Chartis

HAWK

AI in Financial Crime and Compliance

Charting the Path from Pilot to Maturity

Banking Edition

Contents

Executive summary

This report explores how financial institutions are adopting, applying and governing artificial intelligence (AI) across their financial crime, fraud and compliance (FCC) functions. The survey examines the current state of and attitudes toward AI, and its expected and realized benefits. It also considers respondents' expectations for the future, to ascertain where institutions and the industry are on their journey toward AI maturity.

The findings reveal an industry in transition, from early experimentation to strategic deployment.

Survey findings

- **AI adoption in financial crime compliance is moving from pilots to production**. Nearly all banks now deploy AI in some form, with one in three having it fully operationalized in fraud prevention and over 20% in anti-money laundering (AML) transaction monitoring. Traditional machine learning (ML) and natural language processing (NLP) lead the way, but firms are rapidly advancing toward next-generation capabilities: 92% plan to increase their investment in generative AI (GenAI) and 84% in agentic AI over the next two to three years. This surge reflects a transition from experimentation to operational deployment, as banks embed AI more deeply.

- **The impact is growing**. 48% of respondents saved more than $1m through AI in the past year.

- **Most financial institutions and regulators now encourage AI adoption, signaling confidence**. 62% of financial institutions are planning to increase their adoption of AI, and 75% expect their regulators to be pro-AI moving forward.

- **Firms continue to face skills shortages and integration challenges**. 60% ranked internal expertise as their primary business challenge, while 62% ranked poor quality data as their main technical challenge.

- **Agentic AI is seen as a key driver of future progress**. While it can improve investigative depth and automate manual work, it raises questions about auditability, compliance and workforce impact.

The industry is entering a new stage of maturity, embedding intelligence into operations and strengthening risk detection while reinforcing control and oversight. The focus is no longer on whether firms are using AI, but how effectively and safely they are scaling it.

> *'It's going to be hard to keep up if organizations don't start adopting these tools; bad actors certainly will. We need to be using the same kinds of technologies to defend against the threats they're using to attack.'*
>
> **Chief Compliance Officer, Tier 2 Financial Institution**

Introduction and methodology

This study is based on responses from 125 financial institutions across the banking sector.[1] Respondents represent a mix of organization types (see Figure 1), including traditional banks (44%), credit unions and mutual or cooperative banks (28%) and neobanks or digital/challenger banks (28%). Participants were compliance, risk and technology professionals involved in financial crime prevention, AML and fraud management. Data was collected via an online survey to assess firms' current and planned use of AI, its associated benefits, its challenges and firms' investment outlook over the next three years. Several respondents also participated in interviews.

Figure 1: Respondent's organization type

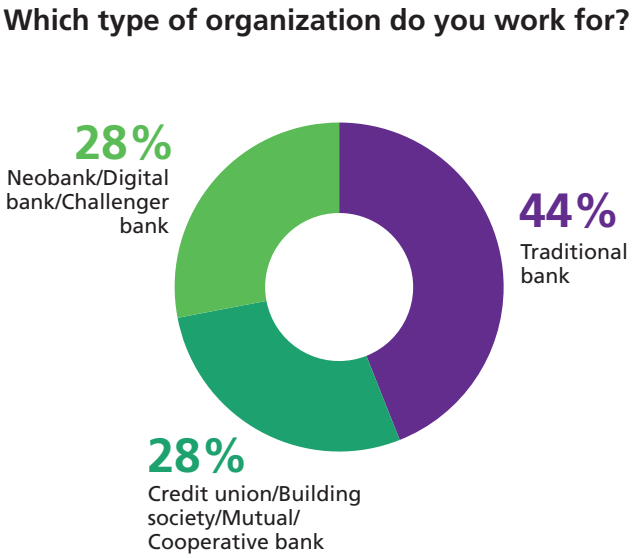

Source: Chartis Research

[1] *Note that not all percentages may add to exactly 100%, due to rounding.*

Current attitudes toward AI

Overview

To shed light on the industry's overall AI mindset, the survey explores how financial institutions perceive and position AI within their anti-financial crime, anti-fraud and compliance functions. Key questions include attitudes toward AI, use of AI, which types of AI are used most, and the degree to which AI is embedded.

AI is widely used and encouraged among respondents

A significant majority of respondents (61%) report that their organization encourages AI use and they themselves actively use it at work (see Figure 2), while another 28% say their organization supports AI but they personally do not yet use it. Only a small minority work in organizations that are neutral to AI (9%) or discourage its use (2%).

Clearly, banks are pro-AI: 89% encourage AI use, according to their compliance staff. Only 1% said their employer discouraged it. However, staff themselves are lagging in engagement, with 28% not yet using AI even though their employer encourages it.

The overall picture is one of widespread institutional endorsement of AI across firms.

Figure 2: Attitudes to and use of AI

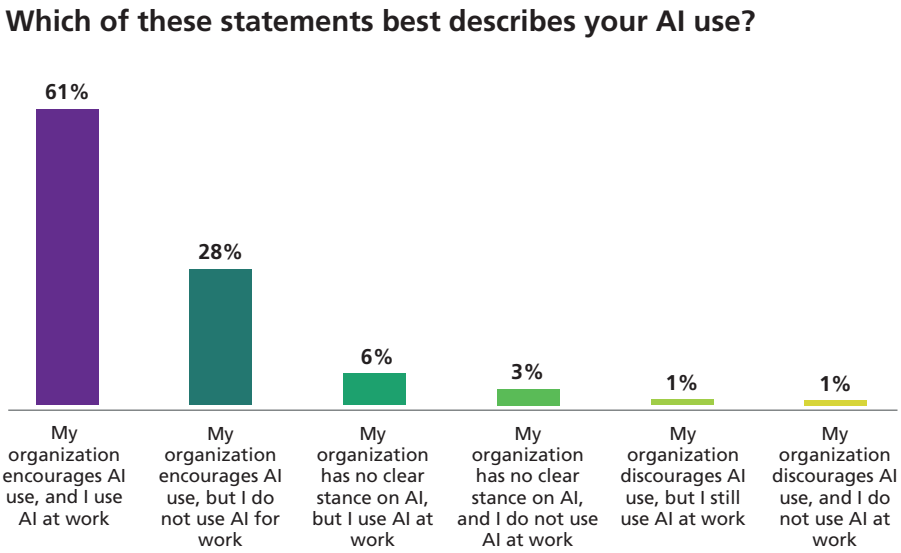

Source: Chartis Research

Almost all organizations are using AI, but many banks remain in the pilot stage of adoption

Nearly half of respondents (48%) report that their organizations are testing or piloting AI (see Figure 3), while another 22% are still in the exploration stage, underscoring that most firms remain in early to mid-adoption phases. Only 16% have AI operational and just 6% are using it at scale. This indicates that while there is strong interest and experimentation around AI, enterprise-wide deployment has not yet occurred.

Figure 3: Stage of AI adoption

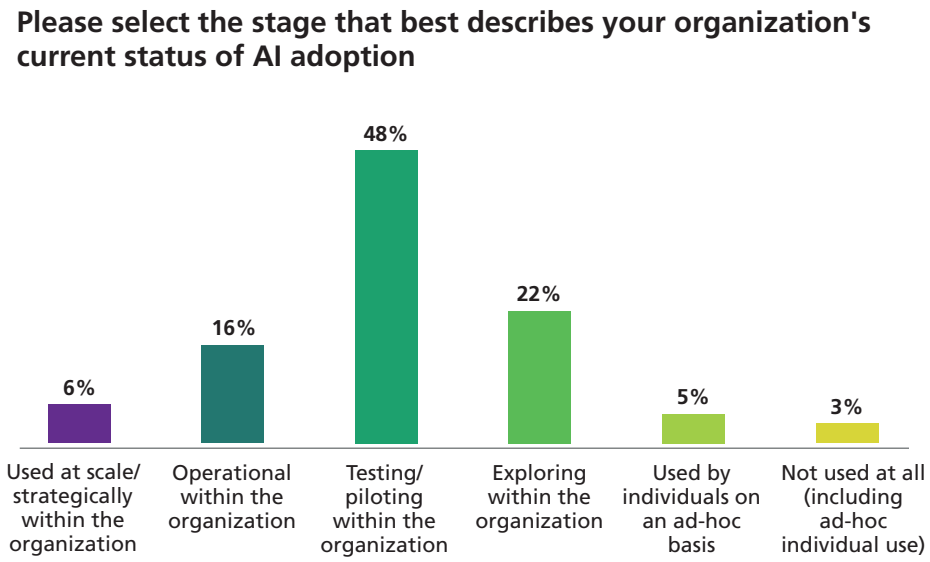

Source: Chartis Research

Fraud has most AI adoption across FCC domains

Banks are actively using AI to some degree across all aspects of their financial crime and compliance programs. Fraud prevention leads the way in AI adoption (see Figure 4), with 100% using AI at some level. This may reflect the relative ease of applying ML to fraud compared with compliance-heavy domains, which must contend with low risk appetites and strict requirements for explainability and model governance. Regulatory reporting is the least advanced area of AI use, demonstrating perhaps that regulation is still an area of concern for most banks.

One in ten firms are already using AI at scale for fraud, compared with 6% for screening and 5% in AML. At the other end of the spectrum, only 7% say that AI is ad hoc in fraud, compared with 26% who do not use AI or only use it ad hoc in AML transaction monitoring. This rises to 48% for regulatory reporting and 35% for case management.

Figure 4: Adoption of AI in different business areas

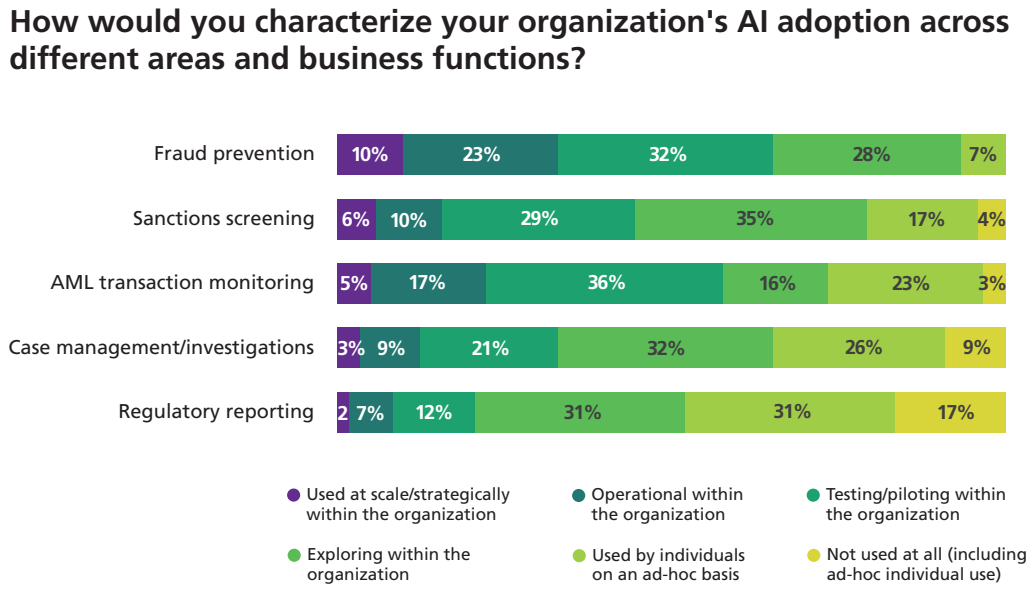

Source: Chartis Research

Regional analysis – Europe and North America lead the way

Adoption in Europe and North America is relatively widespread (see Figures 5 and 6), with organizations piloting and operationally using AI across all financial crime functions.

