cloud security alliance®

anjuna®

Survey Report

The State of Cloud and AI for Financial Services

2026 Industry Survey Report

Condensed editorial rewrite for FSI cybersecurity professionals

Based solely on the uploaded CSAI Foundation / Cloud Security Alliance AI Safety Initiative report.
All data points and in-text citation numbers are preserved from the source document.

The industry has moved from asking whether to adopt AI to deciding how to govern it before autonomy outruns control.

Acknowledgments

Contributors

Troy Leach
A. Orbert
Mark Bower - *Anjuna collaborator*
Keith Hennett
Hillary Baron
Sofia Pogrebynska
Adam Parker, CISO
Ruth Jo. Scheier
James Ross
Matt Browne
Anant Wairagade
Jimmy Barber
Aman Sardana
Fred Budd
Julia Cherashore
Jeff Stapleton
Shiva Gosula

CSA Staff

Graphic Design

Stephen Lumpe
Stephen Smith

About the Sponsor

Anjuna Security gives enterprises the freedom to govern and scale AI agents across different environments with confidence. Built on confidential computing technology, Anjuna's platform provides the central control to ensure AI agents act only as intended, stay within budget, and continuously comply with regulatory frameworks. Every action is evaluated, and every operational risk is managed; credentials never leak, and sensitive data is always controlled and secured. With Anjuna, regulated enterprises can move fast with trusted autonomous AI, efficiently meeting goals without increasing risk and losing security oversight. For more information, visit anjuna.io.

Table of Contents

Executive Summary

The third iteration of Cloud Security Alliance (CSA)'s survey on the adoption of cloud and now AI, has shown that cloud services are effectively universal and several factors of GenAI have moved from experimentation into production at scale. The strategic question is no longer whether institutions will adopt cloud services, Large Language Models (LLM) and agents. It is now whether security, governance, and operating models are evolving fast enough to keep pace to manage the potential risk.

The 2026 survey findings show a sector moving toward agent-mediated financial operations. Sixty-two percent of organizations report deploying AI agents. Among those using agents, 38 percent have granted conditional or high autonomy.

The agentic autonomy is expected to extend to the consumer as well. 85% believe AI agents will initiate and execute financial transactions on behalf of users, and 65 percent believe that shift will require a new authorization model.

Many comments from respondents resonated with a concern that adoption is moving faster than control maturity. 20% of respondents reported known AI-related security incidents, while another 21 percent are unsure whether such incidents have occurred, raising the potential likelihood of prior incidents. The data suggests that deployment without visibility may lead to future struggles to investigate AI-related incidents, properly report and learn from them.

While cloud services are becoming omnipresent with 98.3% use among financial service organizations, traditional cloud security risks have not gone away, only more complicated as AI services intertwine. Third-party and supply-chain risk leads cloud concerns at 55%, followed closely by misconfiguration at 52%. Non-human identities were also recognized as a growing concern for misconfiguration of credentials and absence of monitoring, and policy enforcement.

The consistent theme for both cloud and AI is that the primary risk is a data problem. For example, sensitive data leakage was the top AI security concern at 61 percent, far exceeding concerns about model attacks, prompt injection, or adversarial techniques. In practical terms, the concern is not just that AI systems could be attacked. It is that sensitive financial data can be exposed through ordinary usage patterns such as prompts, chat history, training data, and retrieval-augmented generation connectors. In other words, employees not aware of the risk of their actions or lack of guardrails to prevent such use.

Security priorities are responding to that reality. *Improving identity and access security* (48%), and *integrating AI security into cloud environments* (44%) lead future planned security investments. Yet progress is constrained by the same operational friction seen across the sector: budget limitations and lack of AI expertise (both at 45%), and data-classification gaps at 26 percent.

Financial services has entered an era in which AI systems will participate directly in financial decision-making and execution, powered by the scalability of cloud services already entrenched. The institutions best positioned for that shift will be those that invest now in AI-specific monitoring, non-human and agent identity governance, data controls at the point of AI interaction, and authorization models designed for autonomous systems.

1. AI is in production or expected soon

Over one-third are actively implementing AI in production, with an additional 9 percent at advanced adoption. Only 2% report no AI usage.

43% in active implementation or advanced adoption; Additional 49% in exploration or pilots

2. AI Agents are mainstream

AI agents are gaining autonomy. Customer service, cybersecurity operations, fraud detection, and back-office operations lead examples.

62% deploying AI agents; **93%** of those using agents grant some form of autonomy; **38%** with conditional or high autonomy

3. AI risk is primarily a data problem

Sensitive data leakage through AI interactions is the dominant AI security concern, far exceeding worries about model attacks or adversarial techniques.

61% cite data leakage as top AI risk

4. Cloud services remain universal

Just as in the prior 2023 report, there has been an increase in cloud utilization as majority of software is designed as cloud-native and ability to scale new features more easily.

98.3% operate cloud services; **33%** primarily or fully cloud-based architecture.

5. Strong Senior Leadership Support for Cloud Security

Great understanding of business risk and ability to enable AI-deployment has improved overall executive support.

91% report 'strong' or 'moderate' support

6. Misconfiguration and Third-party Risk remain top cloud risks

Similar concern of human error completed the Top 3 security concerns, showing greater confidence in security tools and supporting technology.

3rd-Party Risk (**55%** and Misconfigurations (**52%**) remain as a Top 3 Security Issues. Ransomware (**9%**) seen as much lower concern compared to 2023 results

7. The industry expects autonomous financial transactions

An overwhelming majority of respondents believe AI agents will initiate and execute payments on behalf of consumers.

85% anticipate agentic payments; **65%** believe a new authorization framework is required

Key Findings

Cloud Adoption: A Universal Baseline with New Complexities

Cloud adoption has reached practical saturation in financial services. Only 1.7 percent of respondents describe their organizations as entirely on-premises. The remaining 98.3 percent spans primarily on-premises with some cloud at 20 percent, hybrid environments at 46 percent, primarily cloud with some on-premises at 14 percent, and fully cloud-based architectures at 19 percent. Hybrid remains the dominant model because legacy core systems still coexist with cloud-native services.

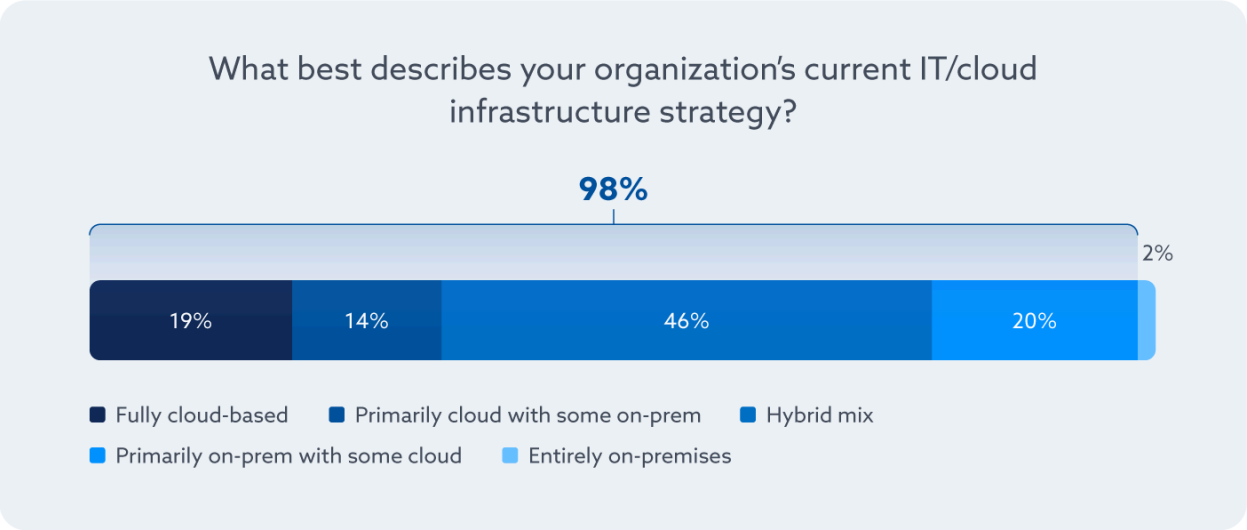

The trend line is clear. In 2020, 91 percent of respondents reported using some form of cloud computing. By 2023, that number had reached 98 percent. The 2026 results confirm that the debate has shifted from adoption to architecture [1][2]. For most institutions, the real work now lies in managing resilience, data locality, integration, and control across mixed environments rather than deciding whether to use cloud at all.

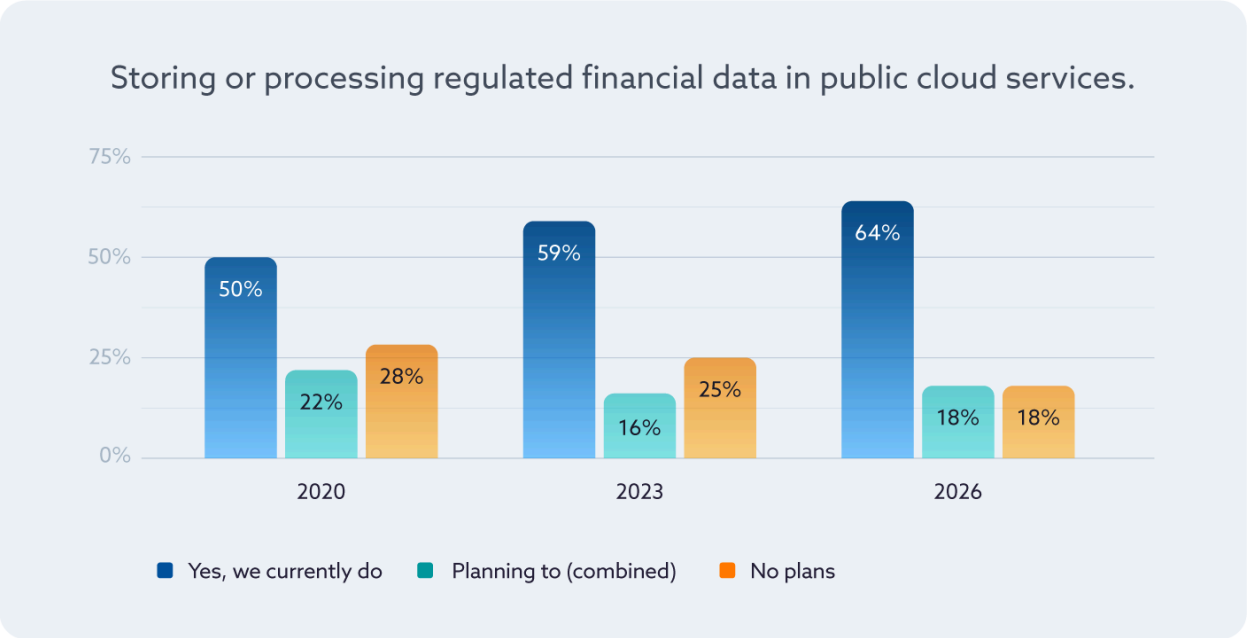

Use of public cloud for regulated financial data continues to rise. Sixty-four percent currently store or process regulated financial data in public cloud services, another 18 percent plan to do so within 18 months, and only 18 percent report no plans to move regulated data into public cloud, down from 25 percent in 2023. Confidence in running business-critical workloads in cloud has also continued to strengthen. The report's message is clear: cloud has become infrastructure, and the governance challenge now sits in how it is used, not whether it is used.

