CSAI Foundation

Cloud Security Alliance AI Safety Initiative

State of Cloud and AI for Financial Services

2026 Industry Survey Report

Unofficial AI-assisted Research

2026-03-31

© 2026 Cloud Security Alliance. Unofficial AI-assisted Research.

© 2026 Cloud Security Alliance. Some rights reserved.

You may download, store, display, view, print, redistribute, and link to this document in its original, unmodified form, provided that attribution to the Cloud Security Alliance is maintained and all trademark and copyright notices remain intact.

This document may not be modified or altered. You may quote portions of the document as permitted by the Fair Use provisions of the United States Copyright Act, provided that attribution is given to the Cloud Security Alliance.

This document may be shared on professional and social media platforms in its original form with attribution.

This document was generated with AI assistance and has not undergone official CSA review and approval processes.

Table of Contents

© 2026 Cloud Security Alliance. Unofficial AI-assisted Research.

Executive Summary

Cloud adoption in financial services is now effectively universal, and artificial intelligence has moved from experimentation into production at scale. The question is no longer whether financial institutions will adopt AI, but whether their security, governance, and operational models are evolving fast enough to manage it. This third iteration of the Cloud Security Alliance's financial services survey, conducted in early 2026 with 340 respondents across the Americas, Europe, the Middle East, Africa, and Asia Pacific, documents both the velocity of this transition and the gaps it has opened.

The 2026 survey results point to an accelerating transition toward agent-mediated financial operations. Sixty-two percent of organizations report deploying AI agents — a figure that encompasses a broad spectrum of implementations, from customer service chatbots with limited decision authority to systems executing trades and compliance actions autonomously. Among those deploying agents, over one-third have granted conditional or higher levels of autonomy. At the same time, 85 percent of respondents believe AI agents will soon initiate and execute financial transactions on behalf of users — though 65 percent insist this will require entirely new authorization models — signaling a significant change in how financial activity may be conducted and authorized.

However, this rapid adoption appears to be outpacing the industry's ability to secure and govern these systems. Nearly one in five organizations (20 percent) report experiencing AI-related security incidents, while an additional 21 percent are unsure whether such incidents have occurred, highlighting a gap in AI security visibility and monitoring. The data suggests that many financial institutions are deploying AI systems without fully understanding their risk exposure, though respondent uncertainty may also reflect organizational silos or the early stage of many deployments.

The nature of that risk is becoming clearer. Data leakage has emerged as the dominant AI security concern, cited by 61 percent of respondents — far exceeding concerns about model attacks, prompt injection, or adversarial techniques. AI risk in financial services is not primarily about system compromise; it is about the unintended exposure of sensitive financial data through normal AI usage patterns, including prompts, chat histories, retrieval-augmented generation connectors, and training data.

Traditional cloud security risks remain unresolved and are increasingly intertwined with AI. Third-party and supply chain risk is now the top cloud security concern, identified by 55 percent of respondents, reflecting both the CrowdStrike outage of July 2024 and growing regulatory pressure from DORA and similar frameworks worldwide. Cloud misconfigurations (52 percent) and identity-related risks — including non-human identities (24 percent) — continue to present persistent challenges, now compounded by the introduction of AI agents as a new class of identity requiring credentials and access control.

Security priorities are evolving accordingly. Improving identity and access management (48 percent) and integrating AI security into cloud environments (44 percent) are now the top investment areas. Yet organizations face significant barriers: budget constraints (45 percent), lack of AI expertise (45 percent), and data classification gaps (26 percent) limit their ability to implement effective controls. The governance landscape remains uneven — while cloud security frameworks such as ISO 27001 (62 percent) and NIST CSF (59 percent) are widely adopted, AI-specific governance frameworks show lower and more fragmented adoption, with NIST AI RMF at 39 percent and CSA's AI Controls Matrix at 24 percent.