Figure 5: Adoption of AI in different business areas (Europe)

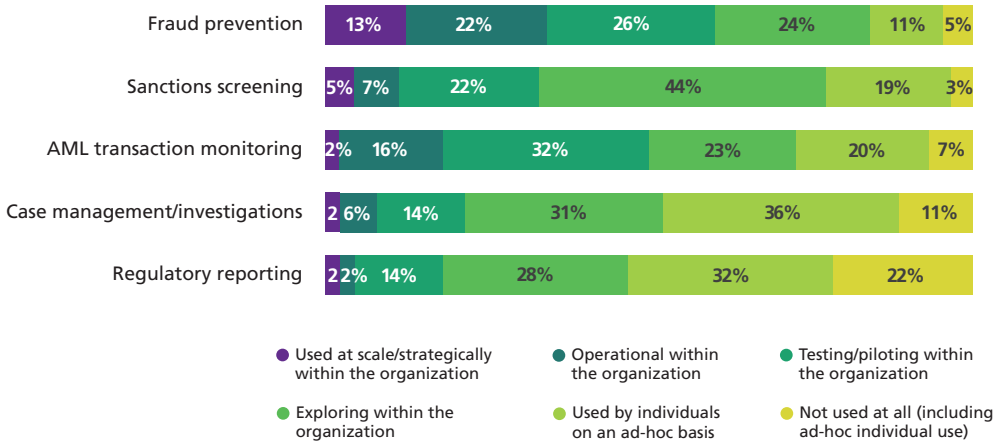

Source: Chartis Research

Figure 6: Adoption of AI in different business areas (North America)

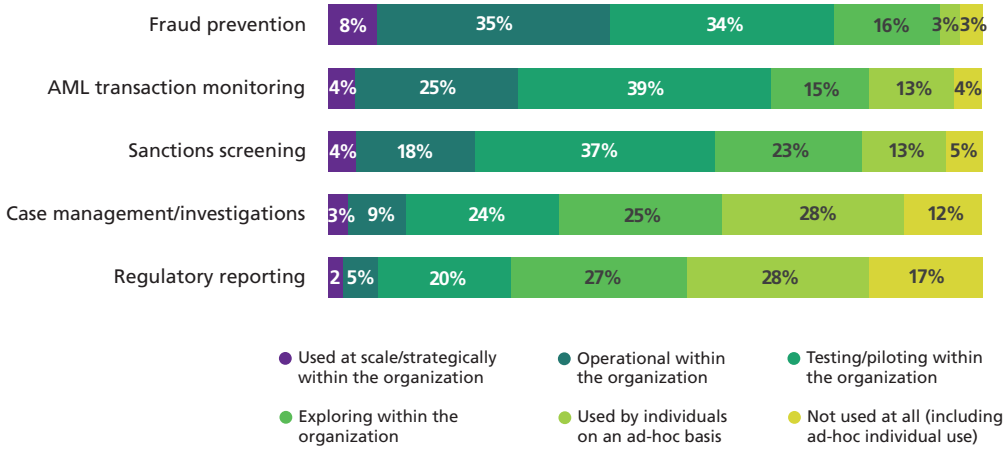

Source: Chartis Research

Latin America is following a developing adoption path but is at an earlier stage than North America (see Figure 7). AI use is concentrated in pilots, with fewer instances of enterprise-wide deployment.

Figure 7: Adoption of AI in different business areas (Latin America)

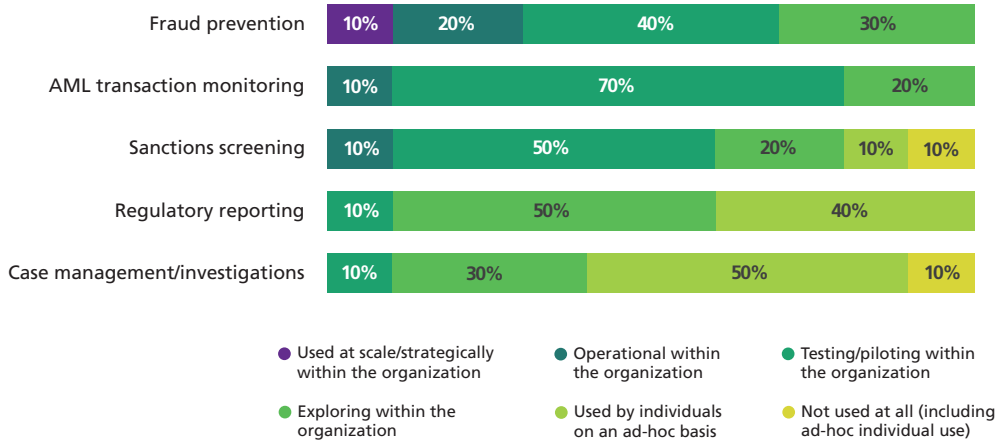

Source: Chartis Research

The Middle East & Africa region shows the lowest maturity overall (see Figure 8), with most firms still exploring or piloting AI; no respondents in this region stated that their firms are using AI at scale or strategically.

Figure 8: Adoption of AI in different business areas (Middle East/Africa)

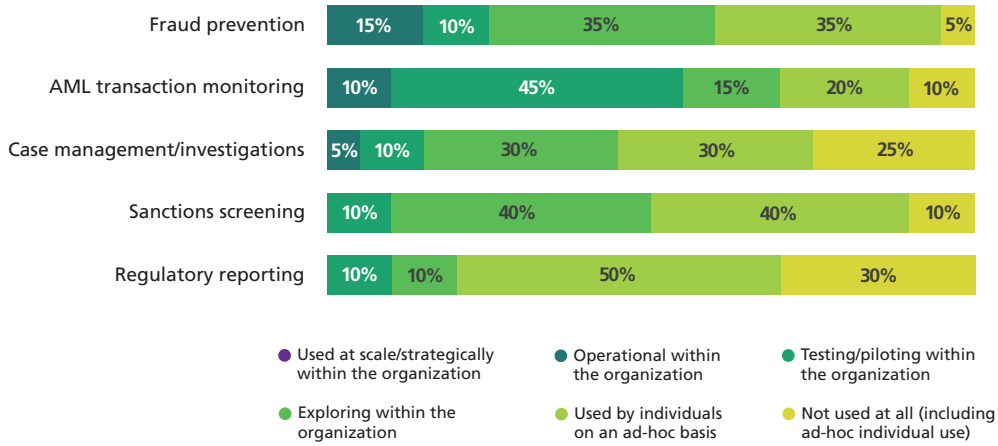

Source: Chartis Research

APAC exhibits a broad but uneven adoption profile (see Figure 9 on page 8), with a focus on fraud, sanctions screening and AML transaction monitoring. Respondents are active in testing and experimenting, which could reflect areas of fast-moving innovation (in Singapore, Hong Kong and Australia, for example) but varying regulatory alignment.

Figure 9: Adoption of AI in different business areas (APAC)

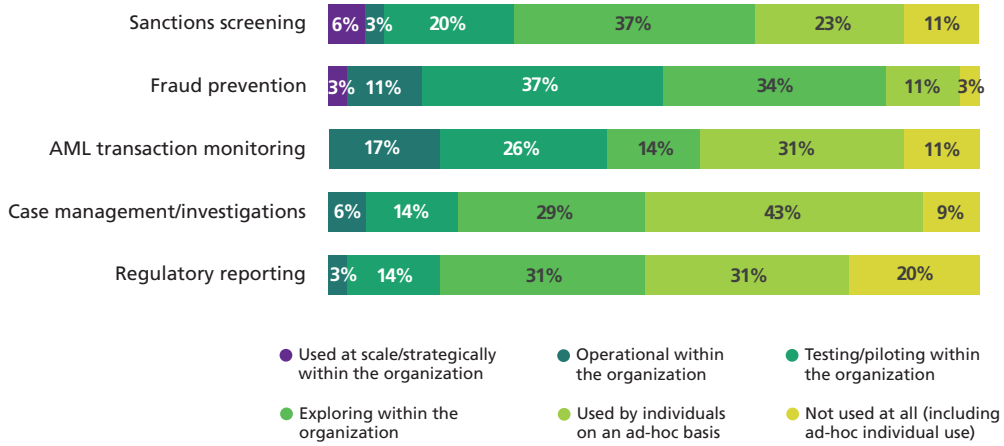

Source: Chartis Research

Machine learning leads the way in AI technique adoption

AI/ML is the dominant technique across all functions (see Figure 10), peaking in case management/investigations (75%). NLP/text analysis is also being applied widely, particularly in case management (61%), regulatory reporting (52%) and sanctions screening (51%), where unstructured data and documentation are central.

Agentic AI adoption is still limited (11-21% of respondents in most areas, and just 4% in regulatory reporting). So far, agentic AI is most likely to be applied in case management/investigations (21%).

The results indicate that ML functions as the backbone, NLP as an enabler for text-heavy tasks and generative/agentic AI are still emerging, at least within the context of financial crime and compliance.

Figure 10: AI techniques used in the business

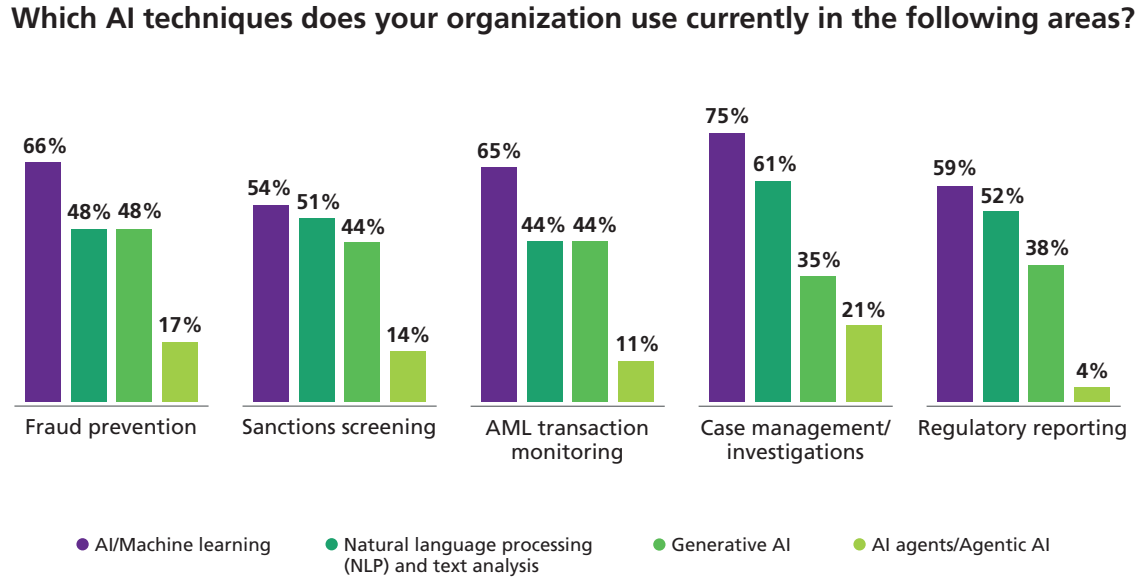

Source: Chartis Research

AI risks and challenges

Overview

Given firms' growing use of AI, the survey also explored the challenges and barriers that continue to influence their adoption and governance of AI in their anti-financial crime and compliance functions. Key questions consider organizational hurdles (including skills gaps, resourcing constraints and stakeholder alignment), as well as technical and regulatory challenges.