Multi-Cloud Strategy: Active Diversification and Exit Planning

Multi-cloud strategy is actively changing. Nearly half of respondents, 48 percent, report changing their cloud service provider strategy in the prior 12 months. Among those organizations, 44 percent migrated workloads between providers, 37 percent introduced exit or contingency plans, 31 percent consolidated providers, 29 percent added a new CSP, and 25 percent reduced reliance on a single provider.

The primary drivers of CSP strategy changes reflect a more disciplined approach to cloud concentration and the priorities shaping financial services. Improved resiliency and availability was the leading driver at 51 percent, narrowly ahead of cost optimization at 50 percent. These are followed by reducing the complexity of managed services (39 percent), regulatory or data sovereignty requirements (38 percent), technical requirements (33 percent), and changes in business needs (32 percent). The prominence of resiliency as a driver reflects the supply-chain concerns, with recent incidents where a single software update from a third-party vendor caused cascading failures across financial institutions worldwide.

Despite the shift toward diversification, 5% of respondents have moved services back to on-premise infrastructure. For most, the direction of travel remains firmly toward cloud. The most common questions are how to distribute workloads across providers to optimize for resilience, cost, and regulatory compliance. The increase of exit and contingency plans by 37 percent of respondents represents a significant maturation from the 2023 survey, where exit planning was identified as an area of weakness. This shift likely reflects, at least in part, DORA's explicit requirements for ICT third-party exit strategies, which became enforceable in January 2025 [4]

The Governance Framework Landscape: Fragmented but Evolving

Financial institutions operate within a dense web of overlapping governance and cybersecurity frameworks, and the 2026 survey captures a snapshot of which frameworks are actually being used to manage AI and cloud security. The results reveal both the enduring dominance of established standards and the early-stage adoption of AI-specific governance tools.

ISO/IEC 27001 leads adoption at 62 percent, confirming its position as the de facto baseline for information security management in global financial services. The NIST Cybersecurity Framework follows at 59 percent, SOC 2 attestations (42 percent). The CSA Cloud Controls Matrix maintains strong adoption at 39 percent, serving as a cloud-specific complement to broader frameworks. Note that for CSA's STAR Level 2, independent assessments for CCM, they must be associated with either an ISO/IEC 27001 or SOC 2 assessment.

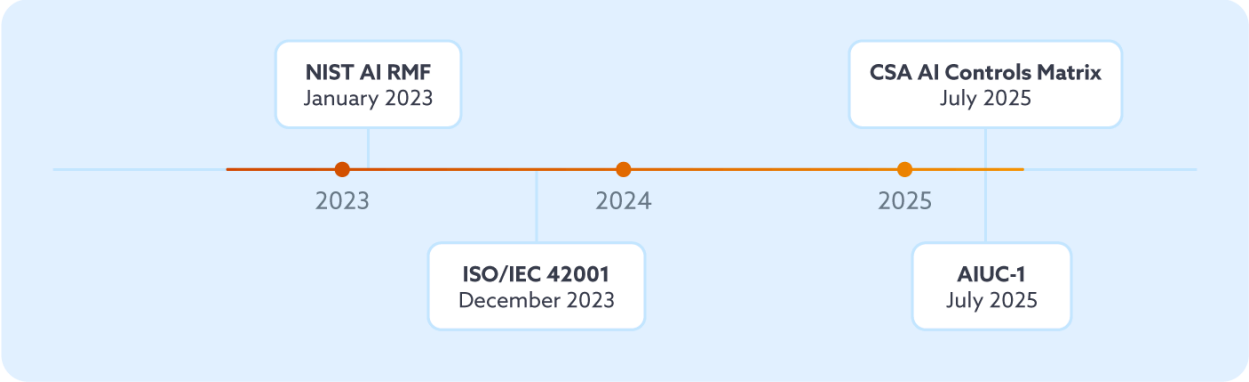

The AI-specific framework landscape is newer and more fragmented. The NIST AI Risk Management Framework (AI RMF) has achieved the highest adoption among AI governance tools at 39 percent, benefiting from both the reputation of NIST and its early release in January 2023. The CSA AI Controls Matrix (AICM), published in July 2025, has been adopted by 24 percent of respondents. ISO/IEC 42001, the first international standard for AI management systems, published in December 2023, shows 21 percent adoption. The AI Unified Compliance framework (AIUC-1), published in July 2025, focused on AI agents had been adopted by 7 percent of respondents.

The "Other" category in the framework question provides a revealing window into the regulatory frameworks that financial institutions must also navigate. Respondents cited DORA, the EU AI Act, GDPR, PCI DSS, NIST 800-53, SOX, COBIT, and various national regulations, underscoring the compliance burden that this highly-regulated industry bears. Several respondents noted that they rely on internal control standards derived from multiple frameworks, suggesting that the industry norm is not necessarily focused-adherence to any single framework but rather a composite approach that maps controls across multiple overlapping requirements.

This fragmentation creates both risk and opportunity. The risk is that organizations invest significant effort in compliance mapping without achieving coherent security outcomes. The opportunity is for integrative frameworks to serve as unifying layers that reduce duplication and align cloud and AI security governance within coherent control structures.

This also is the catalyst for CSA's Compliance Automation Revolution (CAR) to support industry interest in creating more machine-language capabilities to accommodate continuous monitoring, real-time security posture and more easily demonstrated compliance requirements.

Security Posture: Tools, Risks, and Priorities

The security tooling landscape is maturing, but coverage is uneven. SIEM remains the most widely deployed tool at 76 percent. This could reflect regulatory requirements for central log aggregation, audit trails and alerting as well as the customization, such as tailored automated forensics as noted by one respondent.

CSPM follows at 56 percent, XDR for cloud at 40 percent, CNAPP at 39 percent, and CWPP at 28 percent. Respondent write-ins also point to the next wave of tooling, including AI Security Posture Management, DSPM, CTEM, and shadow-AI detection.

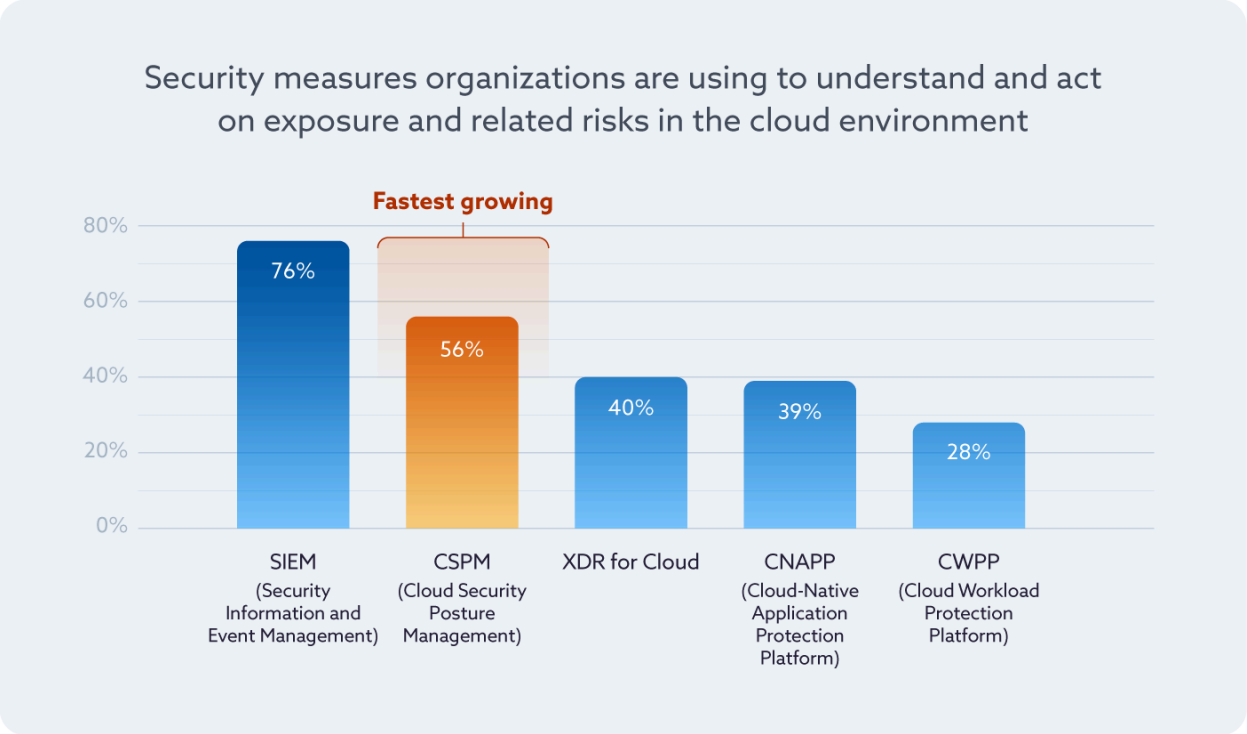

CSPM adoption shows significant adoption in the past three years up from an estimated 25-29%, representing the fastest-growing tool category.

Top Risks

When asked to identify the three greatest security risks to their cloud infrastructure, respondents produced a risk ordering that reflects the current threat environment. Third-party risks and supply chain attacks dominate at 55 percent. Cloud misconfigurations remain the second greatest concern at 52 percent, a persistent challenge that reflects both the complexity of cloud environments and the difficulty of maintaining consistent security configurations across multi-cloud architectures.

Human error follows at 27 percent, insecure non-human identities at 24 percent, insecure APIs at 24 percent, and insecure human identities at 19 percent. One notable shift from prior surveys is the dramatic decline of ransomware as a perceived top-three risk, falling to just 9 percent of respondents — down from a far more prominent position in 2023 when it featured alongside data exfiltration and system sabotage as a leading concern.

This decline warrants cautious interpretation. It may reflect genuine improvements in organizational resilience: better backup and recovery capabilities, improved detection, and ransomware-specific incident response playbooks have reduced expected impact. However, ransomware attacks on financial services continue at high frequency — 65 percent of financial services IT professionals reported being hit in 2024 [21] — and the ICBC LockBit attack demonstrated that consequences remain severe [22]. The more likely

explanation is that ransomware has become a known and managed risk in a way that AI-driven threats and supply chain vulnerabilities have not, which is why those newer categories now dominate the risk register.

The positioning of non-human identity risk at 24 percent represents an important signal. As financial institutions deploy more automated processes, microservices, and now AI agents, the population of non-human identities (service accounts, API keys, machine certificates, and agent credentials) has grown to vastly outnumber human identities. Vendor research suggests that non-human identities significantly outnumber human identities in enterprise environments, with one estimate placing the ratio at approximately 144 to 1 across industries and 96 to 1 in financial services [25]. While precise figures vary by source and methodology, the directional trend, rapid growth in machine identities driven by microservices, automation, and now AI agents is broadly supported.

Eighty-two percent of organizations experienced at least one identity-driven cyberattack in the past year, with the rate even higher in financial services [26]. The survey data suggests that financial services are beginning to recognize non-human identity management as a distinct security challenge, but the scale of the problem demands more urgent attention than the 24 percent figure might suggest.

Security Priorities and Implementation Barriers

Respondents are prioritizing the right areas, even if execution remains uneven. Improving identity and access security leads next-year priorities at 48 percent, followed by integrating AI security into cloud environments at 44 percent. Other priorities include detecting and remediating misconfigurations at 33 percent, improving resiliency and disaster recovery at 30 percent, adapting or improving Zero Trust capabilities at 27 percent, budget optimization at 26 percent, maintaining compliance at 22 percent, managing workload vulnerabilities at 15 percent, implementing cloud detection and response at 15 percent, evaluating attack surface at 14 percent, and preparing for quantum threats at 10 percent [7].