Top Five Insights from the 2026 Survey

#	Insight	Key Data Point
1	**AI is in production — not experimental.** Only 2% of financial institutions report no AI usage. Over one-third are actively implementing AI in production, with an additional 9% at advanced adoption.	98% AI engagement; 43% in active implementation or advanced adoption
2	**AI agents are mainstream — and gaining autonomy.** The majority of financial institutions are deploying AI agents across customer service, cybersecurity, fraud detection, and back-office operations, with over one-third granting conditional or higher autonomy.	62% deploying AI agents; 38% with conditional or high autonomy
3	**The industry expects autonomous financial transactions.** An overwhelming majority of respondents believe AI agents will initiate and execute payments on behalf of consumers — but most insist this will require entirely new authorization models.	85% anticipate agentic payments; 65% require new authorization frameworks
4	**AI risk is primarily a data problem.** Sensitive data leakage through AI interactions — via prompts, chat histories, RAG connectors, and training data — is the dominant AI security concern, far exceeding worries about model attacks or adversarial techniques.	61% cite data leakage as top AI risk
5	**Security observability is lagging adoption.** One in five organizations has confirmed an AI-related security incident. A further one in five cannot determine whether incidents have occurred — pointing to gaps in AI-specific monitoring and detection.	20% confirmed AI incidents; 21% unsure

Taken together, the findings document a transition underway in financial services: the emergence of an operating model in which AI systems — and increasingly AI agents — participate directly in financial decision-making and execution. This transition introduces risks that existing cloud security and governance frameworks were not designed to address, alongside systemic risk dimensions including correlated model failures across institutions using the same foundation models and concentration risk in AI model providers. The institutions best positioned for this next phase will be those that invest in AI-specific security monitoring and controls, identity systems that extend to non-human and agent identities, data governance that operates at the point of AI interaction, and authorization models designed for autonomous systems.

This report presents the complete survey results with longitudinal comparisons to CSA's 2020 and 2023 surveys, contextualizes the findings within the current regulatory environment including DORA, the EU AI Act, and evolving supervisory guidance worldwide, and offers recommendations for financial institutions, cloud service providers, and regulators.

Survey Methodology

The methodology for this report was designed to build upon and extend the longitudinal baseline established by CSA's 2020 Cloud Usage in the Financial Services Sector report and the 2023 State of Financial Services in Cloud report [1][2]. Many questions from the prior surveys were retained to enable direct trend comparisons, while a substantial new section was added to address AI and machine learning adoption, security, governance, and the emergence of agentic AI in financial services.

The survey was fielded between January 15 and March 1, 2026, collecting 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 comprised 27 questions spanning cloud infrastructure strategy, regulatory compliance, security tools and practices, AI adoption maturity, AI-specific risks and governance, AI agent deployment, and organizational demographics.

The respondent pool reflects a global and senior audience. Regionally, 57 percent of respondents operate in the Americas, 48 percent in EMEA, and 39 percent in APAC (percentages exceed 100 because respondents operating in multiple regions could select all that apply). In terms of seniority, 12 percent hold C-level or executive positions, 26 percent serve as directors, 38 percent are managers, and 22 percent are staff-level practitioners. The industry composition is anchored in banking and credit unions at 37 percent, followed by fintech at 19 percent, with insurance, investment management, professional associations, cloud service providers, regulatory agencies, and other financial services roles constituting the remainder. Organization

sizes range from under 100 employees at 17 percent to over 10,000 employees at 29 percent, providing a cross-section of perspectives from community banks and startups to global systemically important financial institutions.

An important qualification: as with any voluntary survey, the results should be interpreted as directional indicators of industry sentiment and maturity rather than statistically precise measurements of the entire financial services sector. Respondents who engage with CSA and its working groups may have higher awareness of cloud security practices than the broader industry. The data is most valuable when examined for trends over time and for the relative priority ordering of concerns, risks, and investment areas.

Survey Results

Cloud Adoption: A Universal Baseline with New Complexities

The adoption of cloud services by financial institutions has reached a plateau of near-universality. Only 1.7 percent of respondents describe their organizations as entirely on-premises, a figure that underscores how thoroughly cloud computing has permeated even the most conservative financial institutions. The remaining 98.3 percent operate across a spectrum from primarily on-premises with some cloud (20 percent) through hybrid environments (46 percent) to primarily cloud with some on-premises (14 percent) and fully cloud-based architectures (19 percent). The hybrid model remains the modal strategy, reflecting the reality that most financial institutions continue to operate legacy core banking systems alongside cloud-native applications and services.

This distribution represents a continuation of the trajectory documented in CSA's prior surveys. In 2020, 91 percent of respondents reported using some form of cloud computing. By 2023, that figure had risen to 98 percent. The 2026 data confirms that cloud adoption in financial services has effectively saturated — the remaining question is not adoption but architecture. The 19 percent of organizations that describe themselves as fully cloud-based represents a notable population of institutions, predominantly fintechs and digitally native firms, that have completed the migration journey entirely. For the majority, however, the hybrid reality persists, driven by regulatory requirements for data residency, the inertia of mainframe-based core systems, and the economic calculus of migration.