Business challenges: expertise is essential

The leading business challenge (see Figure 11) is insufficient expertise/resources (60% rank this number one), showing that many firms struggle to support AI systems during development and after deployment. Next are regulatory concerns (64% rank it their number two concern, with 83% putting it in the top three), reflecting persistent uncertainty over compliance expectations.

Figure 11: Main business challenges in AI adoption

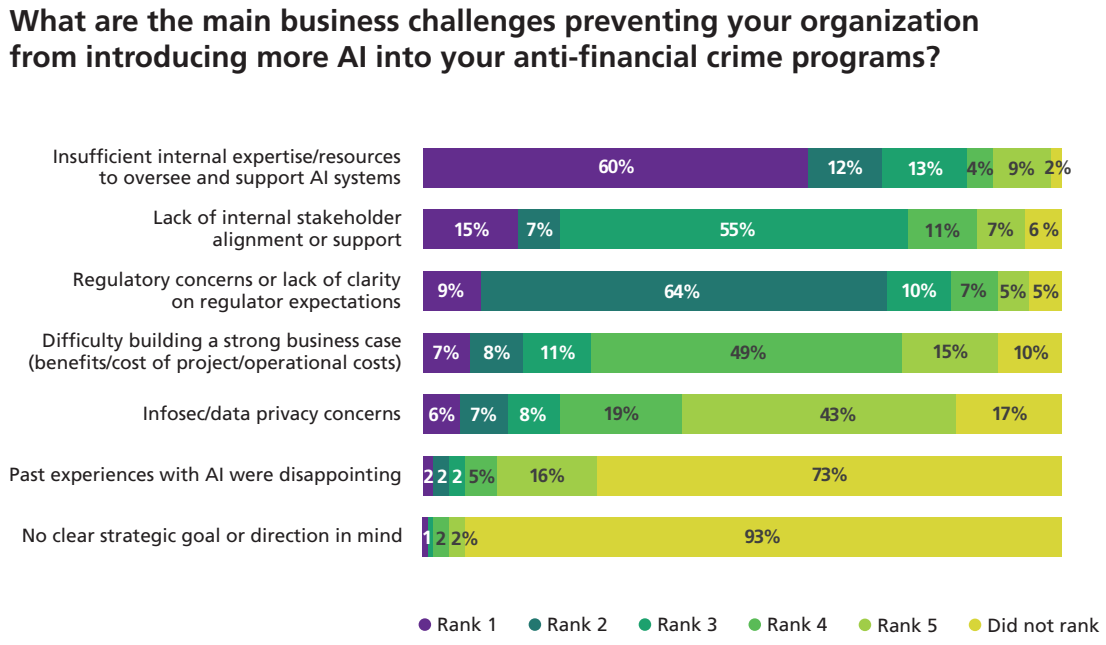

Source: Chartis Research

Changing business challenges: regulation and data

Most business challenges have become less of a concern as firms have adopted AI (see Figure 12). The two clear exceptions to this are regulatory concerns, with 85% saying concerns have increased or stayed the same, and infosec/data privacy concerns with 81%.

Figure 12: Evolution of business challenges as AI matures

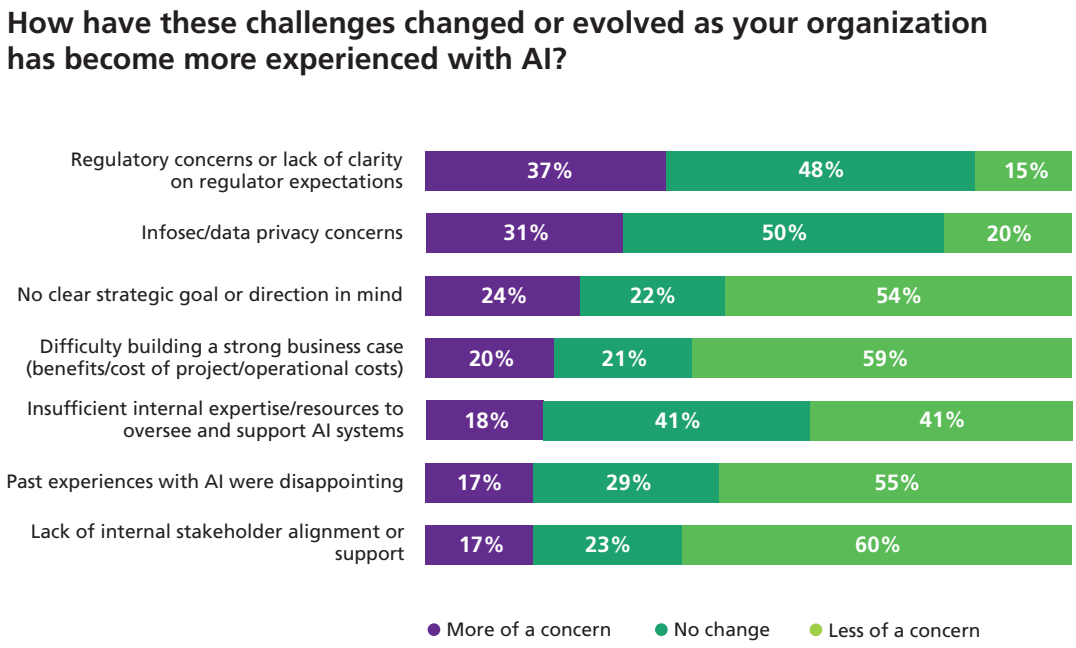

Source: Chartis Research

By contrast, issues such as unclear strategy, business case development and stakeholder alignment diminish over time (with 54%, 59% and 60%, respectively, saying it is less of a concern), suggesting that firms gain confidence and direction once their AI initiatives mature. Notably, insufficient expertise/resources remain a persistent bottleneck (with 41% reporting no change and 18% seeing them as more of a concern), while disappointing past experiences fade in importance as use cases improve (55% said they are less of a concern).

Overall, the data suggests that while internal alignment and strategic focus strengthen with AI maturity, external pressures become more pertinent as business hurdles to long-term adoption.

Technical challenges: data and integration dominate

The leading technical barrier to AI adoption in anti-financial crime programs (see Figure 13 on page 11) is limited or poor-quality training data (62% rank it as their top concern), underscoring the ongoing importance of foundational data's availability and integrity.

Integration with existing systems is the next most pressing technical barrier, ranked second overall by 55% of respondents. However, trust in model outputs stands out as another big ongoing issue, cited among the top five concerns by 84% of firms. This reflects persistent uncertainty around explainability, auditability and the reliability of AI-driven decisions, particularly as models become more complex with the adoption of generative and agentic techniques. Model performance (ranked fifth by 45%) is closely linked, as firms grapple with maintaining accuracy and stability once models are deployed in dynamic, high-volume environments. By contrast, scalability and post-deployment maintenance appear less urgent (unranked by 54% and 58%, respectively), suggesting that most banks are now confident in their infrastructure but continue to struggle with the trust and transparency of AI decisioning.

Overall, the results point to a consistent theme: while infrastructure is improving, data quality, integration and trust remain the biggest obstacles to scaling AI effectively in compliance.

Figure 13: Main technical challenges in AI adoption

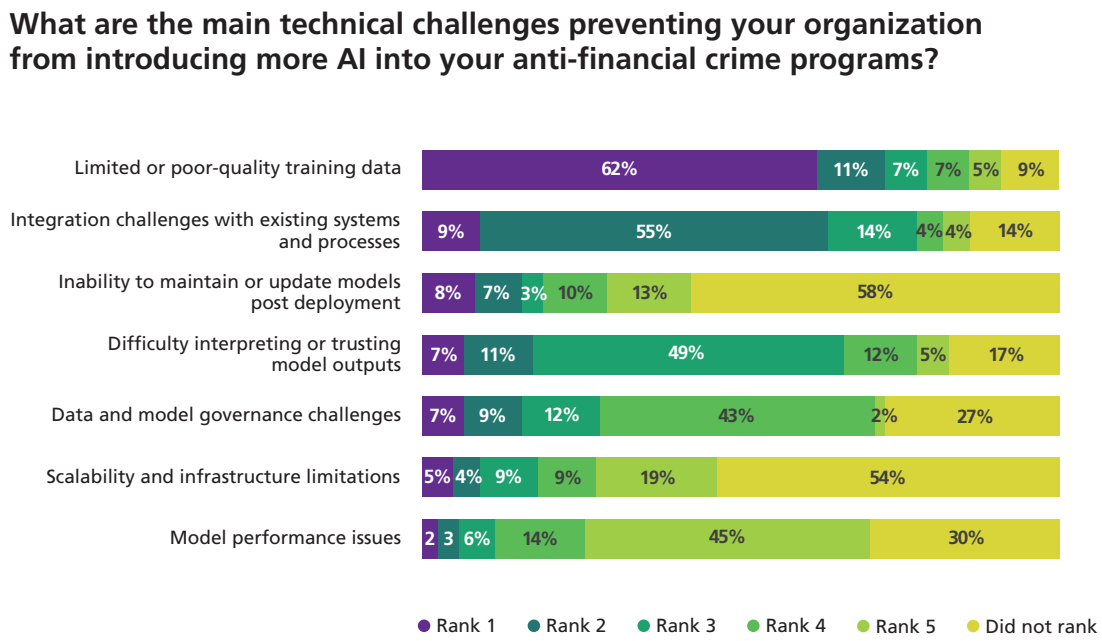

Source: Chartis Research

Changing technical challenges: integration and model management

Although training data was noted as a hurdle preventing organizations from getting started with AI, it becomes less of a concern once adoption is under way. The challenges that grow in concern include integration, trusting model outputs and maintaining and governing models (see Figure 14). Model maintenance is notable, as most respondents did not rank post-deployment model upkeep in their top five concerns, but it is clearly a growing issue.

Figure 14: Evolution of technical challenges as AI matures

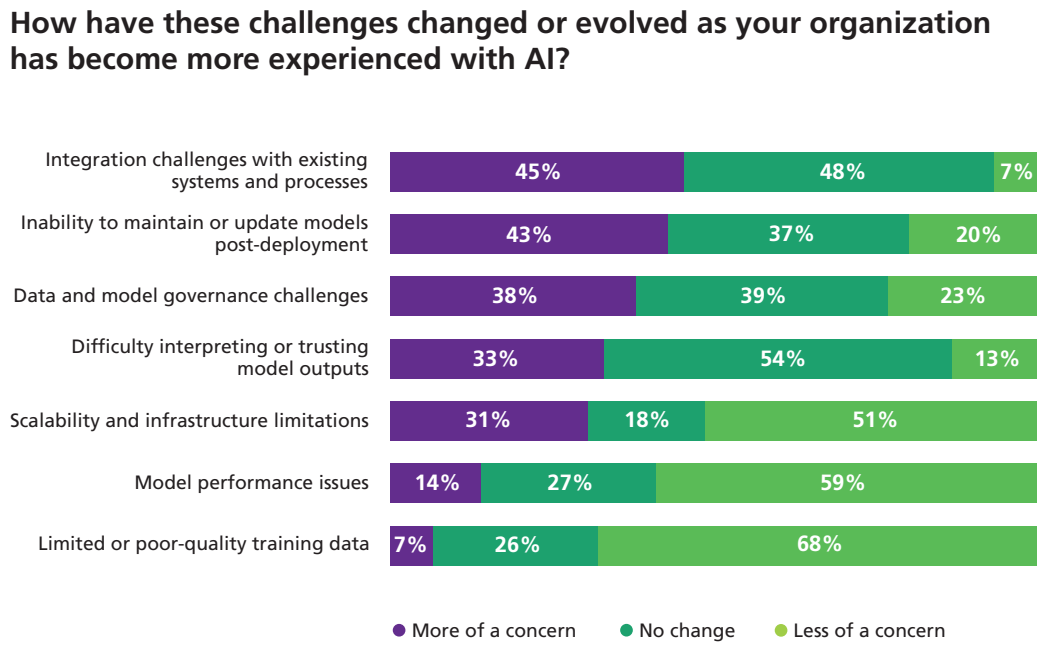

Source: Chartis Research

Data and model governance and trust in model outputs continue to be major challenges, of growing concern for 38% and 33% of respondents, respectively, reflecting the complexity involved in ensuring explainability and regulatory alignment.