The barriers are less about technology availability than organizational capacity. Budget limitations and competing priorities lead at 45 percent. Insufficient staffing or skills in cloud security, IAM, and DevSecOps affect 28 percent. Data-classification and policy-maturity gaps affect 26 percent. Multi-cloud inconsistency stands at 24 percent, legacy integration at 22 percent, single-provider dependence or limited portability at 20 percent, and inadequate visibility at 20 percent.

The skills gap has also moved up the stack. In 2020, the industry was still short on foundational cloud-security capabilities. By 2026, lack of AI and ML expertise at 45 percent has become the more acute constraint. Data classification is especially important because organizations that cannot reliably label and govern their data cannot confidently control what AI systems can access, train on, or reveal. Senior leadership support is generally positive, with 43 percent describing it as strong and 38 percent as moderate. While 16 percent still report limited or no meaningful support the results show a positive directional shift from the prior survey and that executive awareness of cloud security as a business imperative and part of expected infrastructure has improved.

81% of respondents report **strong or moderate** senior leadership support for cloud security

Implementation Barriers: Budget, Skills, and Policy Maturity

The greatest barriers to implementing new cloud security capabilities paint a picture of an industry constrained less by technology availability than by organizational capacity. Budget limitations and competing priorities lead at 45 percent, a persistent challenge that reflects the tension between security investment and revenue-generating initiatives. Insufficient staffing or skills in cloud security, IAM, and DevSecOps follow at 28 percent. Data classification and policy maturity gaps affect 26 percent, multicloud complexity and inconsistent controls across providers 24 percent, difficulty integrating with legacy systems 22 percent, dependence on a single CSP or limited portability 20 percent, and inadequate visibility across cloud environments 20 percent.

Greatest barriers to implementing new cloud security capabilities.

45% Budget limitations / competing priorities

28% Insufficient staffing or skills

26% Data classification / policy maturity gaps

24% Multi-cloud complexity / inconsistent control implementation across providers

22% Difficulty integrating with legacy systems or on-prem dependencies

20% Dependence on a single CSP / limited portability

20% Inadequate visibility / telemetry across cloud environments

The skills gap has also migrated up the technology stack since the prior surveys. In 2020, the shortage was primarily in foundational cloud security competencies — configuring cloud-native controls, understanding shared responsibility models, managing cloud IAM. By 2026, the AI/ML expertise gap (45 percent) surpassed the cloud security staffing gap (28 percent) as the more acute constraint, and the talent market for AI security professionals resembles the cloud security market nearly a decade prior: intense competition for practitioners who combine machine learning expertise with security architecture knowledge and financial services domain understanding.

The data classification gap at 26 percent deserves particular attention in the context of AI adoption.

Organizations that cannot reliably classify and label their data face compounded challenges when deploying AI systems that process, learn from, and generate outputs based on that data. Without clear data classification policies, it becomes impossible to enforce appropriate controls on what data AI systems can access, what they can include in training, and what they can expose in outputs. This gap is a root cause of the data leakage concerns that dominate the AI risk landscape.

Senior leadership support for cloud security is generally strong, with 43 percent of respondents describing it as strong support and 38 percent as moderate. However, 16 percent report limited or no meaningful support, a concerning figure given the scale of investment required to address the security challenges.

AI Adoption: From Experimentation to Production at Scale

The 2026 survey introduced new AI adoption questions and the results show an industry that has moved well beyond AI experimentation. When asked about the adoption of third-party and ML within business process, approximately 91% reported some use or exploration. Fourteen percent are in early exploration, 35 percent in piloting, 34 percent in active production implementation, and 9 percent in advanced adoption. The pattern for in-house AI development is nearly identical: 17 percent early exploration, 33 percent piloting, 35 percent active implementation, and 9 percent advanced adoption.

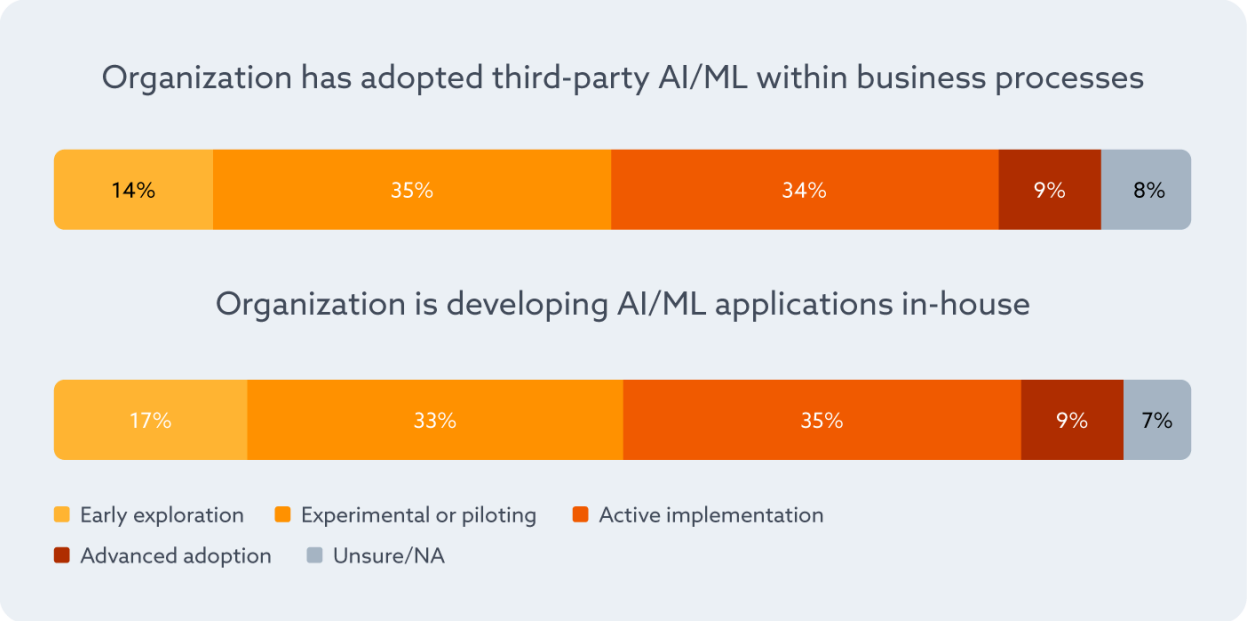

That symmetry suggests that institutions are pursuing a dual strategy: buying AI capabilities from vendors while also building proprietary applications aligned to their own regulatory and operational needs.

These figures represent a transformative shift for an industry that was tentatively discussing AI use cases as recently as 2023. Several factors have accelerated adoption since the prior survey: the maturation of LLM capabilities, the availability of enterprise-grade AI platforms , and growing evidence that AI can deliver measurable improvements in fraud detection, customer service efficiency, and operations.

Consulting firms have also published optimistic projections for AI in banking. For example, Accenture estimates 22 to 30 percent productivity gains for early adopters [8], while McKinsey projects $200 to $340 billion in annual potential [9] – though these figures represent vendor forecasts rather than observed outcomes and should be interpreted with the commercial interests of their authors in mind.

The institution examples underscore the acceleration. JPMorgan Chase allocated $19.8 billion to technology in 2026, including $1.2 billion for AI, and reports more than 200,000 employees using its LLM Suite [10]. Bank of America reports more than 210,000 associates using its generative AI assistant, 90 percent workforce AI adoption, and more than a 50 percent reduction in IT service desk queries [11]. Goldman Sachs deployed autonomous AI software engineers across its 12,000 developers in July 2025 [12]. Citigroup deployed proprietary generative AI tools to more than 150,000 employees across 80 countries [13].

AI Agents: The Autonomy Frontier

The most consequential finding in the survey may be the breadth of AI agent deployment. Sixty-two percent of respondents report that their organization is using AI agents. Because the survey defined agents broadly, that figure includes both more capable autonomous systems and simpler task-oriented bots. Even so, the adoption signal is strong. Only 27 percent report no agent use, and 11 percent are unsure, which may itself signal shadow deployment outside formal governance.

Also worth mentioning is that the visibility into SaaS vendors' use of AI features can bypass TPRM reviews and detection. In March 2026, Obsidian Security observed 69,749 interactions between users and corporate data with SaaS embedded AI features, most of which would have gone unnoticed without in-browser detection. Respondents may not have considered those third-party use cases when completing the survey.

Among organizations deploying AI agents, the use cases span the full breadth of financial services operations. Customer service and support, including AI chatbots for customer inquiries, leads at 63 percent. Cybersecurity and IT operations, encompassing AI-driven threat detection and system monitoring, follows at 47 percent. Internal employee assistive tools for coding, writing, and analytics are

deployed by 46 percent. Back-office automation handles routine administrative and IT tasks at 44 percent. Fraud detection and AML compliance monitoring runs at 41 percent. Personalized advice and marketing through roboadvisors and AI-driven customer insights accounts for 38 percent. Credit underwriting and risk analysis, an area with significant regulatory implications for model explainability, is at 22 percent. Trading and investment management, including algorithmic trading, reaches 20 percent. Regulatory compliance automation accounts for 20 percent. More advanced regulatory rules, such as the classifications within the EU AI Act and similar global standards may influence the adoption for the latter categories as maturity develops around demonstrating controls for bias, ethics and other requirements.

The level of autonomy granted to AI agents reveals a spectrum of trust that financial institutions are navigating in real time. Only 7 percent maintain a no-autonomy posture in which humans make all decisions. The majority, at 55 percent, operate under limited autonomy where AI performs tasks but with active human oversight — the classic "human in the loop" model. Conditional autonomy, where AI acts independently in low-risk scenarios under defined guardrails, applies to 33 percent. And 5 percent have granted high autonomy, allowing AI to operate independently for critical actions.

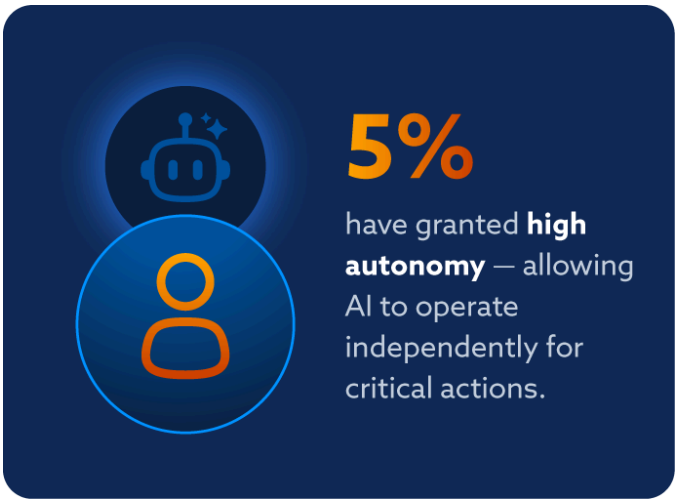

That final number is small in absolute terms, but strategically important compared to where respondents were just three years ago. It suggests that some institutions have already crossed from AI assistance to AI action.

Agentic Payments: A New Authorization Paradigm

The coming shift for agentic AI use is especially visible in payments. Asked whether consumers will use AI agents to initiate and execute payment transactions without direct real-time participation, only 5 percent said no. 85 percent of respondents anticipate agent-driven payments. The majority (65%) however, believe this type of payment will require an entirely new model for authorization.