The placement of regulated financial data in public cloud has continued its steady expansion. Sixty-four percent of respondents currently store or process regulated financial data in public cloud services, with an additional 18 percent planning to do so within 18 months. Only 18 percent report no plans to place regulated data in public cloud, down from 25 percent in 2023. When these figures are compared longitudinally, the

trajectory is clear: from 59 percent storing regulated data in cloud in 2023, the industry has moved to 64 percent in 2026, with the planning pipeline suggesting that more than 80 percent of financial institutions will have regulated data in public cloud by 2028.

Confidence in using cloud for business-critical workloads has also continued to strengthen. In the 2023 survey, 32 percent of organizations reported having a majority of their business-critical workloads (50 percent or more) in production at cloud providers, nearly double the 17 percent recorded in 2020. The 2026 data shows this trend continuing as financial institutions move past the experimentation phase and commit core operations to cloud infrastructure.

Multi-Cloud Strategy: Active Diversification and Emerging Exit Planning

The multi-cloud landscape in financial services is undergoing active transformation. Nearly half of respondents (48 percent) report changing their cloud service provider strategy within the past twelve months, a remarkably high rate of strategic adjustment that reflects the combined influence of regulatory pressure, cost optimization efforts, and resilience requirements.

Among those who changed their CSP strategy, the most common action was migrating workloads between cloud providers, reported by 44 percent. This was followed by introducing exit or contingency plans (37 percent), consolidating providers (31 percent), adding a new CSP (29 percent), and reducing reliance on a single provider (25 percent). The coexistence of diversification and consolidation may appear contradictory, but it reflects the nuanced reality of multi-cloud strategy: organizations are simultaneously rationalizing their portfolio of cloud relationships while ensuring they are not critically dependent on any single provider.

The primary drivers of CSP strategy changes reveal the priorities shaping financial services cloud architecture. Improved resiliency and availability leads 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 post-CrowdStrike environment, in which a single software update from a third-party vendor caused cascading failures across financial institutions worldwide in July 2024, demonstrating in concrete terms the systemic risk that regulators had been warning about for years [3].

Despite the shift toward diversification, only 5 percent of respondents have moved services back to on-premises infrastructure. The direction of travel remains firmly toward cloud — the question is how to distribute workloads across providers to optimize for resilience, cost, and regulatory compliance. The introduction 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, reflecting its influence both within the United States and increasingly as a reference model internationally. SOC 2 attestations are used by 42 percent, driven by the need to demonstrate assurance to customers and partners. The CSA Cloud Controls Matrix maintains strong adoption at 39 percent, serving as a cloud-specific complement to broader frameworks.

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 NIST brand and its early release in January 2023. The CSA AI Controls Matrix (AICM), which extends the cloud controls framework to address AI-specific risks, 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 — a notable figure for a relatively young standard. The AI Unified Compliance framework (AIUC-1) trails at 7 percent.

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 financial institutions bear. Several respondents noted that they rely on internal control standards derived from multiple frameworks, suggesting that the industry norm is not adherence to a 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 — whether the NIST AI RMF, ISO/IEC 42001, the CSA AI Controls Matrix, or institution-specific composites — to serve as unifying layers that reduce duplication and align cloud and AI security governance within coherent control structures.

Security Posture: Tools, Risks, and Priorities

The security tooling landscape in financial services cloud environments continues to mature, and the longitudinal trajectory from 2023 to 2026 reveals significant but uneven investment. Security Information and Event Management (SIEM) systems are the most widely deployed, used by 76 percent of respondents,

reflecting the centrality of log aggregation and alerting in regulated environments where audit trails are mandatory. Cloud Security Posture Management (CSPM) tools follow at 56 percent — up from an estimated 25 to 29 percent in 2023 industry surveys — representing the fastest-growing tool category and reflecting regulatory imperatives from PCI DSS v4.0 and DORA. Extended Detection and Response (XDR) for cloud environments is used by 40 percent, up from fewer than 5 percent in 2023 per Gartner estimates, driven by SOC consolidation around platforms like CrowdStrike Falcon and Microsoft Sentinel. Cloud-Native Application Protection Platforms (CNAPP) are at 39 percent and Cloud Workload Protection Platforms (CWPP) at 28 percent.

The juxtaposition of 56 percent CSPM adoption with 52 percent citing misconfiguration as a top risk reveals a critical coverage gap: the margin between risk exposure and detection capability is effectively zero for the fraction of organizations without CSPM. In a financial services context, operating a cloud environment without continuous posture monitoring is analogous to running a trading operation without position risk monitoring — the risk is identified, the tool to manage it is widely available, yet a meaningful proportion of organizations remain uncovered.