By contrast, 59% cite model performance and 68% cite training data quality as diminishing concerns. This may indicate that firms are growing in maturity and confidence in their AI development. It is also worth noting that while data quality is considered the most significant of all the technical challenges, it is diminishing most as a concern. Scalability and infrastructure limitations have also eased for many, with 51% citing them as less of a concern, suggesting that technical capacity is improving.

Overall, the challenge profile has shifted: basic performance and data quality issues fade with experience, while integration, governance and lifecycle management are emerging as the technical hurdles to large-scale adoption.

AI benefits and savings

Overview

The survey also examined the value that AI can deliver across anti-financial crime, anti-fraud and compliance functions, along with how those benefits are evolving as adoption matures. Among the subjects addressed are the areas where firms are realizing gains, from improved detection accuracy and faster investigations to efficiency and cost reduction, and how these outcomes compare with initial expectations. By analyzing achieved and anticipated savings, we can highlight the real-world impact of AI investments, and how organizations can maximize their return as they progress toward AI maturity.

Firms wanted increased risk coverage and threat reduction from AI…

The top expected benefits from AI (see Figure 15) centered on increased risk coverage (64% ranking it first) and improved detection accuracy (76% ranking it in the top three). Firms also emphasized faster alert triage and investigation and the ability to scale operations with existing resources, underscoring a focus on improving analytical reach and efficiency rather than purely reducing cost.

Figure 15: Top five expected benefits of AI

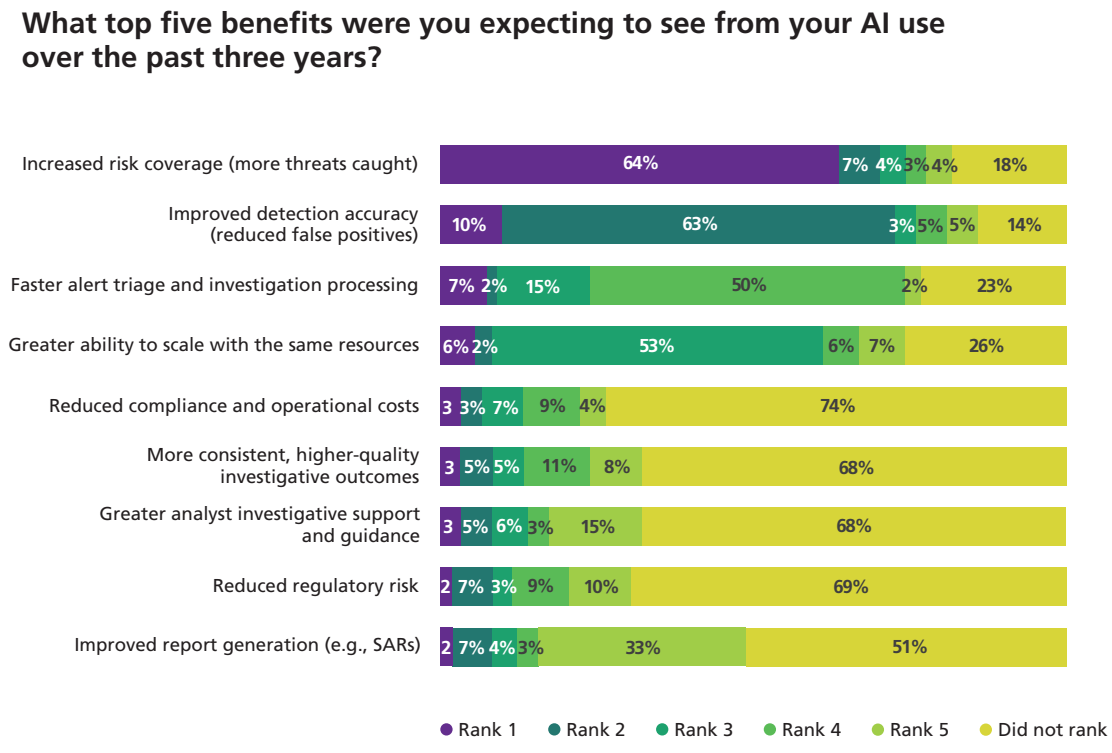

Source: Chartis Research

These priorities show that banks initially pursued AI not to cut headcount or expenses, but to detect more risk and spend more time on high-value investigative work. Cost/regulatory risk reduction and report generation ranked lower among expected benefits, suggesting that early AI strategies were mainly about enhancing detection capability and operational agility.

'We're hopeful that AI will take on more of the drudge work – the repetitive, manual tasks that slow our teams down. We don't yet know exactly what to expect, but the goal is clear: to make people's lives easier, not more complicated.'

Chief Compliance Officer,
Tier 1 Financial Institution

… but the realized cost reduction and benefits from AI have often been in efficiency gains

Compared to expectations, the realized benefits of AI show a shift toward efficiency gains (see Figure 16). The most widely achieved outcomes were faster alert triage and investigation processing (ranked in the top three by 74%) and improved detection accuracy (ranked in the top three by 66%), reflecting AI's tangible impact on day-to-day operational performance. Reduced compliance and operational costs followed as the third most realized benefit (64% in the top three), indicating that efficiency improvements are now beginning to translate into measurable savings.

Figure 16: Top five realized benefits of AI

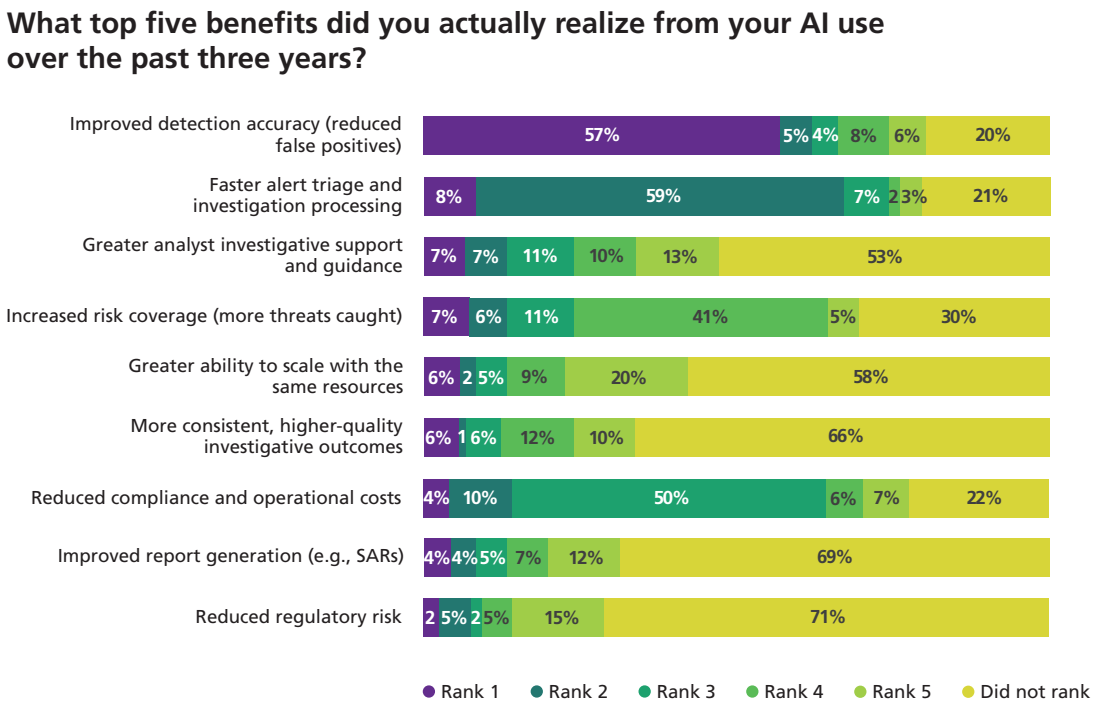

Source: Chartis Research

By contrast, qualitative gains such as analyst investigative support and more consistent investigative outcomes were not ranked among the top five by 53% and 66% of respondents, respectively, while regulatory risk reduction and reporting improvements were rarely cited at all.

Taken together, the biggest gaps between expected and realized benefits show where AI delivers the most value (see Figure 17 on page 14). The findings suggest that AI has delivered most strongly in improving investigative accuracy and speed but has yet to meet expectations around other benefits, such as scalability and report generation. This reinforces earlier insights: adoption is still in the testing and operational phases, and while performance gains are tangible, such strategic benefits as regulatory assurance remain a work in progress.

Figure 17: AI expectations vs. realities

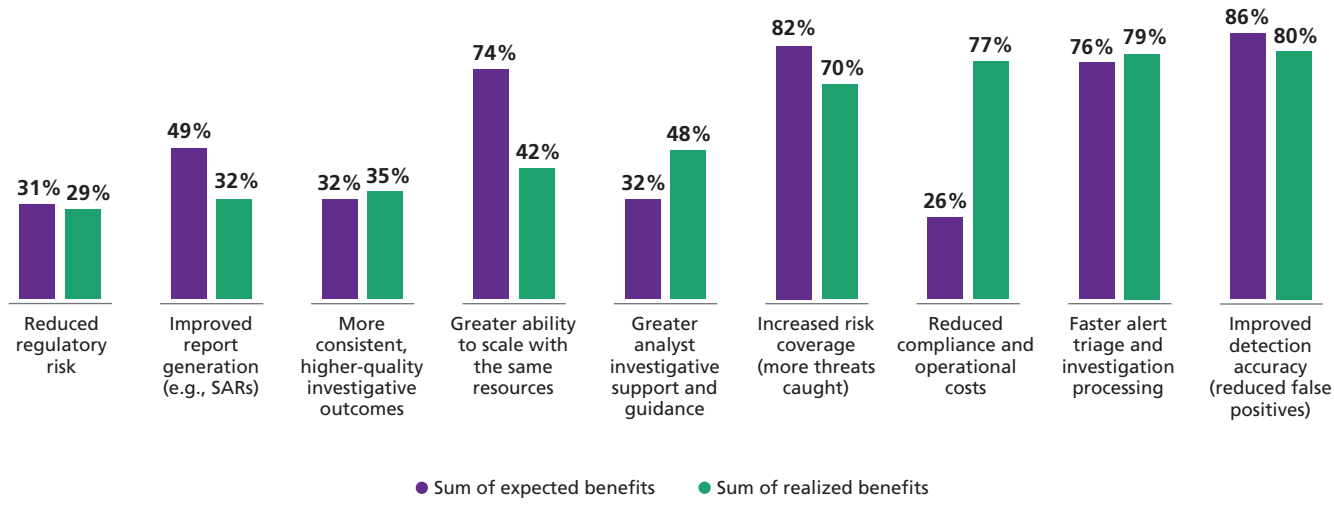

Source: Chartis Research

Cost reduction outperformed expectations: only 26% expected it, but 77% realized it. However, ability to scale with the same resources was notably low (74% vs. 42%), indicating some struggles with throughput or capacity scaling.