The significance is not just that AI agents may transact. It is that existing payment instruments and authentication models were designed around a human being present to confirm details of the transaction. They were not designed for a delegated software agent that can negotiate, select, and execute purchases on a consumer's behalf. The survey response indicates a belief the industry will need a new trust anchor before agentic commerce can possibly scale. Trusted Execution Environments in confidential computing

infrastructure presents a potential approach to remove reliance on traditional software-only trust, delivering a hardware-backed measurable trust for security critical agents or reliant services..

However, that transition is already underway. Visa launched Intelligent Commerce in October 2025 with the Trusted Agent Protocol (TAP) and by December 2025 had completed hundreds of secure agent-initiated transactions with more than 100 partners [14]. Mastercard launched Agent Pay in April 2025, rolled it out to all U.S. cardholders by the 2025 holiday season, and partnered with Santander on Europe's first live end-to-end AI-agent payment in early 2026 [15]. Stripe launched a Machine Payments Protocol in March 2026 [16], and Google announced Agent Payments Protocol in September 2025 with 60 collaborating organizations [17]. Edgar, Dunn & Co. projected the agentic commerce market at $135 billion in 2025 and $1.7 trillion by 2030 [18]. Yet, in a *Future of Money* survey by Accenture, 87 percent of payments leaders and CTOs believe trust will be the biggest barrier, and 78 percent expect fraud to increase materially [19].

AI Security Risks

Concerns over data loss or access to sensitive data seemed to dominate respondents' concerns when asked to identify the three greatest security risks when using AI. The primary themes from feedback seemed to center around confidentiality, authorization, and auditability for how that data was being protected.

The top response by a significant margin was leakage of sensitive data that may arise from prompts, files, chat history or other means was cited by 61 percent of respondents This finding is consistent with broader industry research indicating that the most immediate AI risk for enterprises is not adversarial attack on AI systems but rather the uncontrolled exposure of sensitive data through normal use of those systems. One CISO mentioned known incidents where staff were not aware of AI policies within their financial organization and unintentionally shared customer records in a public LLM for reconciliation, which was quickly detected by the security team.

The second-tier risks center on identity, access, and data exposure through AI architecture. Excessive permissions or weak authorization for AI agents and tools were cited by 33 percent, unauthorized access or exfiltration via RAG connectors or retrieval tools by 27 percent, and sensitive data disclosure via model behavior including training data leakage and inference attacks by 23 percent. Credential or secret exposure through API keys, tokens, system prompts, and plugins was flagged by 19 percent.

Model and application security risks occupy a third tier. Prompt injection or jailbreaks causing harmful actions was cited by 19 percent, data poisoning or integrity attacks on training pipelines by 10 percent, and model theft or extraction of proprietary capabilities by 3 percent. Systems and communications risks, including tool or API invocation tampering and MCP-related threats, were identified by 9 percent. This signals either growing confidence in the maturation of the models and input/output guardrails or at least greater concentration on immediate risk for other activities.

Governance risks form a significant cluster. Lack of auditability, traceability, or monitoring for AI-driven actions was cited by 23 percent — nearly matching the data disclosure concerns and suggesting that financial institutions are deeply concerned about their ability to explain and audit AI decisions, a requirement that is both regulatory and operational. Regulatory or policy non-compliance from AI-driven processing was flagged by 16 percent, and third-party or supply chain risk from AI providers by 10 percent.

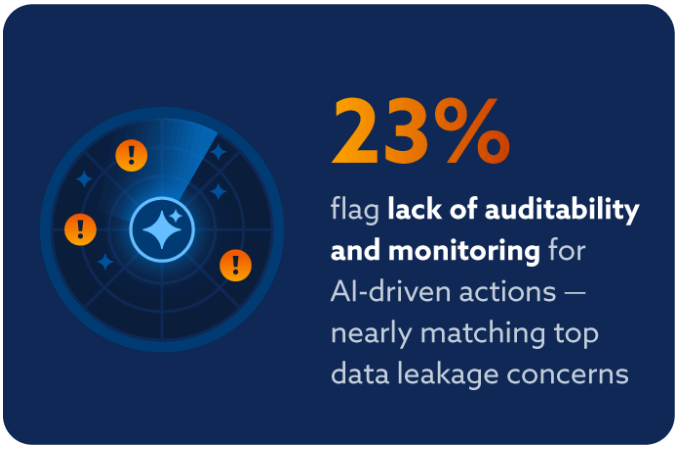

Social engineering amplified by AI, including deepfakes and AI-powered spear-phishing, concerns 16 percent of respondents, while automated or AI-driven attacks that evade traditional detection affect 12 percent. The deepfake concern is grounded in real incidents: the February 2024 case in which a finance worker at engineering firm Arup was deceived by a deepfake video call impersonating senior executives into transferring $25 million demonstrates the scale of financial loss that AI-enabled social engineering can inflict [27]. Industry estimates suggest deepfake-enabled fraud losses may have exceeded $200 million globally in early 2025, while Deloitte has projected that AI-enabled fraud could reach $40 billion annually in the United States by 2027 — though such vendor forecasts should be treated as directional estimates rather than precise predictions.

AI Security Incidents: An Opaque Problem?

One commonly raised concern was the feeling that observability is lagging. Twenty percent report AI-related incidents or breaches, while another 21 percent are unsure whether such incidents have occurred. That uncertainty signals may be a warning that there is an absence of AI-specific monitoring and detection. Traditional monitoring tools may not properly identify some of the generative AI actions or correctly observe and report on newer attack factors.

It was also recognized that nearly a third of the survey respondents skipped this question. Several reasons could account for this but regulatory requirements around reporting incidents may have also influenced the number of responses.

"Our organization strengthened cloud and AI security by embedding security early into the SDLC using policy as code continuous misconfiguration monitoring and automated threat detection. A major challenge was balancing rapid AI experimentation with data protection and regulatory compliance. We address this by enforcing least privilege. IM model access, logging and secure data pipelines. As a result, we reduced cloud security incidents by over 40% and enabled teams to deploy AI workloads faster without compromising security demonstrating that strong governance can coexist with innovation"

Survey respondent,
2026 CSA Financial Services Survey

AI Adoption Barriers: Privacy, Skills and Trust

The barriers to adoption mirror those findings. Concerns about data privacy and security lead at 54 percent. Lack of AI and ML expertise stands at 45 percent. Lack of trust in AI outputs also stands at 45 percent. Unclear regulatory requirements affect 36 percent, unclear ROI 29 percent, budget constraints 25 percent, and integration with legacy systems 22 percent. Organizational resistance is only 5 percent, which suggests that leadership support is generally present even where operational readiness is not. The open-ended comments in the report reinforce the same themes: shadow AI, governance lag, and the growing appeal of centralized AI gateways that allow innovation while maintaining policy enforcement.

Open-ended Comment Opportunity

Several themes emerge repeatedly from the open-ended final question that allowed respondents to share additional thoughts or cover topics not as thoroughly addressed in the questions.

Absence of a well-understood AI strategy for the company, challenge of shadow AI, with employees using unsanctioned AI tools on sensitive financial data; the difficulty of maintaining security governance in the face of rapid AI were comments that were repeated.

One respondent captured the sentiment of many: "AI/ML adoption within my organization is exploding as many groups have started developing AI agents to fit their needs. Governance is one area that is lacking, but eventually it will be coming in the near future."

Yet that sentiment is not universal. One respondent from Nepal referenced the disparity of cost to utilize cloud and the gap in knowledge that limits the adoption of AI in certain regions of the world.

Another described an emerging architectural pattern: "Our strategy is rooted in Zero Trust architecture, ensuring that as we scale our cloud footprint, identity remains our primary perimeter. The biggest challenge we face is the 'shadow AI' phenomenon — employees using unsanctioned LLMs. To combat this, we've implemented a centralized AI gateway that allows us to harness innovation while maintaining strict data loss prevention standards."

Others identify similar early adopter challenges for cloud and mobile devices, stating that the maturity of the security solutions, clear strategy and governance expectations helped to accelerate use and expect the same for AI.

Analysis: Seven Defining Trends

The sections that follow present no new data. Rather, they apply a second analytical lens — CSA's AI agent paired with human expert review — to surface additional industry trends within the data already shared in this report.

1. The AI Security Integration Imperative

The survey's central message is that financial institutions have deployed AI faster than they have secured it. AI security integration is already a top-three investment priority at 44 percent, yet confirmed AI incidents stand at 20 percent and another 21 percent of respondents cannot say whether incidents have occurred. At the same time, 61 percent cite data leakage as the top AI risk. That combination points to a classic control-lag problem.

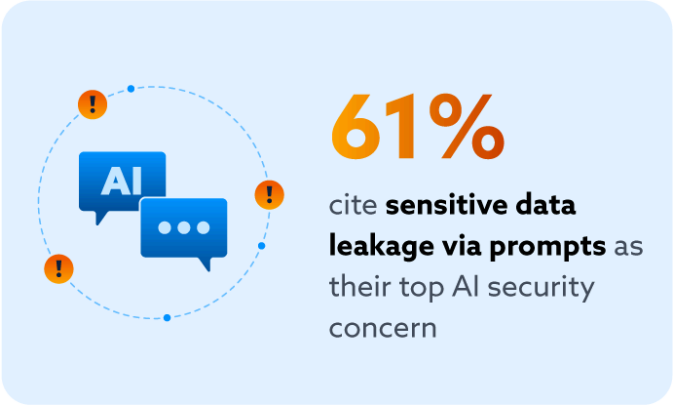

Traditional cloud-security stacks such as CSPM, CWPP, CNAPP, and SIEM were built for virtual machines, containers, and APIs. They do not inherently address prompt-layer leakage, vector-store exposure, model-in-use compromise, RAG permission drift, or AI-generated output risk. The appearance of AISPM, AI-focused DSPM, and shadow-AI detection in respondent write-ins shows that the market is beginning to build the missing layer. In financial services, this integration challenge is amplified by regulatory expectations around model risk management. Existing frameworks such as SR 11-7 remain influential, but they were built for statistical models, not context-dependent, probabilistic large language models [29]. With AI-driven attack risk rising exponentially [45] with increasingly powerful AI models, risk mitigation to the AI risk-control systems themselves will emerge as an important consideration. Confidential computing, as one example, presents an opportunity to limit risks to AI security mechanisms themselves.

2. Third-Party Risk in the Age of AI

Third-party and supply-chain risk has become the sector's leading cloud security concern at 55 percent. That is not theoretical. The July 2024 CrowdStrike incident caused an estimated 8.5 million Windows devices to crash worldwide, with banking-sector losses estimated at $1.15 billion [3]. It was followed by a 15-hour AWS outage in October 2025 affecting more than 4 million users, a Google Cloud IAM failure in June 2025 that disrupted Lloyds Bank, Bank of Scotland, Coinbase, and Robinhood, and multiple Azure

regional outages. Between August 2024 and August 2025, AWS, Azure, and Google Cloud experienced more than 100 combined service outages [20].

Ransomware remains part of the same story. Sixty-five percent of financial services IT professionals reported being hit in 2024, with average recovery costs of $2.58 million [21]. The LockBit attack on ICBC Financial Services disrupted Treasury market operations [22], and the Infosys McCamish breach exposed the data of 6.5 million individuals, including Bank of America customers [23].