The write-in responses to the security tools question surface emerging categories: AI Security Posture Management (AISPM), Data Security Posture Management (DSPM), Continuous Threat Exposure Management (CTEM), and shadow AI detection tools. These responses, while representing a small minority, signal the direction of tool evolution as organizations recognize that AI workloads require purpose-built security observability.

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, a finding that would have been improbable five years ago when cloud misconfigurations and insider threats topped most surveys. 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 — essentially equivalent to API security and ahead of human identity risk — 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 institutions are beginning to recognize non-human identity management as a distinct security challenge, but the scale of the problem — particularly as AI agents add yet another category of non-human identity requiring credentials and permissions — demands more urgent attention than the 24 percent figure might suggest.

Security Priorities: Identity and AI Lead the Agenda

When asked to identify their top three cloud infrastructure security priorities for the next twelve months, respondents revealed an agenda dominated by identity management and AI security integration. Improving identity and access security leads at 48 percent, a priority that encompasses both human and non-human identity management, privileged access controls, and the challenge of governing AI agent credentials. AI security integration follows closely at 44 percent, marking the first time an AI-specific capability has appeared among the top security investment priorities in CSA's financial services surveys.

The remainder of the priority list includes detecting and remediating cloud misconfigurations (33 percent), improving resiliency and disaster recovery (30 percent), adapting or improving Zero Trust capabilities (27 percent), budget optimization (26 percent), maintaining compliance (22 percent), managing workload vulnerabilities (15 percent), implementing cloud detection and response (15 percent), evaluating attack surface (14 percent), and preparing for quantum threats (10 percent).

The 27 percent prioritizing Zero Trust represents continued momentum for an approach that has been gaining traction since the 2023 survey, where Zero Trust was the topic of greatest interest for 72 percent of respondents. The shift from interest to implementation is notable — organizations are moving past conceptual alignment with Zero Trust principles and investing in the architectural changes required to implement continuous verification, least-privilege access, and microsegmentation in cloud environments.

The quantum preparation priority at 10 percent, while small, is significant for financial services. Cryptographic agility and post-quantum migration planning are particularly urgent for institutions that process long-lived financial instruments, where data encrypted today may need to remain confidential for decades. FS-ISAC's September 2025 guidance urging global coordination for post-quantum cryptography migration underscores the industry-level concern [7].

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 follows at 28 percent. Data classification and policy maturity gaps affect 26 percent, multi-cloud 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.

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) has 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 circa 2018: 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 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 identified in this survey. The organizations most likely to struggle with AI security governance are those where leadership engagement is weakest, as AI security requires cross-functional coordination between security, data science, compliance, and business teams that only executive sponsorship can drive effectively.

AI Adoption: From Experimentation to Production at Scale

The 2026 survey marks the first time CSA has included comprehensive questions on AI adoption in its financial services research, and the results reveal an industry that has moved well beyond experimentation. When asked about the adoption of third-party AI and ML technologies in business processes, only 2 percent

of respondents report that their organization is not using AI at all. Fourteen percent are in early exploration, 35 percent are in experimental or piloting stages, 34 percent are actively implementing AI in production, and 9 percent describe themselves as advanced adopters with AI deeply integrated into operations.

The pattern is nearly identical for the development of applications with AI and ML technologies: 17 percent early exploration, 33 percent experimental or piloting, 35 percent active implementation, and 9 percent advanced adoption. The close alignment between third-party AI adoption and in-house AI development suggests that financial institutions are pursuing a dual strategy — consuming AI capabilities from vendors while simultaneously building proprietary AI applications tailored to their specific regulatory and operational requirements.

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 large language model capabilities, the availability of enterprise-grade AI platforms from major cloud providers, competitive pressure from fintech firms that are AI-native by design, and growing evidence that AI can deliver measurable improvements in fraud detection, customer service efficiency, and operational productivity. Consulting firms have published optimistic projections for AI in banking — 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 scale of investment by individual institutions underscores this acceleration. JPMorgan Chase allocated $19.8 billion for technology in 2026, with $1.2 billion earmarked specifically for AI, and its LLM Suite now serves over 200,000 employees — the largest AI deployment on Wall Street [10]. Bank of America reports that over 210,000 associates use its generative AI assistant, achieving 90 percent workforce AI adoption and reducing IT service desk queries by more than 50 percent [11]. Goldman Sachs became the first major bank to deploy autonomous AI software engineers across its 12,000-strong developer workforce in July 2025, with CIO Marco Argenti describing the initiative as building a "hybrid workforce" [12]. Citigroup has deployed proprietary generative AI tools to over 150,000 employees across 80 countries [13]. These are not pilot programs; they represent production-scale AI integration across the world's largest financial institutions.