Cost savings in AML from AI remain modest, but are expected to grow

While banks say that immediate costs savings were not high on their agendas for AI, 71% have already seen cost savings in AML (see Figure 18). More interestingly, 94% say that they expect to see cost savings in the next 12 months, with over half expecting to save more than $5m.

Clearly firms expect significant savings from AI, increasingly seeing it as a driver of substantial future efficiency and cost reduction once adoption matures.

Figure 18: Annual savings from using AI for AML

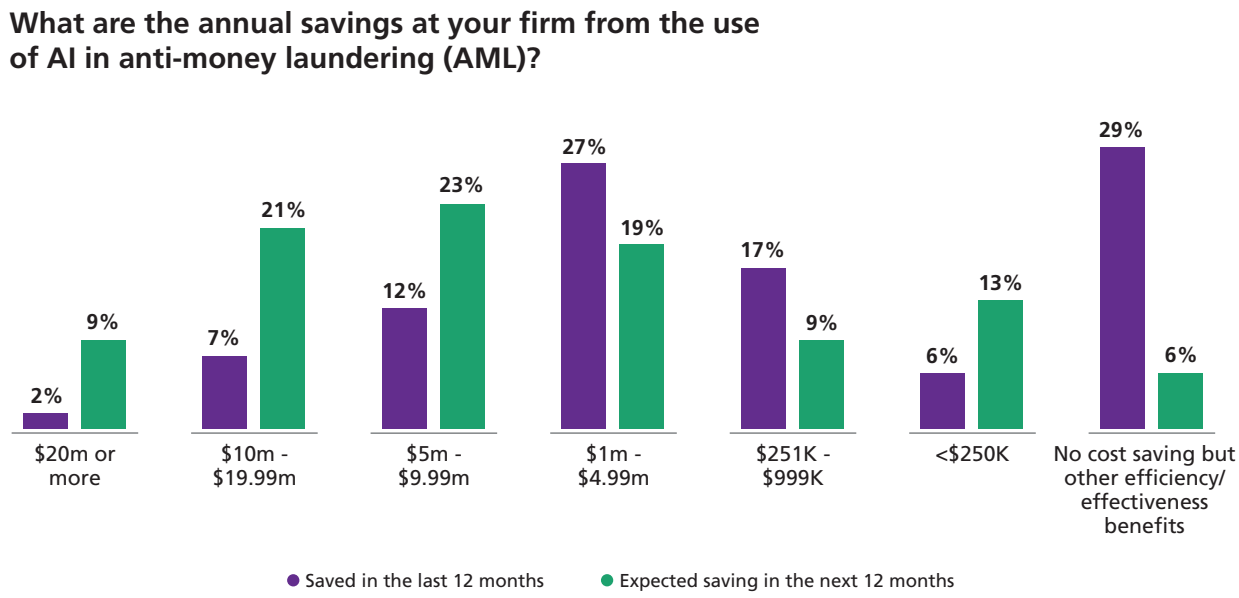

Source: Chartis Research

The future of AI in anti-financial crime, anti-fraud and compliance

Overview

The survey also examined how financial institutions expect their adoption of AI to evolve in the next two to three years, in terms of technology investment, regulatory attitudes and organizational readiness. Key questions concerned such trends as the rise of generative and agentic AI, growing expectations of regulatory support and shifting views on governance, accountability and workforce impact.

By capturing these forward-looking perspectives, we can better outline the next phase of firms' journey toward AI maturity, as they begin to deploy AI at scale.

AI is the path forward in anti-financial crime

Expectations for the next two to three years are positive (see Figure 19), with 62% of banks pushing ahead with AI adoption in financial crime and compliance. A third of respondents anticipate gradual adoption in limited areas and 28% expect significant acceleration and wider embedding of AI.

Another 30% of respondents expect to maintain their current levels of AI adoption. Only a very small percentage of respondents expect to scale back (3%) or avoid adoption (1%). Overall, the outlook suggests momentum will continue, with most firms either deepening or cautiously expanding their AI use.

Figure 19: Anticipated change in AI adoption

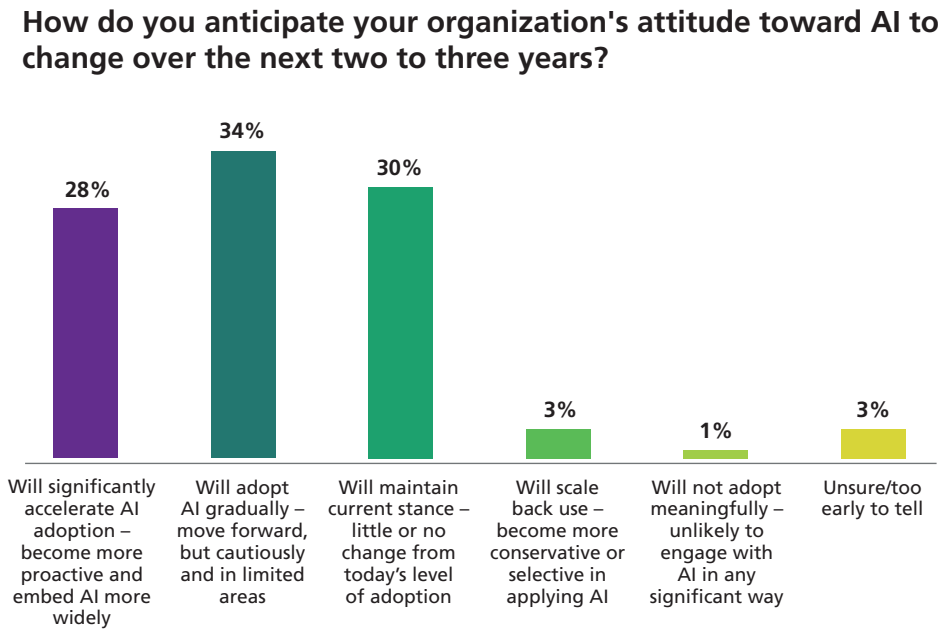

Source: Chartis Research

Firms expect regulators to become more supportive of AI

Although regulatory concerns and a lack of clarity on regulators' expectations were mentioned as being growing concerns for banks, clearly there is optimism: 60% say that they expect regulators to become more supportive of AI (see Figure 20).

Figure 20: Perceived change in regulators' attitudes to AI

Which of these statements best reflects your view as to how financial regulators' attitudes towards AI will change over the next two to three years?

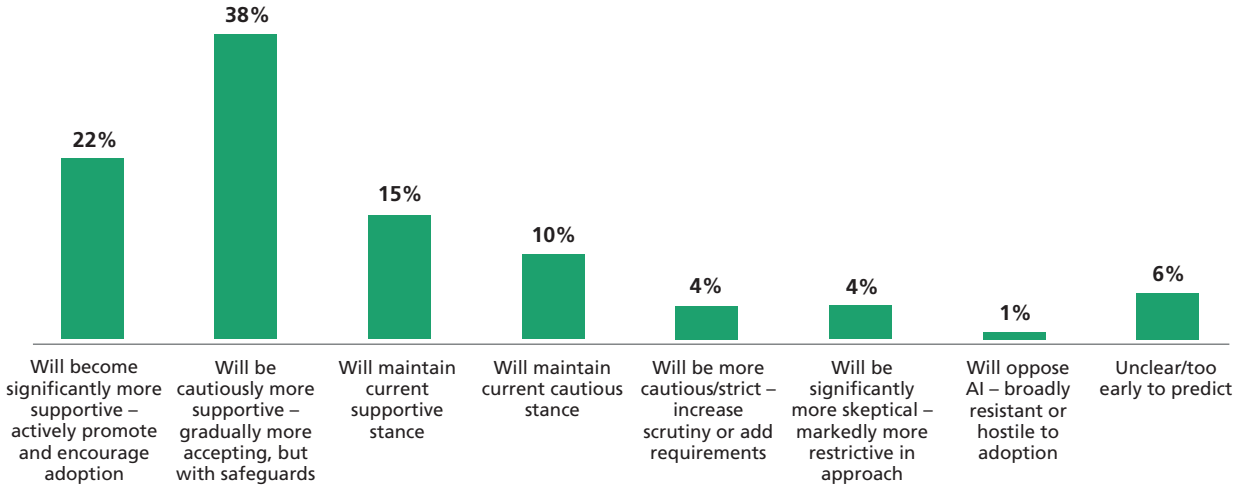

Source: Chartis Research

A smaller group anticipates little change, with 15% of respondents expecting the current supportive stance to continue and 10% expecting the current cautious stance to persist. In interviews, respondents often considered the EU AI Act to be helpful guidance, but they also emphasized its complexity. Some believed regulation in the US to be lagging, noting the difficulty of planning in an environment where there is considerable uncertainty.

Only 8% of respondents expect stricter or more restrictive approaches, and just 1% foresee outright opposition. Overall, there is some optimism that regulators will increasingly enable AI adoption in a controlled manner.

The future: a regional perspective

Table 1 on page 17 summarizes respondents' anticipated change in adoption, by region. In those markets where we have seen greater and more aggressive adoption and testing of AI, outlined in green, the expectation is that adoption will continue to be material and, in some cases, will be relatively aggressive. Overall, the responses align with the mood that Chartis has observed in the industry – a shift to more cautious but continued adoption focused on embedding AI in processes and use cases with a clearer ROI, rather than deploying it everywhere in a rush to identify potential areas of benefit.

In those regions outlined in red – primarily Africa and the Middle East, where adoption has been slower – there is greater skepticism and slower deployment. This presents an opportunity to take use cases identified in other regions and enable them in more cautious areas, with the potential for 'leapfrogging'.

As for how regulators are perceived with respect to AI, we can see a correlation between firms' approach to AI and the regulatory climate in which they operate.

Table 1: Anticipated change in AI adoption, by region

	Africa	Asia	Europe	Latin America	Middle East	North America	Pacific (non-Asia)
Unsure/too early to tell	0%	4%	3%	0%	10%	1%	0%
Will not adopt meaningfully – unlikely to engage with AI in any significant way	0%	0%	0%	0%	0%	2%	0%
Will scale back use – become more conservative or selective in applying AI	0%	8%	3%	0%	0%	2%	0%
Will maintain current stance – little or no change from today's level of adoption	0%	32%	37%	20%	70%	19%	40%
Will adopt AI gradually – move forward, but cautiously and in limited areas	100%	44%	36%	60%	10%	43%	30%
Will significantly accelerate AI adoption – become more proactive and embed AI more widely	0%	12%	21%	20%	10%	33%	30%

Source: Chartis Research

The green outlines in Table 2 on page 18 highlight areas with a more supportive regulatory approach, and these map to those areas reporting a more aggressive and proven deployment of AI. With a more supportive climate and a shift to embedded and explainable outcomes, we can reasonably project continued AI growth and deployment.

This correlation also holds in regions where institutions are more cautious (outlined in red). There is an opportunity here for the vendor community to assist regulators and institutions in these regions to become more comfortable in the capabilities of and controls around AI deployment.