DORA turns this risk into a harder supervisory obligation. Effective January 17, 2025, it requires more rigorous ICT third-party inventory, risk assessment, exit planning, and concentration-risk management across approximately 22,000 EU financial entities. By November 2025, 19 Critical ICT Third-Party Providers had been designated for direct EU oversight, with penalties of up to 1 percent of average daily worldwide turnover [4][24]. AI adds another layer because institutions increasingly depend not only on third-party infrastructure, but also on third-party model behavior, updates, training choices, and security controls.

Institutions increasingly depend not only on third-party infrastructure, but also on third-party **model behavior, updates, training choices, and security controls.**

3. The Coming Wave of Agentic Finance

With 62 percent already deploying agents, 38 percent of agent users granting conditional or high autonomy, and 85 percent anticipating autonomous AI payments, financial services is approaching a new operating model. AI systems are moving from advisory roles into active roles within financial processes. The law remains clear on one point: liability stays with the regulated institution. Delegation to an AI agent does not transfer accountability to the model provider, cloud vendor, or software agent.

What remains unsettled is the governance model that sits beneath that legal reality. Institutions will need defensible answers to questions about agent identity, delegated authority, revocation, auditability, and machine-enforceable guardrails. CSA's agentic IAM work provides one answer: agent identity should be treated as a first-class IAM object with verifiable credentials, time-bounded permissions, and auditable logs tied to the human or organizational principal on whose behalf the agent acts [6]. The MAESTRO framework extends that into a broader security model spanning identity, tool authorization, inter-agent communication, and oversight [30].

4. Regulatory Convergence and Persistent Divergence

The global regulatory picture is directionally aligned but tactically fragmented. Europe currently has the clearest combined posture through DORA and the EU AI Act. The EU AI Act entered into force on August 1, 2024. Prohibited practices applied from February 2, 2025. General-purpose AI obligations applied from August 2, 2025. High-risk AI obligations, including those affecting credit scoring and insurance-risk assessment, take full effect on August 2, 2026, with penalties of up to EUR 35 million or 7 percent of global annual turnover [28].

In the United States, the framework remains more distributed and principles-based. SR 11-7 still anchors model-risk expectations, while third-party risk guidance continues to come from multiple prudential regulators [29]. In Asia-Pacific, the report points to APRA's CPS 230 becoming effective July 1, 2025, the Monetary Authority of Singapore's November 2025 AI risk consultation, and the Basel Committee's December 2025 third-party risk principles [31][36][37]. The implication for global firms is practical rather than philosophical: they must translate shared principles into controls that satisfy several frameworks at once.

The implication for global firms is **practical rather than philosophical:** they must translate shared principles into controls that satisfy several frameworks at once.

5. The Technical Security Gap: From Guardrails to AI Data Pipelines

The strongest technical insight in the report is that traditional security controls do not fully address the AI data pipeline. The combination of 61 percent concern over data leakage, 27 percent concern over RAG exfiltration, and 19 percent concern over credential exposure points directly to a new control surface. In a common enterprise RAG architecture, risk appears at ingestion, retrieval, generation, storage, and tool use.

At ingestion, poisoned or malicious documents can persist in the vector store and later shape model behavior. At retrieval, similarity search does not inherently respect document-level authorization, which can allow users to receive content they are not entitled to see. At generation, the model can reveal sensitive retrieved data or produce outputs influenced by adversarial content. Vector databases also

represent a comparatively immature part of the security stack, often lacking the row-level security, field-level encryption, and auditability expected of traditional data platforms.

For agentic architectures, Model Context Protocol and similar tool-use patterns add yet another layer: tool poisoning, overprivileged connectors, unsafe tool descriptions, and weak authorization boundaries. The report's control direction is sound and practical: input and output guardrails, retrieval-layer permission enforcement, least-privilege review for agent-connected tools, and monitoring for drift, anomalous outputs, and extraction attempts [30][32]. It also highlights confidential computing as a promising path for deploying AI on regulated data while reducing exposure to the surrounding infrastructure layer.

6. Systemic Risk: Beyond the Individual Institution

The report is also right to elevate systemic risk. If many institutions rely on the same foundation models, the same cloud concentrations, and similar AI-enabled decision patterns, their errors may become correlated. That creates a stability problem, not just a firm-level problem. The Financial Stability Board highlighted this issue in October 2025 [40].

The same logic applies to markets, credit, and interconnections. AI-driven trading systems trained on similar data may amplify market moves. Credit models based on similar external AI services may tighten or loosen credit in the same direction at the same time. High-autonomy agents acting across interconnected financial networks could propagate mistakes faster than human intervention can contain them. This is why the 5 percent of respondents already granting high autonomy should be read not only as a governance signal, but as an early systemic-risk indicator.

7. From Cloud Security to AI-Cloud Security

The final analytical shift is the most strategic: cloud security and AI security are no longer separable disciplines in financial services. The top cloud risk, third-party concentration at 55 percent, and the top AI risk, data leakage at 61 percent, are not two unrelated problems. They are two expressions of the same underlying challenge: sensitive financial data is now moving through more complex combinations of cloud services, external providers, AI models, retrieval systems, and software agents than the original control architectures were designed to manage.

Observability, supply chain accountability and automated responses emerge as key principles as AI-cloud security evolves into a unified discipline, with shared governance, shared telemetry, and shared responsibility [5].

Opportunities and Recommendations

For Financial Institutions

Financial institutions should integrate AI risk management into existing cloud-security governance rather than treat AI as a separate side program. Data classification should be treated as a prerequisite control. Institutions cannot govern AI well if they do not know which data is sensitive, where it resides, and which systems can reach it.

Identity and access management for non-human and agent identities should be a near-term priority. With 62 percent of organizations deploying AI agents and non-human identity risk ranking among the top cloud concerns, institutions need verifiable agent credentials, scoped permissions, behavioral monitoring, auditable action logs, and automated credential rotation.

Organizations should also conduct AI-specific threat modeling using resources such as the OWASP Top 10 for LLM Applications, MITRE ATLAS, and CSA's MAESTRO framework [30][32]. In practical terms, four technical controls deserve immediate attention: deploy input and output guardrails; enforce access controls at the RAG retrieval layer; review agent-connected tools for least privilege and time-bounded authorization; and monitor models and agents for drift, anomalous outputs, and extraction attempts.

Institutions should also maintain formal inventories of material AI systems, including shadow AI, and extend exit and contingency planning to critical AI services.

For Cloud Service Providers

Cloud service providers serving financial institutions should increase transparency around service changes, model updates, and security events. Financial institutions are accountable for the behavior and resilience of services they do not fully control. Short-notice changes that might be tolerable in consumer software can create material governance problems in regulated environments.

Providers should also invest in financial-services-specific AI governance features that support audit trails, model lineage, explainability, privacy-preserving deployment, and stronger operational transparency. The 54 percent of respondents citing data privacy and security concerns as the top barrier to AI adoption represent both a control problem and a market signal.

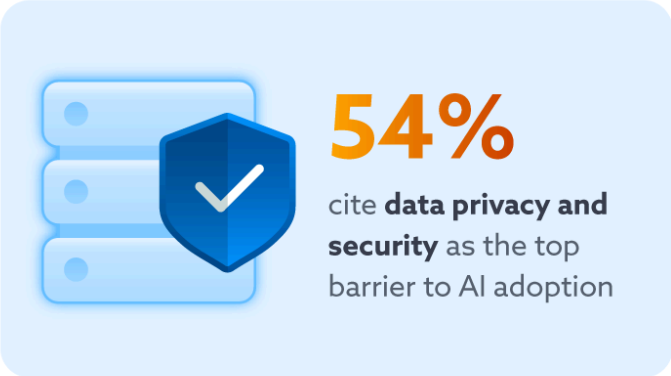

For Regulators

Regulators can reduce both confusion and risk by clarifying how existing supervisory expectations apply to AI. Thirty-six percent of respondents cite unclear regulatory requirements as a barrier to adoption. That ambiguity can slow useful deployment, but it can also encourage institutions to move ahead without mature guardrails.

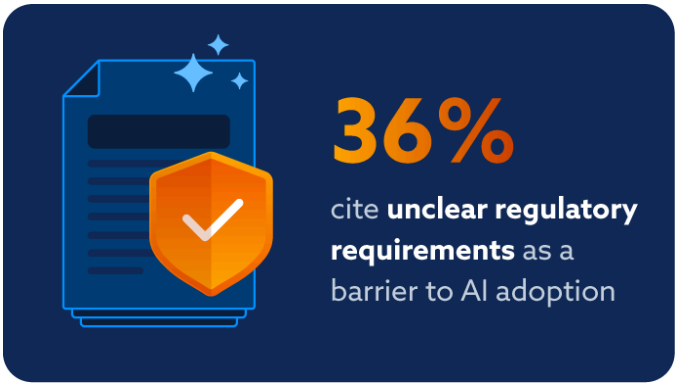

Several actions would improve supervisory effectiveness. Regulators should develop AI-specific examination procedures and train examiners in AI and ML validation. They should consider requiring inventories of AI systems used in material functions, coordinate internationally on AI-related concentration risk, and evaluate whether concentration-risk oversight should extend beyond cloud infrastructure to foundation-model providers. They should also incorporate AI failure scenarios into stress testing and supervisory analysis.

Conclusion

The 2026 survey captures a sector, historically known as slower adopters in the middle of a real transition to leading in the deployment of the latest technological advancements. Cloud has become infrastructure. AI is becoming operations. The question now is how much autonomy institutions and those handling sensitive financial and personal data are willing to delegate to AI systems, and whether security and governance can mature before that autonomy deepens.

The data suggests a disciplined but uneasy story. Ninety-two percent or more are clearly using AI and exploring the latest capabilities. Sixty-two percent have deployed AI agents. Eighty-five percent expect AI to conduct financial transactions. At the same time, 55 percent identify third-party risk as the leading cloud concern, 61 percent identify data leakage as the leading AI concern, and 20 percent report AI-related security incidents while another 21 percent are unsure whether such incidents have occurred.

The institutions most likely to benefit from AI will not be the ones that move slowest. They will be the ones that build controls fast and maintain the trust of consumers, regulators and other stakeholders. Respondents see visibility of AI-specific monitoring, agent-aware identity governance, data classification and policy enforcement at the point of AI interaction, retrieval-layer authorization, and other guardrails mandatory that can operate at production speed.

In financial services, trust remains the business model. The next decade will depend on whether that trust can be extended to systems that increasingly act without a human in the process.

Full Survey Results

Are you involved in the security of cloud and/or AI services at your organization?

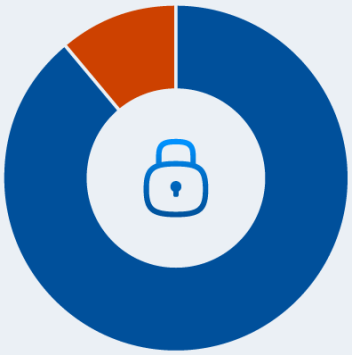

89% Yes

11% No

What best describes your organization's current IT/cloud infrastructure strategy?