AI Agents: The Autonomy Frontier

The most consequential finding in the 2026 survey may be the extent to which financial institutions have deployed AI agents — systems capable of performing tasks with varying degrees of autonomy. Sixty-two percent of respondents report that their organization is using AI agents. This figure should be interpreted with an important caveat: the survey defined agents broadly to include any AI system performing tasks with varying degrees of independence, which encompasses both sophisticated autonomous systems and simpler chatbots with decision-tree logic. The true prevalence of agentic AI systems with meaningful autonomous

capability is likely lower. Nevertheless, the breadth of adoption is significant. Only 27 percent are not using agents, and 11 percent are unsure — a response that may itself indicate shadow AI agent deployment within business units that security teams have not identified.

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 robo-advisors 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.

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.

The 5 percent figure for high-autonomy AI — representing approximately 10 respondents among agent-deploying organizations — should be interpreted cautiously given the small absolute number, but it signals that some financial institutions have begun delegating critical operational decisions to autonomous AI systems. These organizations have made a calculated determination that certain functions are better served by autonomous AI than by human-mediated processes. While the survey does not specify which functions have been delegated, the write-in comments suggest that high-autonomy applications are concentrated in algorithmic trading, real-time fraud scoring, and security operations center triage — domains where the speed of machine decision-making provides a material advantage and where human latency creates unacceptable risk. The historical parallel to algorithmic trading is instructive: when financial institutions first deployed automated trading algorithms in the 1990s and 2000s, many did so without adequate circuit breakers, kill switches, or human oversight protocols. Knight Capital's $440 million algorithmic trading loss in 45 minutes in 2012 catalyzed the development of robust governance frameworks for automated financial decision-making. AI agents in financial services today occupy a similar position: the technology is deployed, the risks are present, and the governance frameworks are still forming.

Agentic Payments: A New Authorization Paradigm

The survey asked respondents whether they believe consumers will use AI agents to initiate and execute payment transactions without their direct participation in real time, Only 5 percent said no. Twenty percent said yes outright. Sixty-five percent said yes, but likely with a new model for authorization. Ten percent were unsure.

The combined 85 percent who anticipate AI-driven autonomous payments represents an industry that sees fundamental change coming to the authorization and authentication models that have governed financial transactions for decades. The current card payment system — built on the assumption that a human cardholder initiates, reviews, and approves each transaction — is not designed for an environment in which an AI agent acting on a consumer's behalf negotiates terms, selects merchants, and executes purchases autonomously. The 65 percent who specify "a new model for authorization" recognize that the existing infrastructure must evolve.

This evolution is already underway and moving rapidly. Visa launched its Intelligent Commerce initiative in October 2025 with the Trusted Agent Protocol, an open framework for distinguishing legitimate AI agents from malicious bots, and by December 2025 had completed hundreds of secure agent-initiated transactions in controlled environments with over 100 partners [14]. Mastercard unveiled Agent Pay in April 2025 using enhanced tokenization technology, completed rollout to all U.S. cardholders by the 2025 holiday season, and in early 2026 partnered with Banco Santander to execute Europe's first live end-to-end payment by an AI agent within a regulated banking framework [15]. Stripe launched its Machine Payments Protocol in March 2026, enabling AI agents to make programmatic payments including stablecoins and fiat currency [16]. Google announced its Agent Payments Protocol (AP2) in September 2025 with 60 collaborating organizations including Adyen, American Express, PayPal, and Revolut [17]. Edgar, Dunn & Co. projects the agentic commerce total addressable market at $135 billion in 2025, growing to $1.7 trillion by 2030 [18].

The security implications are significant and largely unresolved. Eighty-seven percent of financial institution CTOs and heads of payments surveyed believe trust will be the most significant barrier to agentic payments adoption, and 78 percent expect fraud will increase significantly due to agentic commerce [19]. Agent-initiated transactions require new approaches to identity verification, transaction authorization, spending limits, liability allocation, and fraud detection that account for the delegation chain from human consumer to AI agent to merchant. Critically, agentic commerce removes many signals that merchants currently rely upon for fraud detection — customer IP addresses, browsing patterns, and behavioral biometrics — meaning that legitimate agent transactions may resemble the very bot activity that existing fraud systems are designed to block.