Table 2: Perceived change in regulators' attitudes to AI, by region

	Africa	Asia	Europe	Latin America	Middle East	North America	Pacific (non-Asia)
Unclear/too early to predict	0%	4%	2%	10%	10%	3%	0%
Will oppose AI – broadly resistant or hostile to adoption	0%	4%	0%	0%	0%	3%	0%
Will be significantly more skeptical – markedly more restrictive in approach	30%	0%	4%	0%	20%	3%	10%
Will be more cautious/stricter – increase scrutiny or add requirements	40%	8%	6%	0%	30%	5%	0%
Will maintain current cautious stance	10%	16%	12%	10%	10%	13%	20%
Will maintain current supportive stance	0%	8%	28%	0%	10%	6%	20%
Will be cautiously more supportive – gradually more accepting, but with safeguards	20%	44%	34%	50%	10%	41%	40%
Will become significantly more supportive – actively promote and encourage adoption	0%	16%	13%	30%	10%	25%	10%

Source: Chartis Research

Firms expect AI to have significant positive effects over the next three years

Banks overwhelmingly expect further positive impact from AI on effectiveness over the next two to three years (see Figure 21). Overall, 86% anticipate positive effects from ML, 83% from GenAI, and 76% each from NLP and agentic AI.

Figure 21: Impact of AI on effectiveness in the next two to three years

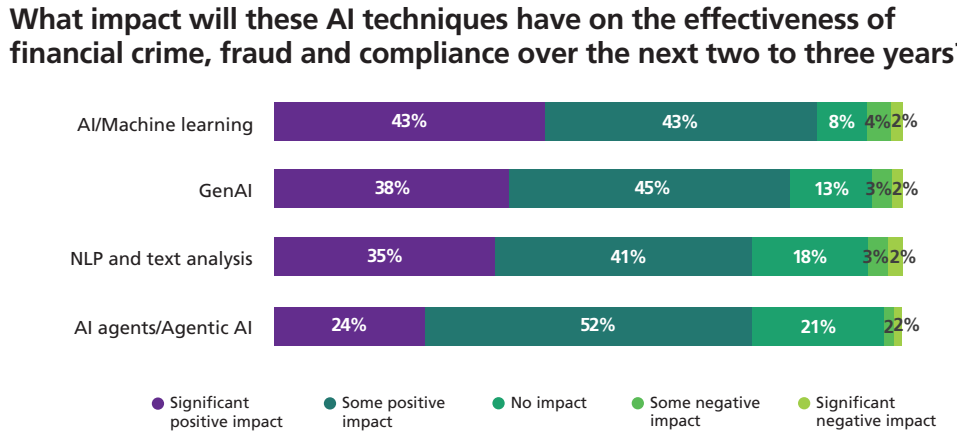

Source: Chartis Research

ML leads in expected impact, with 43% of respondents predicting a significant positive effect and another 43% a somewhat positive one, reinforcing its foundational position in driving accuracy and consistency in detection. GenAI (38% significant impact) and NLP (35%) follow closely, both viewed as key enablers for managing unstructured data and enhancing investigative workflows through contextual analysis and summarization.

By contrast, agentic AI (though newer) elicits strong optimism, with 24% expecting a significant and 52% a moderate positive impact. Its lower effectiveness ranking reflects its current role in efficiency-driven use cases such as workflow orchestration, task automation and analyst assistance, rather than direct improvements in detection accuracy. Notably, negative expectations are minimal (4%), indicating widespread confidence that agentic AI will add long-term value as its capabilities mature.

AI investment is expected to grow by more than 25%

Organizations overwhelmingly expect to increase their AI investment over the next two to three years (see Figure 22), although responses differ across technologies. For AI as a whole, 82% of respondents anticipate growth of more than 25%.

Figure 22: Expected change in AI investment

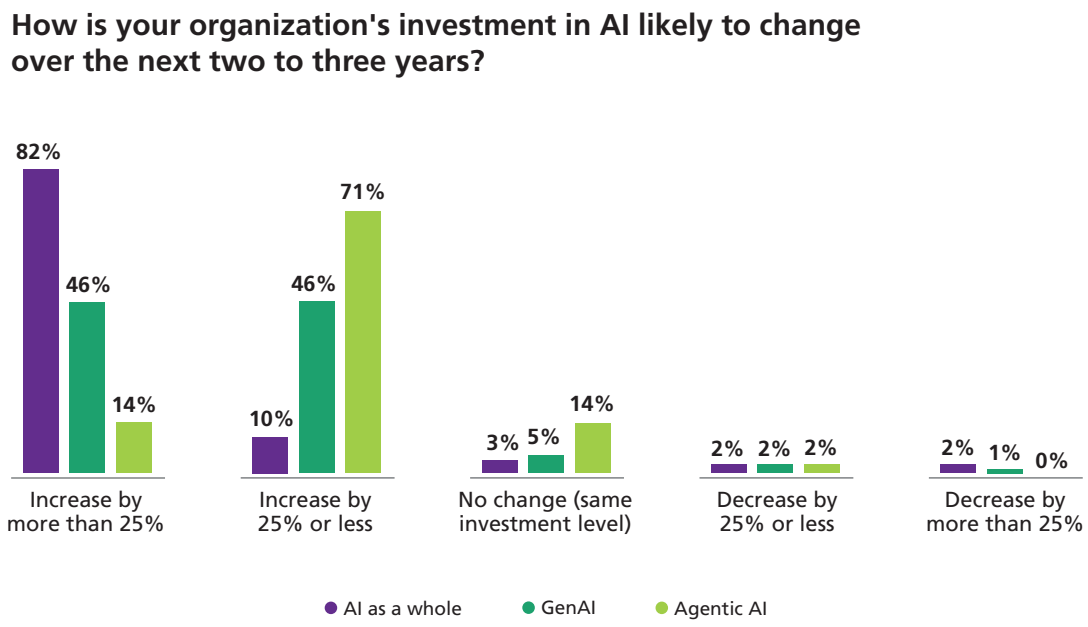

Source: Chartis Research

Respondents also show confidence in GenAI, with 92% saying that investment will increase. These respondents are split evenly between >25% increases and ≤25% increases (46% of respondents for both), reflecting both enthusiasm and measured rollout.

For agentic AI, 71% of respondents expect modest increases (≤25%) and 14% predict larger jumps, indicating that it is currently used more in proof of value/proof of concept rollouts. Very few respondents expect spending cuts across any category.

This suggests broad commitment to scaling AI, with GenAI models driving near-term investment and agentic AI still at an exploratory funding stage.

Generative AI's potential

Overview

Generative AI (GenAI) refers to a class of AI systems that can create new content, including text, images, code or data, based on patterns learned from existing information.

Unlike traditional ML models, which focus on prediction or classification, GenAI can generate original outputs, summarize complex data and support reasoning or creative tasks.

The survey explores how firms are beginning to apply GenAI in these contexts, the benefits already observed, and the risks and governance challenges that accompany its use. It also considers how GenAI differs from more established forms of AI (moving from detection and prediction to contextual understanding and content creation) and what that shift means for firms' compliance functions.

The impact of GenAI has been mixed so far

The impact of GenAI over the past two to three years has been mixed (see Figure 23).

While 28% of respondents report significant change, saying it has meaningfully reshaped team workflows, an almost equal number (27%) describe the effect of GenAI as minimal, with a limited real-world impact. Another 20% say it has had a moderate impact (enhanced but not transformed some processes), with 20% indicating it has had no impact at all.

This split highlights the uneven maturity of GenAI. Some organizations are already realizing tangible benefits, while others remain cautious or unconvinced, underscoring both the potential and the uncertainty that surround its role in compliance and anti-financial crime functions.

Figure 23: The impact of GenAI

Source: Chartis Research

GenAI's impact focuses on operational efficiency

Respondents considered GenAI's greatest potential to be in improving data and document processing, areas where automation can directly reduce manual effort, with 58% giving this their highest ranking (see Figure 24 on page 21). 51% ranked improved investigative efficiency and guidance second. Model training impacts both effectiveness and efficiency, and 47% of respondents ranked this third in importance.

Less common benefits include faster regulatory interpretation (66% of respondents did not rank) and enhanced internal training or onboarding (67% did not rank).

Overall, firms expect GenAI's near-term impact to center on operational efficiency, investigative insight and model optimization, rather than on analytics-focused tasks. This aligns with the broader trend of task-level adoption before enterprise-scale transformation.

Figure 24: Biggest potential benefits of GenAI adoption

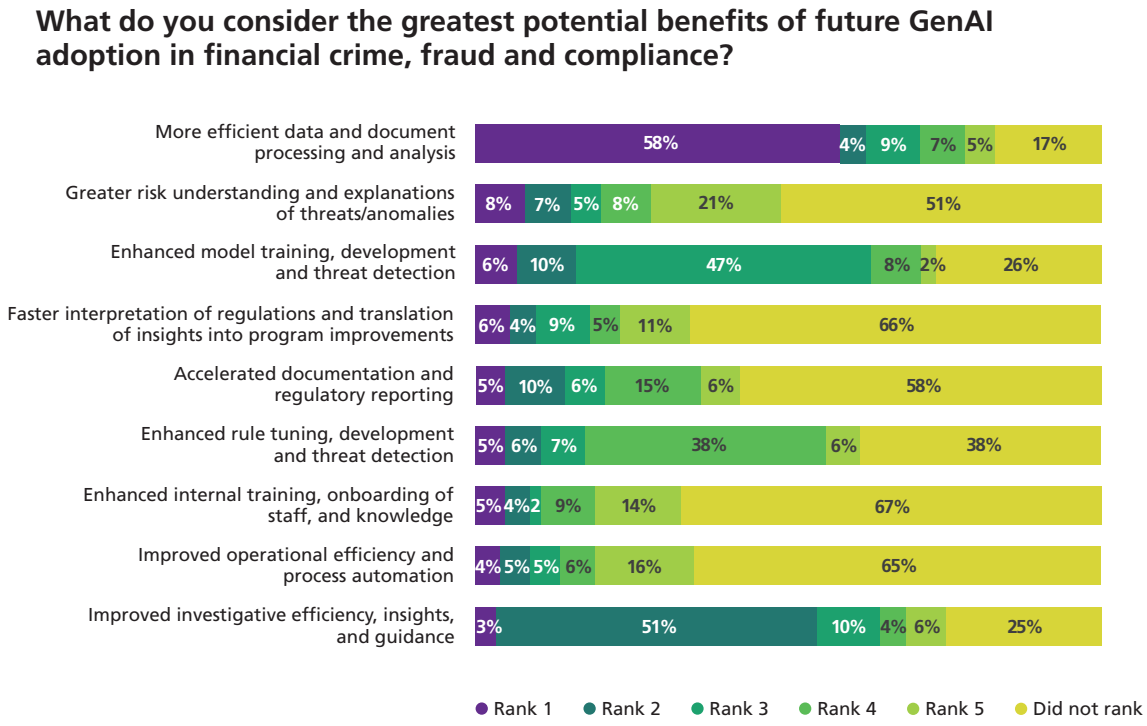

Source: Chartis Research

GenAI risks: trust is key

The most significant perceived risks surrounding future GenAI adoption include hallucinations, data security, and explainability and auditability (see Figure 25).