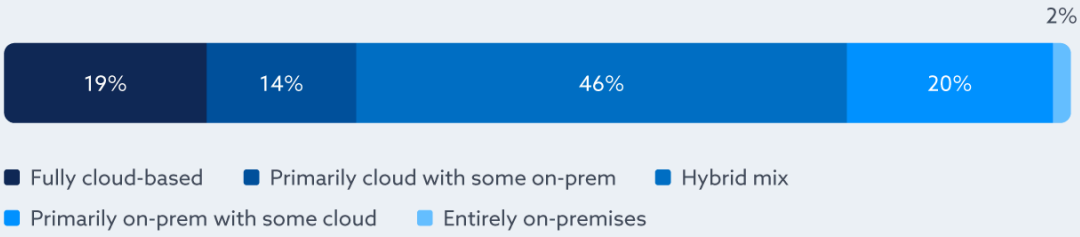

Does your organization store or process regulated financial data or other highly sensitive data in public cloud services?

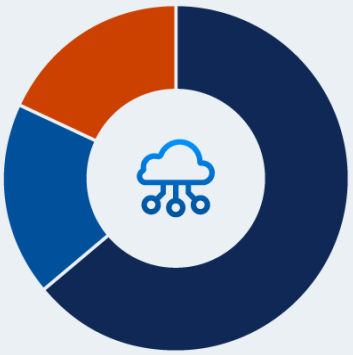

64% Yes, we currently do

18% Not yet, but planning to within 18 months

18% No, and no plans

Which governance or cybersecurity frameworks does your organization use to manage and validate AI or cloud security?

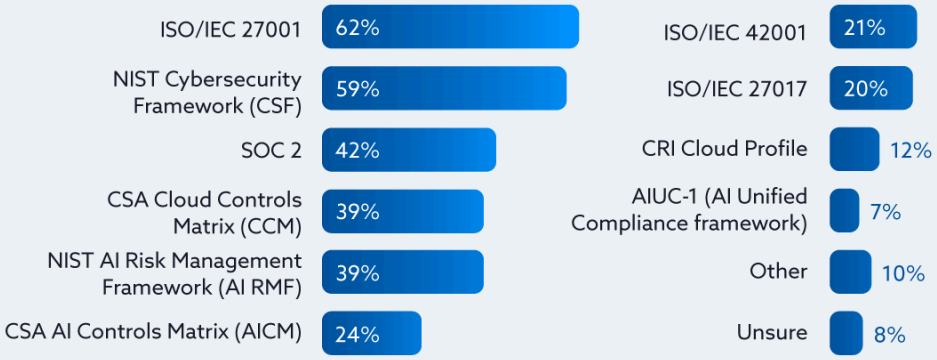

What security measures is your organization using to understand and act on exposure and related risks in your cloud environment?

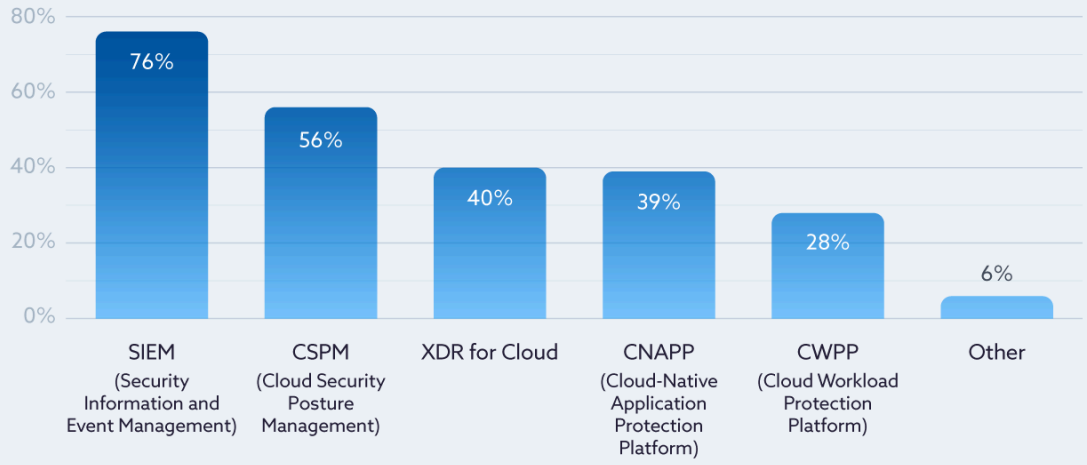

Has your multi-cloud strategy changed in the past three years?

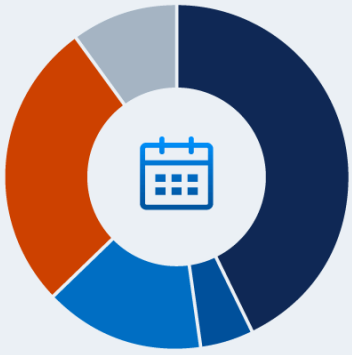

43% Yes, diversified with additional multi-cloud services

5% Yes, moved more services to on-prem

15% Yes, but only reduced the number of multi-cloud providers

27% No significant change

10% We do not have multi-cloud architecture

Are you involved in the security of cloud and/or AI services at your organization?

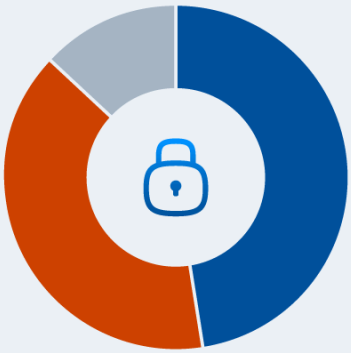

48% Yes

40% No

13% Unsure

What security measures is your organization using to understand and act on exposure and related risks in your cloud environment?

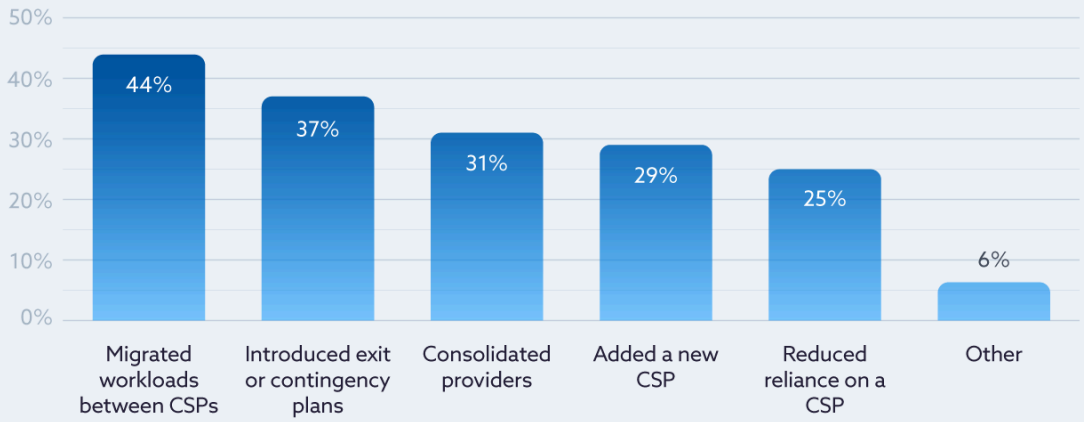

What were the primary drivers for changing the cloud service provider strategy?

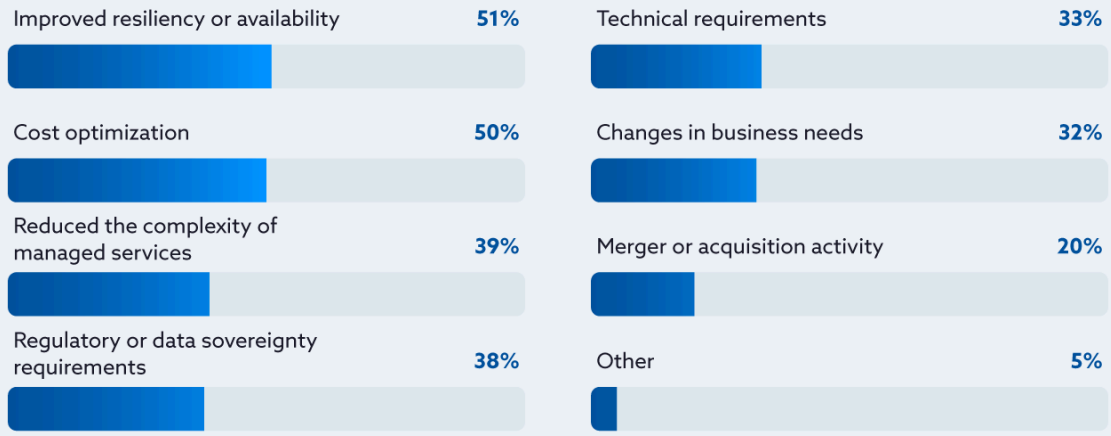

What do you consider the three greatest security risks to your organization's cloud infrastructure?

What are the greatest three barriers to implementing new cloud security capabilities in your organization?

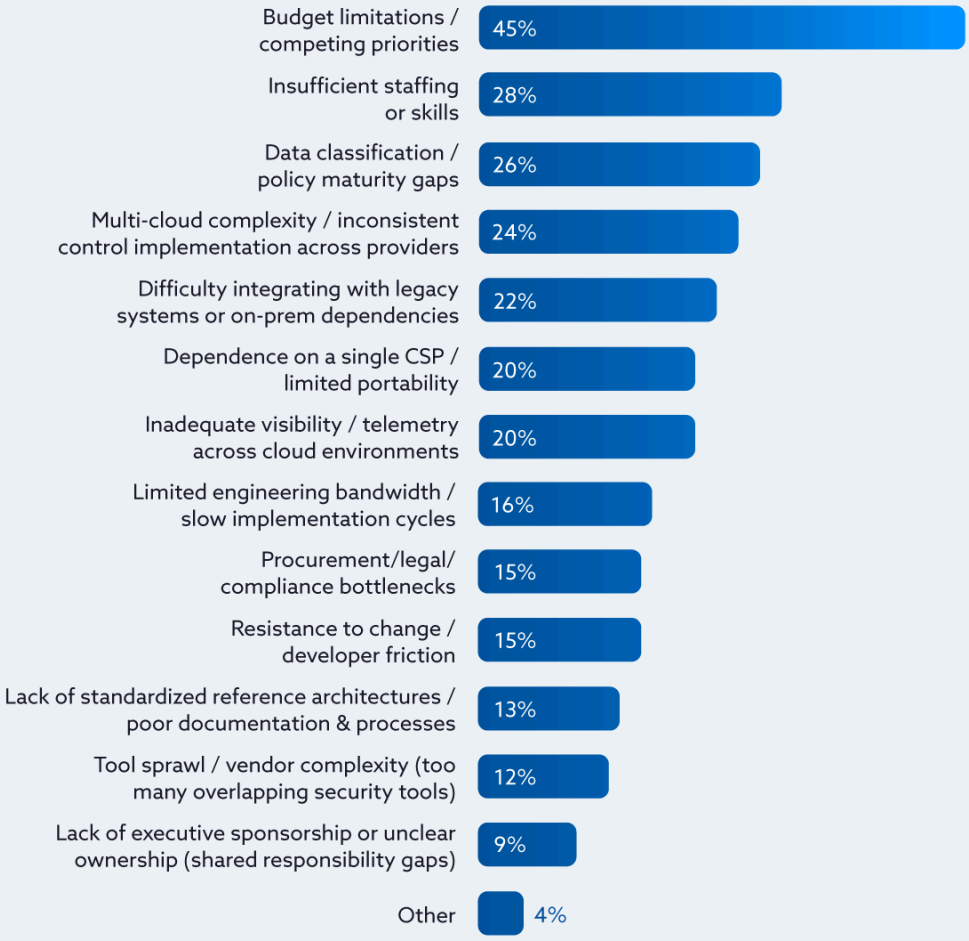

How would you describe senior leadership's level of support for cloud security in your organization?

