third-party concentration risk	Risk arising from reliance on a limited number of external providers, technologies or services, including common dependencies that can expose multiple activities or firms to the same disruption or failure and limit substitutability.

* This glossary is for the purposes of this survey only.