Figure 25: Biggest risks of adopting GenAI

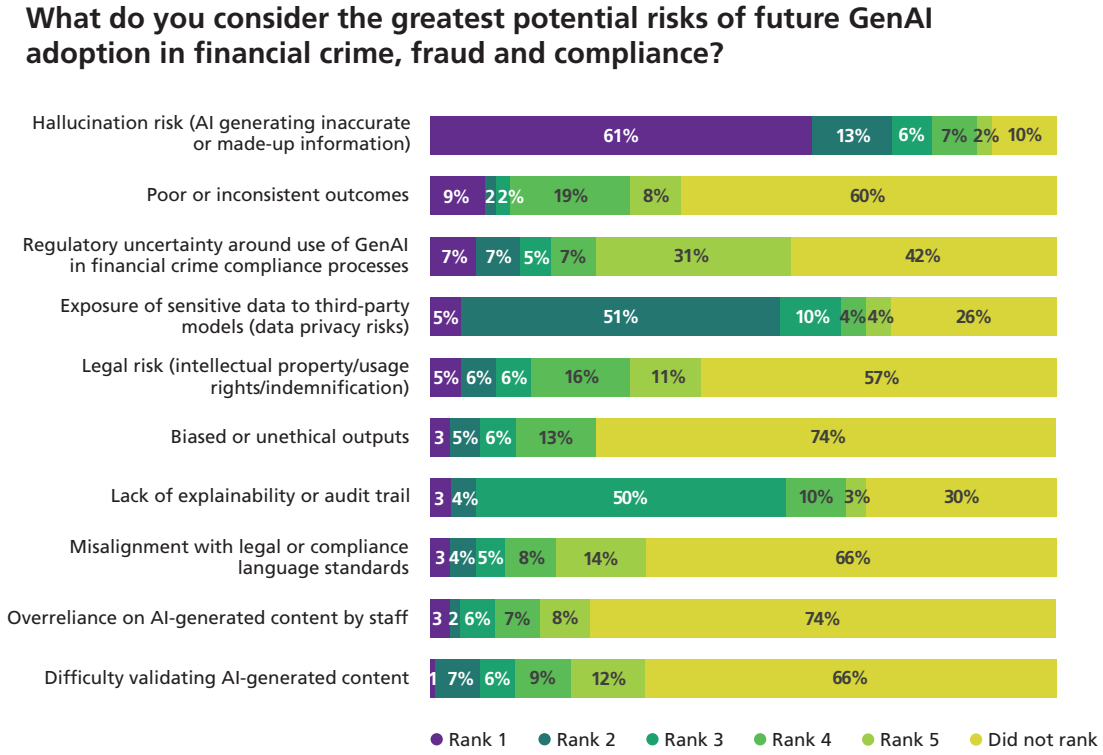

Source: Chartis Research

Hallucination risk is the most dominant concern, with 61% of respondents citing this as their number one risk. Close behind, with 51% of respondents ranking it second, is the exposure of sensitive data to third-party models. This highlights the tension between leveraging GenAI systems and maintaining compliance with stringent data privacy obligations, especially under regimes like the General Data Protection Regulation (GDPR) or banking secrecy laws. Notably, this runs counter to overall AI business challenges, where data privacy was less of a factor, indicating that this is a relatively unique challenge for GenAI.

Regulatory uncertainty (which 57% of respondents selected as one of their top five concerns) is also a factor, indicating persistent ambiguity around how supervisors will treat GenAI-enabled decision-making and documentation. Institutions appear open to GenAI's potential but see governance, control and auditability as issues.

Agentic AI's potential

Overview

Agentic AI refers to systems designed to act with a degree of autonomy toward defined goals. Unlike traditional or GenAI, which respond to prompts or static inputs, agentic AI can potentially plan, execute and adapt sequences of actions based on feedback and changing context.

The survey explores how institutions view the potential of agentic AI to transform analytical and investigative workflows, and where its use could deliver the greatest efficiency gains. It also considers the broader implications of this shift.

Although relatively new, agentic AI has the capacity for transformation

Views on agentic AI are optimistic (see Figure 26): 45% of respondents believe it will enhance current processes without fully transforming compliance, while 33% expect it to drive a major shift in how work is done. For a relatively new and untested technology, this seems like a strong endorsement.

A smaller group of respondents (15%) remain skeptical, seeing agentic AI as overhyped, and 7% are undecided.

Figure 26: Will agentic AI transform the anti-financial crime and compliance function?

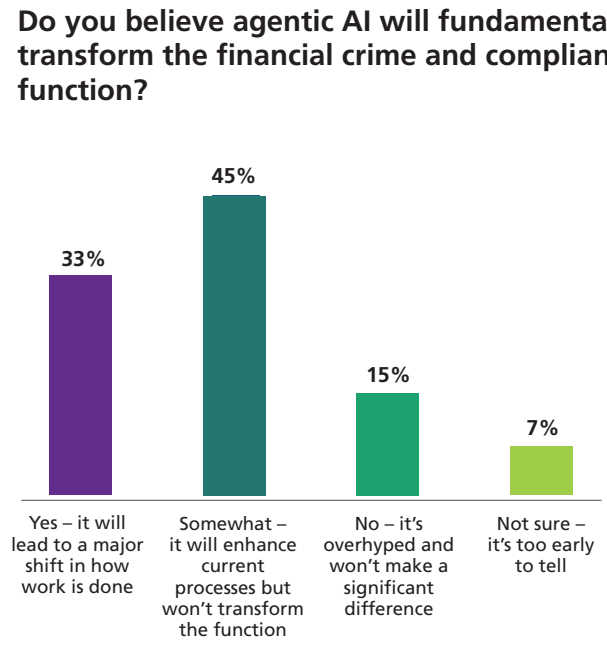

Source: Chartis Research

Agentic AI will have the most impact on investigations, reporting and research

As Figure 27 illustrates, respondents expect agentic AI to have the most transformative impact on case investigations (61% rank this first) and narrative/suspicious activity report (SAR) drafting (55% rank this second). These areas typically require sustained reasoning and/or contextual understanding, as does background research and context gathering, which 58% of respondents ranked in their top three. These findings align with agentic AI's core strengths: dynamic task orchestration and multi-step logic.

Figure 27: Compliance activities that will be transformed by agentic AI

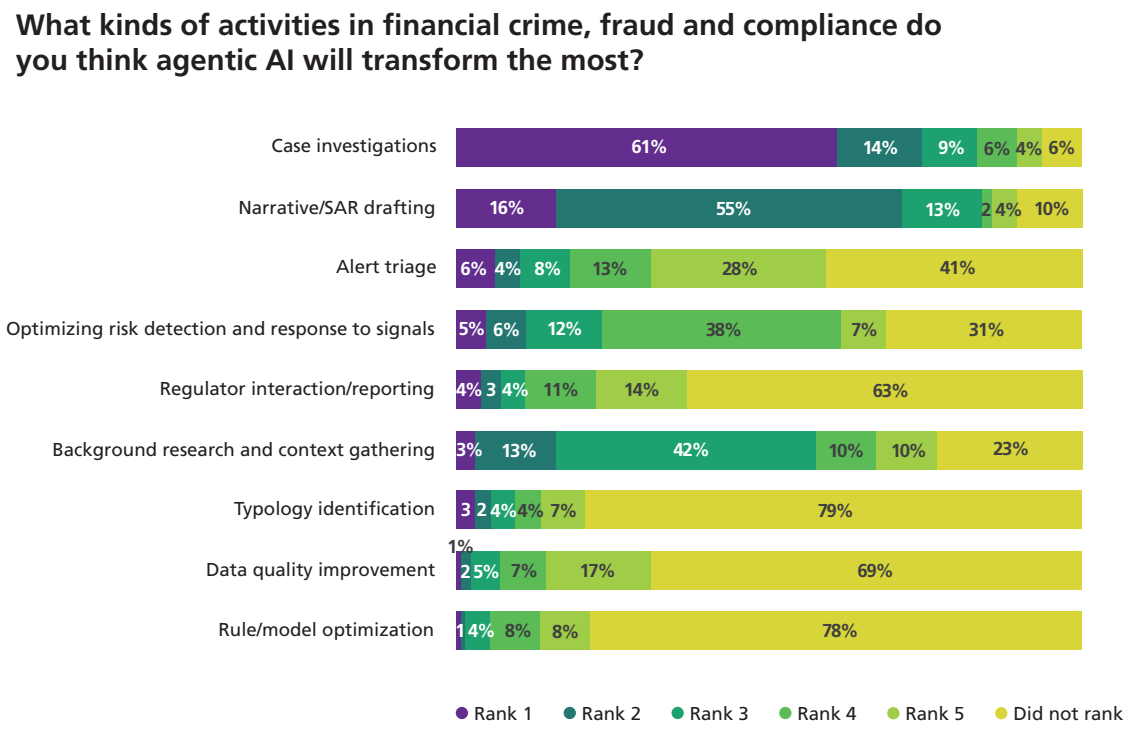

Source: Chartis Research

In contrast, respondents viewed typology identification and rule/model optimization as less likely to be transformed. This is probably due to their reliance on human domain expertise and the need for regulatory transparency – or possibly because solutions in this space are, as of now, relatively undeveloped.

Overall, the results suggest that firms expect agentic AI to enhance efficiency and automation, rather than focus on detection logic. This reinforces the narrative that next-generation AI will serve as a co-pilot, augmenting analyst workflows rather than replacing human judgment.

Firms expect agentic AI to lead to workforce optimization

The majority view is that agentic AI will lead to moderate headcount reductions, with some roles eliminated but core staffing maintained (see Figure 28). A further 10% of respondents expect significant reductions, while 33% foresee no major change, suggesting that agentic AI will cause a shift in how work is performed without reducing overall team size.

Figure 28: Will agentic AI have an impact on compliance headcount?

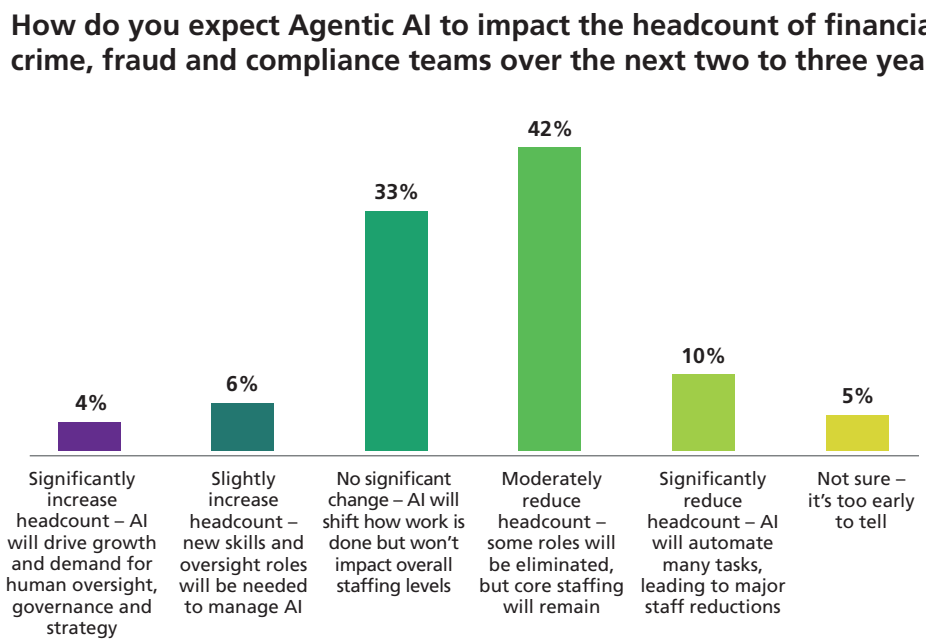

Source: Chartis Research

Only a small minority of respondents anticipate headcount growth, with 4% expecting it to be significant and 6% believing it will be slight, reflecting the limited belief that AI will drive new oversight and governance roles at scale.