What are your organization's top 3 cloud infrastructure security priorities over the next 12 months?

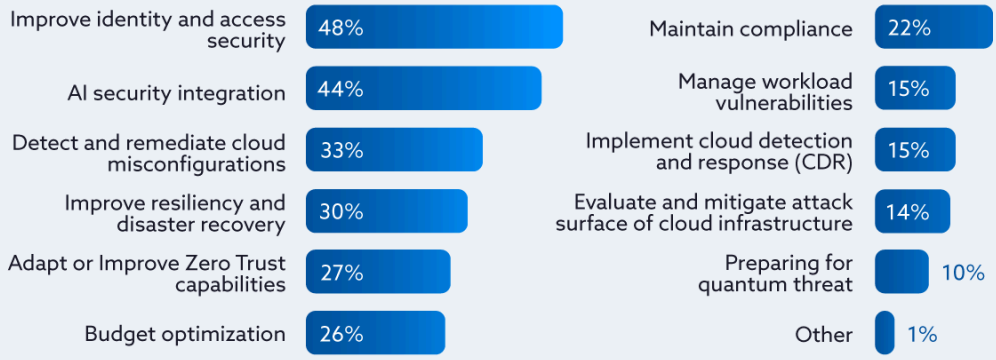

Where is your organization in the adoption of third-party AI/ML technologies in business processes?

Where is your organization in the development of applications with AI/ML technologies?

Early exploration
Experimental or piloting
Active implementation
Advanced adoption
Unsure/NA

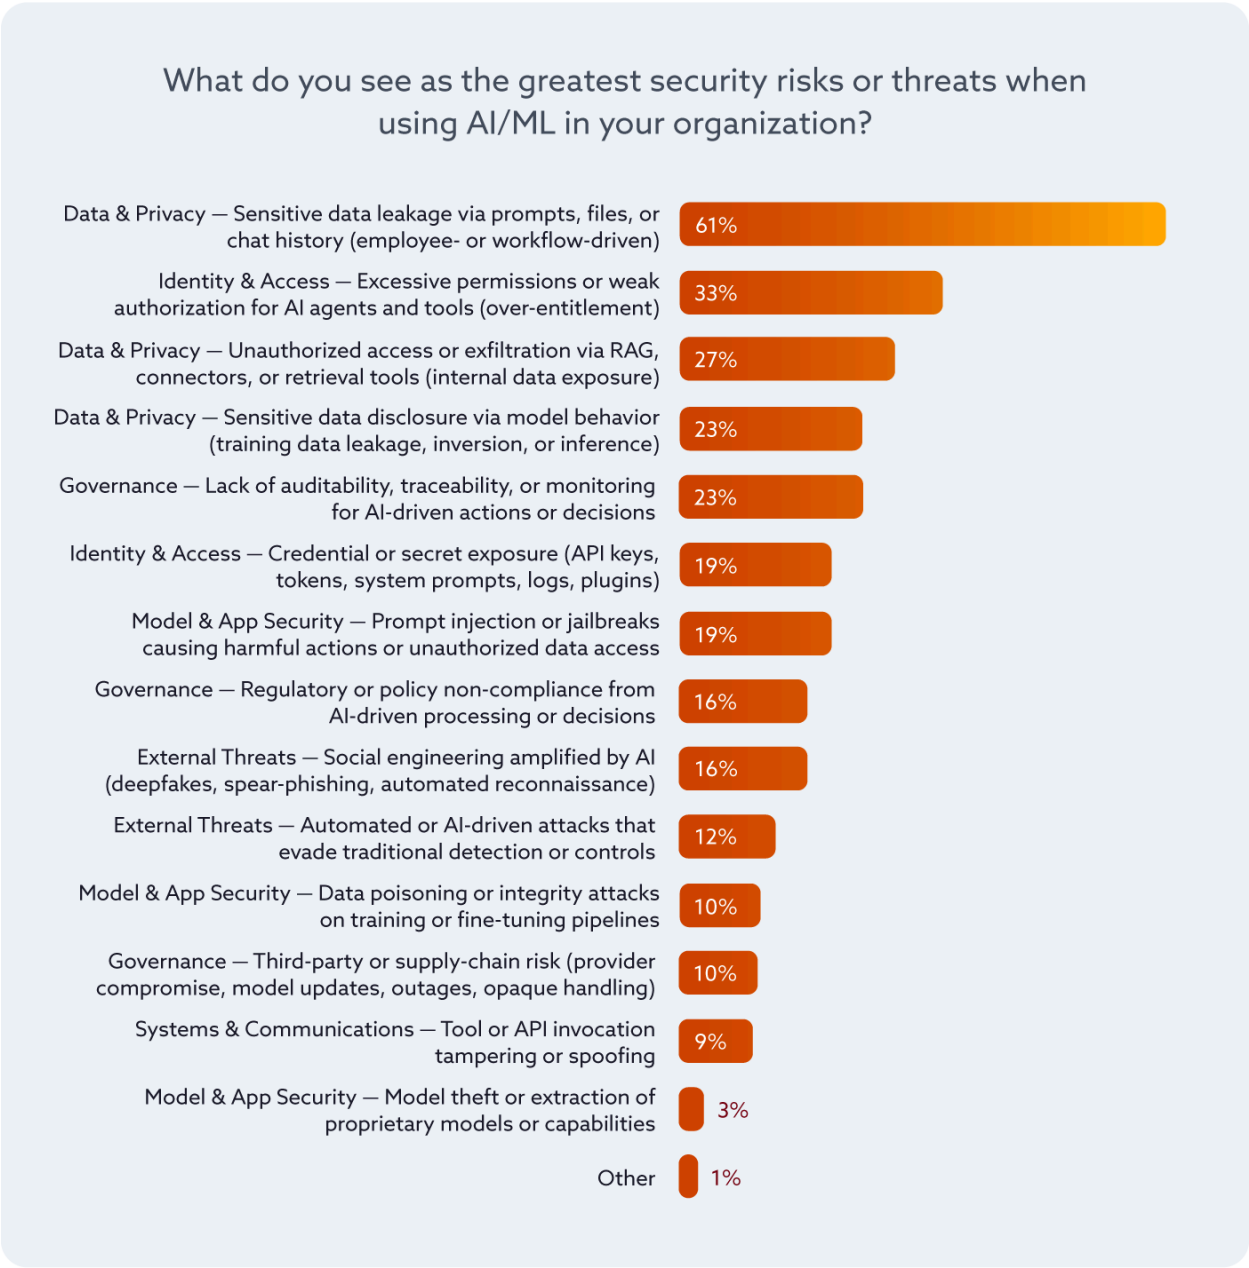

Has your organization experienced any security incidents or breaches related to your AI/ML systems or use of AI?

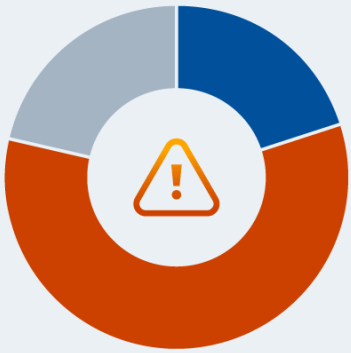

20% Yes

58% No

21% Unsure

What are the biggest challenges or barriers your organization faces in adopting AI/ML solutions for business purposes?

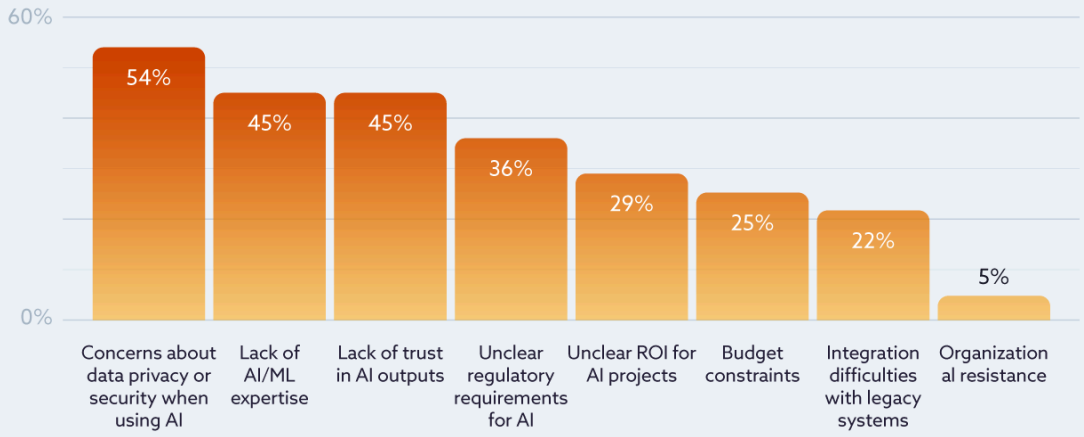

Is your organization using AI agents?

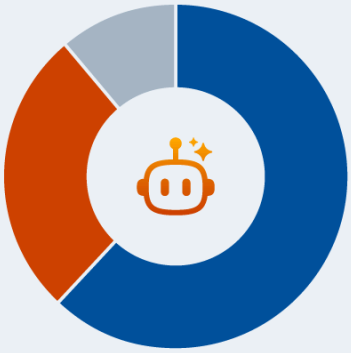

62% Yes

27% No

11% Unsure

In which areas does your organization employ AI agents or AI-driven automation?

What is the highest level of autonomy your organization's AI agents have in performing tasks?

- No Autonomy – Humans make all decisions
- Limited Autonomy – AI performs tasks with human oversight
- Conditional Autonomy – AI acts in low-risk cases under guardrails
- High Autonomy – AI operates independently for critical actions

In the near future, do you believe consumers will be using AI agents to initiate and execute payment transactions without their direct participation in real-time?

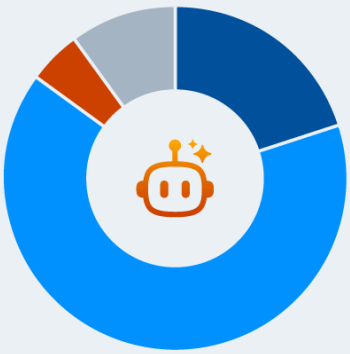

20% Yes

65% Yes but likely with a new model for authorization

5% No

10% Unsure

Appendix: About CSA's Financial Services Program

The Cloud Security Alliance Financial Services Working Group remains the primary venue for industry collaboration on cloud and AI security, regulatory engagement, and best-practice development for the sector. The initiatives most directly tied to this report include the AI Controls Matrix, the MAESTRO framework for multi-agent threat model security, the STAR program, and CSA research on agentic AI identity and access management.

Priority areas for future work identified in the report include agent identity and authorization models, AI-specific incident detection and response, cross-jurisdictional compliance mapping, and financial-services-specific AI security training.

Survey Methodology

This report builds on CSA's 2020 and 2023 financial services cloud surveys and retains many of the earlier questions to support longitudinal comparison [1][2]. It adds a substantial new section focused on AI and machine learning adoption, security, governance, and the emergence of agentic AI in financial services.

The survey was fielded from January 15 to March 1, 2026 and collected 340 responses from professionals involved in cloud computing, AI, cybersecurity, compliance, and risk management within financial services organizations worldwide. Respondents were recruited through CSA's membership network, the Financial Services Working Group, partner organizations, and industry events. The survey included 27 questions covering cloud infrastructure strategy, regulatory compliance, security tooling and practices, AI adoption maturity, AI-specific risks and governance, AI agent deployment, and organizational demographics.

The respondent pool was global and comparatively senior. Regionally, 57 percent of respondents operate in the Americas, 48 percent in EMEA, and 39 percent in APAC. In seniority, 12 percent hold C-level or executive roles, 26 percent are directors, 38 percent are managers, and 22 percent are staff-level practitioners. Banking and credit unions represent 37 percent of respondents, fintech 19 percent, and the balance comes from insurance, investment management, professional associations, cloud service providers, regulators, and other financial services roles. Organization size ranges from under 100 employees at 17 percent to more than 10,000 at 29 percent.

As with any voluntary survey, the results should be read as directional indicators rather than statistically precise measures of the entire industry. Their greatest value lies in trend direction, relative priority, and the consistency of risk signals across cloud and AI domains.

Demographics

The regional distribution reflects a more balanced global sample than the 2023 survey, with increased participation from EMEA and APAC. The multi-select regional question reflects the multinational operations of many respondents.

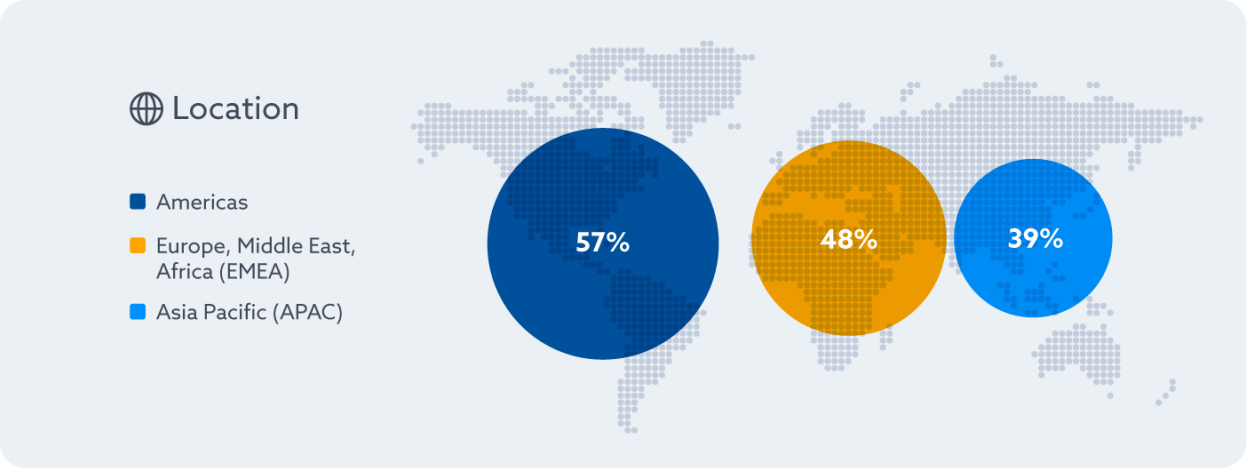

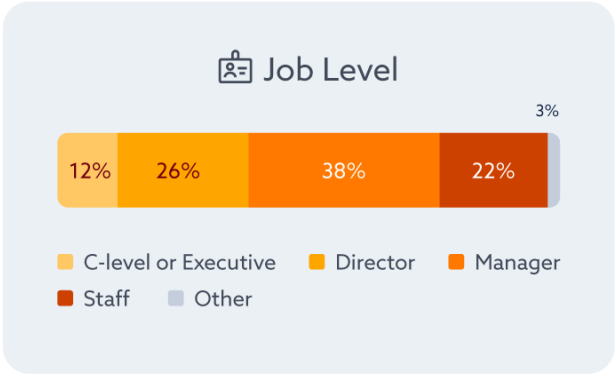

Industry Segment

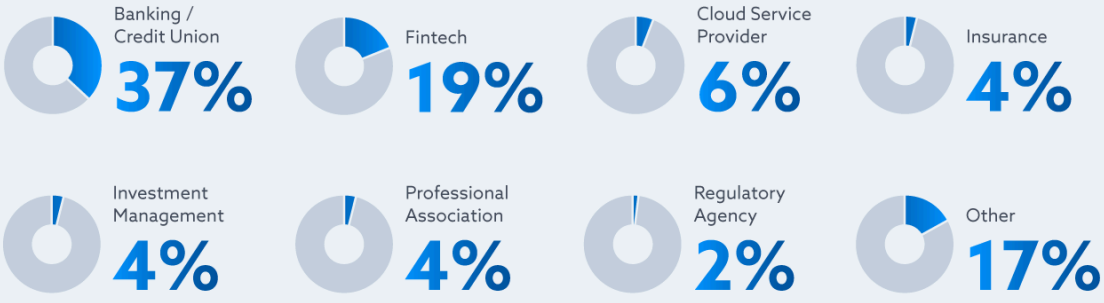

Involvement in Cloud/AI Security

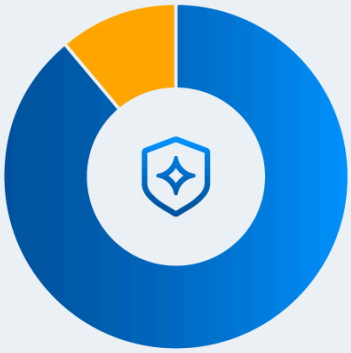

References

[1] Cloud Security Alliance. "Cloud Usage in the Financial Services Sector." 2020.

[2] Cloud Security Alliance. "State of Financial Services in Cloud." June 2023.

[3] CrowdStrike. "Preliminary Post Incident Review (PIR) — Content Configuration Update Impacting the Falcon Sensor." July 2024.

[4] European Parliament and Council of the European Union. "Regulation (EU) 2022/2554 — Digital Operational Resilience Act (DORA)." Official Journal of the European Union, December 2022. Enforcement effective January 17, 2025.

[5] Cloud Security Alliance. "AI Controls Matrix (AICM) v1.0." July 2025.

[6] Cloud Security Alliance AI Safety Initiative. "Agentic AI Identity and Access Management: A New Approach." August 2025.

[7] FS-ISAC. "FS-ISAC Urges Global Coordination for Migration to Post-Quantum Cryptography in Financial Services." September 25, 2025.

[8] Accenture. "Generative AI in Banking: How to Deploy It and the Workforce Impact." 2024.

[9] McKinsey & Company. "Capturing the Full Value of Generative AI in Banking." 2024.

[10] CNBC. "JPMorgan Blueprint for an AI-Connected Megabank." September 2025; Prism News. "JPMorgan Boosts 2026 Tech Budget to $19.8 Billion with $1.2 Billion for AI." 2026.

[11] Bank of America Newsroom. "AI Adoption by BofA's Global Workforce Improves Productivity." April 2025.

[12] CNBC. "Goldman Sachs Autonomous Coder Pilot Marks Major AI Milestone." July 11, 2025; Fast Company. "Meet Devin, Goldman Sachs' New AI Software Engineer." 2025.

[13] Bloomberg. "Citigroup Assembles Banking Team Focused on AI Infrastructure." February 25, 2026.

[14] Visa Investor Relations. "Visa and Partners Complete Secure AI Transactions, Setting the Stage for Mainstream Adoption in 2026." December 18, 2025.

[15] Mastercard. "Mastercard Unveils Agent Pay: Pioneering Agentic Payments Technology." April 2025; Mastercard. "Santander and Mastercard Complete Europe's First Live End-to-End Payment Executed by an AI Agent." 2026.

[16] Stripe / Paradigm. "Machine Payments Protocol (MPP) for AI Agent Transactions." March 2026.

[17] Google Cloud. "Announcing Agent Payments Protocol (AP2)." September 16, 2025.

[18] Edgar, Dunn & Co., cited in Payments Dive. "Visa, Mastercard Race in Agentic AI Commerce Payments." 2025.

[19] Accenture Banking Blog. "Agentic Payments in Commerce." 2026.

[20] IncidentHub. "Major Cloud Outages 2025." 2025; TechTarget. "Cloud Outages Expected to Be the New Normal in 2026." 2025.

[21] Inveni IT. "Ransomware Attacks on Financial Services 2024." 2024; FS-ISAC. "Navigating Cyber 2025." 2025.

[22] CNBC. "ICBC, the World's Biggest Bank, Hit by Ransomware Cyberattack." November 10, 2023.

[23] Cybersecurity Dive. "Bank of America Customer Data Exposed in Infosys McCamish Breach." 2024; SecurityWeek. "Infosys to Pay $17.5 Million in Settlement." 2025.

[24] Morgan Lewis. "DORA: EU Regulators Announce List of Critical ICT Third-Party Providers." November 2025.

[25] Cybersecurity Tribe. "Research Reveals 44% Growth in NHIs from 2024 to 2025." 2025; Help Net Security. "Identity Security Permissions Sprawl." December 2025.

[26] ConductorOne. "Key Takeaways from the 2025 Financial Identity Security Report." 2025.

[27] CNN Business. "Finance Worker Pays Out $25 Million After Video Call with Deepfake 'Chief Financial Officer'." February 4, 2024.

[28] European Parliament and Council of the European Union. "Regulation (EU) 2024/1689 — Artificial Intelligence Act." Official Journal of the European Union, July 2024.

[29] Board of Governors of the Federal Reserve System. "Supervisory Letter SR 11-7: Guidance on Model Risk Management." April 4, 2011.

[30] Cloud Security Alliance. "MAESTRO: Multi-Agent Environment, Security, Threat, Risk, and Outcome Framework." February 2025.

[31] Australian Prudential Regulation Authority (APRA). "Prudential Standard CPS 230 — Operational Risk Management." Effective July 1, 2025.

[32] OWASP. "OWASP Top 10 for Large Language Model Applications." 2025; MITRE. "ATLAS: Adversarial Threat Landscape for Artificial-Intelligence Systems." 2024.

[33] FS-ISAC. "Guidance on the Future State of Generative AI in Financial Services." March 24, 2025.

[34] FS-ISAC. "Cyber Fraud Prevention Framework." April 1, 2025.

[35] NIST. "AI Risk Management Framework (AI RMF 1.0)." January 2023.

[36] Monetary Authority of Singapore. "Consultation Paper on Guidelines on Artificial Intelligence Risk Management." November 13, 2025.

[37] Basel Committee on Banking Supervision. "Principles for the Sound Management of Third-Party Risk (BCBS d605)." December 10, 2025.

[38] IBM. "Cost of a Data Breach Report 2024 — Financial Industry." 2024.

[39] U.S. Government Accountability Office. "GAO-25-107197: Artificial Intelligence — Use and Oversight in Financial Services." May 2025.

[40] Financial Stability Board. "Monitoring Adoption of Artificial Intelligence and Related Technologies." October 2025.

[41] Gartner. "Worldwide End-User Spending on Information Security to Total $213 Billion in 2025." July 2025.

[42] WEF / Accenture. "Artificial Intelligence in Financial Services 2025." 2025.

[43] UK PRA. "Supervisory Statement SS1/23: Model Risk Management Principles for Banks." Effective May 17, 2024.

[44] FCA. "FCA 2026 Payments Regulatory Priorities Report." March 2026.

[45] HBR, "AI Is Reshaping Cyber Risk. Boards Need to Manage the Threat., April 2026.