Overall, expectations lean toward efficiency-driven workforce optimization rather than expansion. This highlights the importance of reskilling and redeployment strategies that can balance the automation of repetitive tasks, while ensuring that human expertise remains part of the process.

'The goal with these tools isn't to cut jobs, it's to repurpose people. We want to redeploy talent to higher-value work, not drive reductions through automation.'

Chief Compliance Officer, Tier 2 Financial Institution

Agentic risks – regulation and auditability

The most pressing perceived risk of agentic AI adoption (see Figure 29 on page 25) is a concern about meeting regulatory and auditability requirements (60% of respondents rank this their greatest perceived risk), reflecting concerns over governance and accountability.

While teams do not expect major headcount changes, they do fear job displacement and a loss of human expertise. Job displacement and over-reliance on automation follows, with 51% ranking this as their second greatest concern, highlighting organizational unease about losing human expertise or oversight as AI becomes more independent. Increased operational complexity and cost ranked third for 49% of respondents, indicating that many are concerned about the potential price point and added complications of agentic AI.

Figure 29: Biggest risks of adopting agentic AI

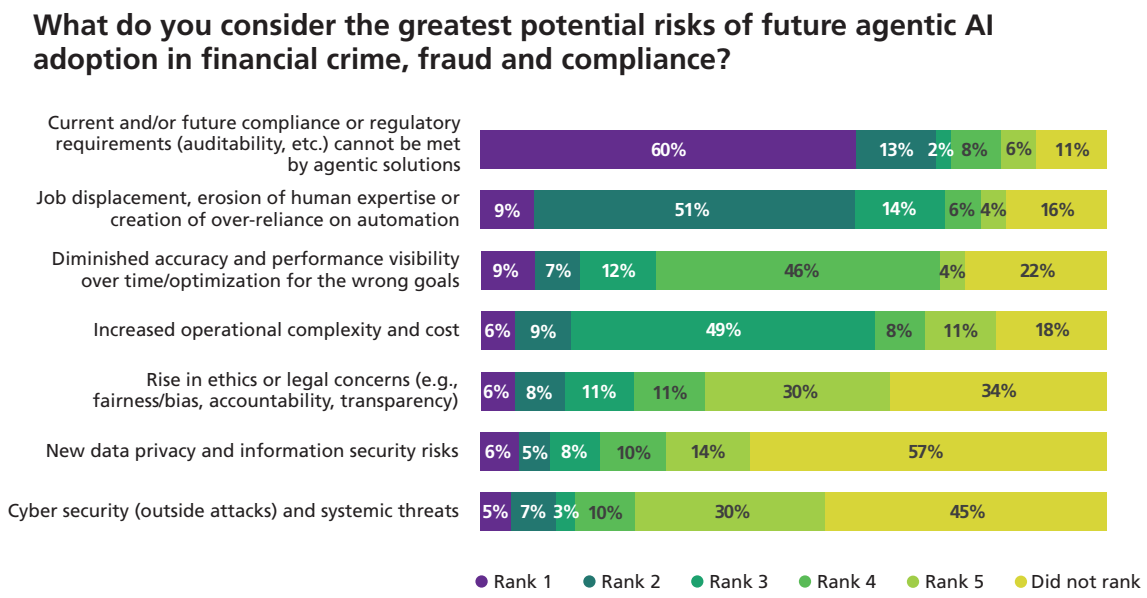

Source: Chartis Research

Other concerns included diminished accuracy (the fear that agentic AI could 'drift' from its intended goals) and performance visibility, with 46% of respondents ranking this fourth. Few respondents selected data privacy or information security risks as one of their top five concerns, with 57% not ranking these issues at all. This indicates some confidence that existing cybersecurity frameworks are robust enough for agentic AI. This is ironic, given that agentic systems are often powered by GenAI, and firms are concerned about GenAI's data privacy. This may indicate an opportunity to educate the industry about agentic technologies and how they might work in a financial crime and compliance context.

Overall, respondents view agentic AI as promising efficiency but raising questions about control, accountability and regulatory compliance. Unlike GenAI, whose risks skew toward output quality and explainability, respondents see agentic AI as a challenge to govern and oversee.

Conclusion

AI adoption in anti-financial crime, fraud prevention and compliance is entering a more mature phase, moving from experimentation toward more targeted, outcome-driven integration. Most firms are now testing or piloting AI with a clear path toward scaling, with adoption strongest in fraud prevention and increasingly extending into AML transaction monitoring, sanctions screening and case management.

While traditional AI techniques such as ML and NLP continue to anchor most use cases, generative and agentic AI are emerging rapidly as new engines of investigative efficiency, decision support and model optimization. These next-generation tools are already approaching adoption levels comparable to established technologies, reflecting growing confidence in their ability to deliver real-world impact across financial crime and compliance workflows.

The business case for AI is strengthening, with firms reporting measurable gains in efficiency, scalability and cost savings. However, regulatory uncertainty, data privacy concerns and limited internal expertise remain the main constraints to wider deployment. On the technical side, challenges around integration, explainability and governance persist but are being steadily reduced as operational experience deepens and governance frameworks evolve.

Encouragingly, most respondents expect increased regulatory support and higher AI investment over the next two to three years. This signals a sector moving toward responsible expansion in their use of AI, where the emphasis will shift from adoption to optimization and assurance, embedding AI as a critical enabler more deeply into financial crime risk management.

Next steps for financial institutions

Banks should move from experimentation to more structured deployment of AI.

Key priorities include:

- **Taking a holistic approach to AI**. Aligning strategy and technology across business, risk and compliance functions to ensure that AI initiatives are consistent, scalable and grounded in governance and compliance frameworks.
- **Strengthening data and model governance**. Standardizing data quality, lineage and validation frameworks to support explainable and auditable AI outputs.
- **Targeting proven use cases**. Prioritizing applications in alert triage, fraud detection and case management, where value and control are measurable.
- **Embedding human oversight**. Maintaining 'human-in-the-loop' review and clear accountability for AI-driven decisions.
- **Investing in capability and coordination**. Building internal AI literacy, formalizing cross-functional governance and engaging with regulators and industry peers.
- **Preparing for agentic AI**. Beginning the controlled testing of autonomous investigative and drafting tools within established compliance frameworks.

Whatever steps are taken, it is clear that the industry is undergoing a journey, shifting from the question of whether it should use AI to how.

Demographics

The survey has a relatively even distribution of firm sizes (see Figure 30).

Figure 30: Respondent organizations' assets under management

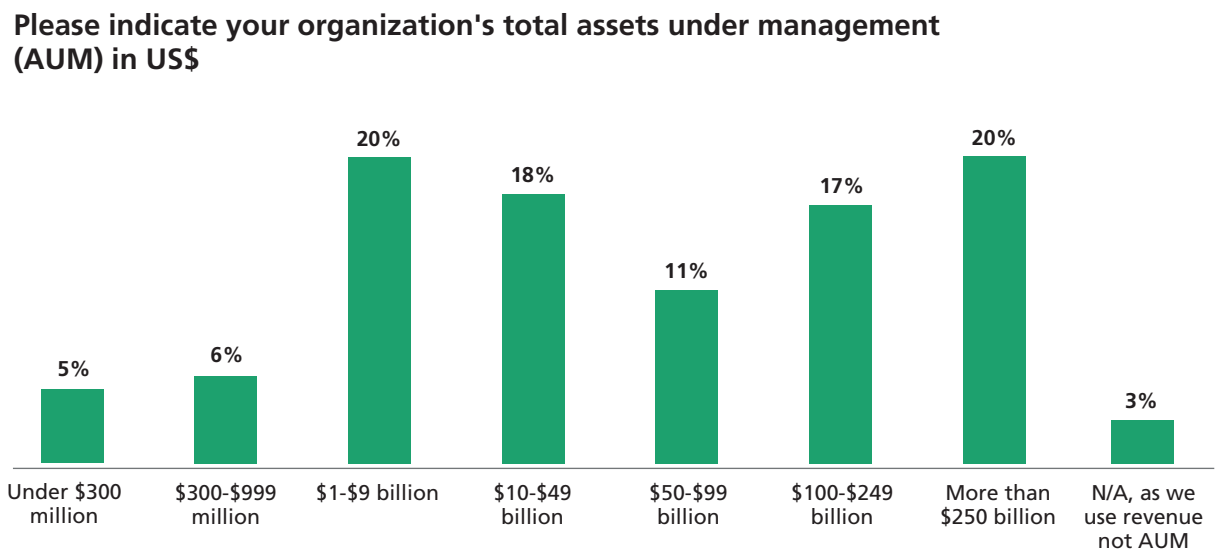

Source: Chartis Research

Most respondents are in risk and technology/IT roles (see Figure 31), at the analyst or director level (see Figure 32).

Figure 31: Respondents' job function

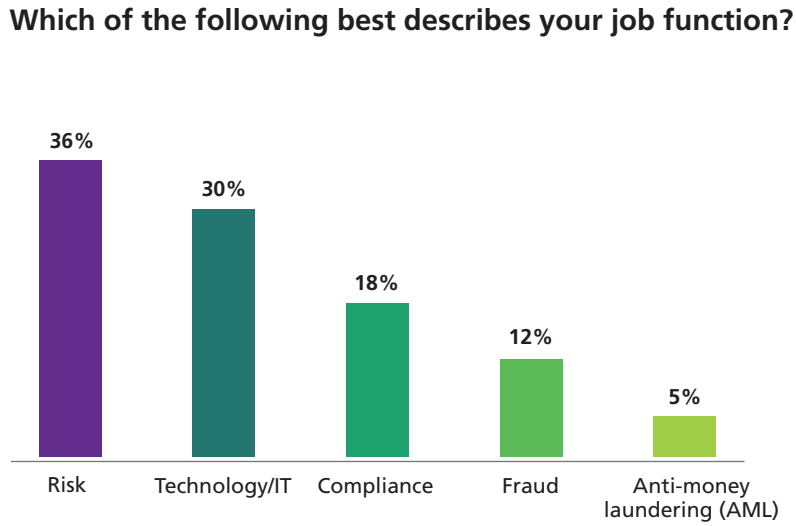

Source: Chartis Research

Figure 32: Respondents' job level

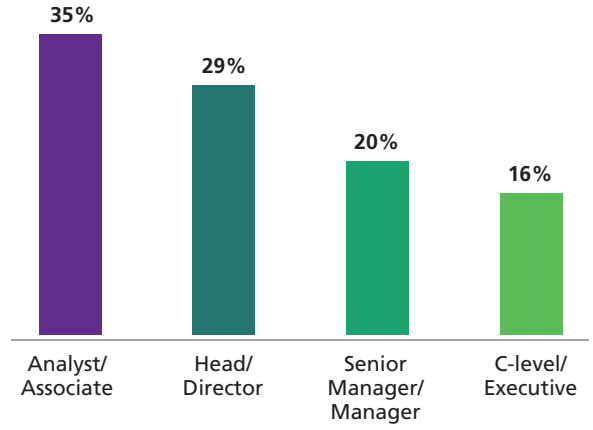

Source: Chartis Research

Respondents are distributed across the globe (see Figure 33), with most in the US and Europe.

Figure 33: Location of respondent organizations' HQ

Source: Chartis Research

Chartis

HAWK