AI Security Risks: Data Leakage Dominates the Threat Landscape

When respondents were asked to identify the three greatest security risks when using AI in their organization, the results revealed a threat landscape dominated by data protection concerns. Sensitive data leakage via prompts, files, or chat history — whether driven by employees or automated workflows — was cited by 61 percent of respondents, making it the single most frequently identified AI risk by a significant margin. 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.

The second-tier risks center on identity, access, and data exposure through AI architecture. Excessive permissions or weak authorization for AI agents and tools was 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.

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.

External threat amplification rounds out the risk picture. 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: A Growing but Opaque Problem

Twenty percent of respondents report that their organization has experienced security incidents or breaches related to AI or ML systems, including data leakage via AI services, model compromise, and similar events. This figure is striking on its own, but the full picture is more concerning: an additional 21 percent responded "unsure," indicating that a significant fraction of financial institutions lack the monitoring, detection, or reporting capabilities to determine whether AI-related security incidents have occurred.

The 21 percent of respondents who are unsure whether their organization has experienced AI-related security incidents may indicate gaps in AI security monitoring and detection capabilities, though respondent uncertainty could also reflect organizational silos or the early stage of many AI deployments. Taken together with the 20 percent confirmed incident rate, the data points to AI security observability as an area requiring further investment. Financial institutions that cannot detect AI-related incidents cannot learn from them, report them to regulators, or implement controls to prevent recurrence. This observability gap is compounded by the novelty of AI-specific attack vectors: traditional security monitoring tools were not designed to detect prompt injection, training data poisoning, or unauthorized data exposure through generative AI interactions.

AI Adoption Barriers: Privacy, Skills, and Trust

The barriers to AI adoption in financial services mirror the broader security concerns but add dimensions specific to AI technology. Concerns about data privacy or security when using AI lead at 54 percent, confirming that the data leakage fears identified in the risk question translate directly into adoption hesitancy. Lack of AI/ML expertise is cited by 45 percent, a workforce challenge that compounds the existing cloud security skills gap. Lack of trust in AI outputs, encompassing accuracy, drift, bias, and unexplainable decisions, affects 45 percent of respondents — a finding with particular significance for financial services, where model outputs may need to be explained to regulators, customers, and courts.

Unclear regulatory requirements for AI slow adoption for 36 percent, reflecting the fragmented and evolving global regulatory landscape. Unclear ROI for AI projects affects 29 percent, budget constraints 25 percent, and integration difficulties with legacy systems 22 percent. Organizational resistance, including lack of executive buy-in for AI, was cited by only 5 percent, suggesting that leadership is generally supportive of AI adoption even as organizations struggle with the practical challenges of implementation.

The trust gap at 45 percent has profound implications for regulated financial services. When AI systems make credit decisions, flag transactions as fraudulent, or assess insurance risk, both regulators and consumers expect those decisions to be explainable and auditable. The EU AI Act classifies credit scoring and insurance underwriting as high-risk AI applications requiring conformity assessments, human oversight,

and detailed technical documentation, with full compliance required by August 2, 2026 [28]. Financial institutions that cannot demonstrate the reliability and explainability of their AI outputs face both regulatory sanction and reputational risk.

Leadership Support and Organizational Maturity

Senior leadership support for cloud security initiatives is generally positive, with 43 percent of respondents describing it as strong and 38 percent as moderate. Limited support is reported by 13 percent, and no meaningful support by 3 percent, with the remaining 3 percent unsure. These figures suggest that executive awareness of cloud security as a business imperative has continued to improve, though the 16 percent reporting limited or no support represent organizations at elevated risk as they navigate the combined challenges of cloud and AI security.

The open-ended comments from respondents provide qualitative texture to the quantitative data. Several themes emerge repeatedly: the 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 experimentation; the tension between innovation speed and compliance requirements; and the emerging problem of agentic identity management, where AI agents require credentials, permissions, and authorization that existing IAM systems were not designed to support.

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." 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." This centralized AI gateway approach — channeling all LLM usage through a monitored, policy-enforced interface rather than attempting to prohibit usage outright — is emerging as the dominant security architecture pattern for financial services AI governance, mirroring the evolution from device prohibition to mobile device management in the early 2010s.

Analysis: Five Defining Trends

1. The AI Security Integration Imperative

The convergence of AI security integration as a top-three priority (44 percent) with the high rate of AI incidents (20 percent confirmed, 21 percent uncertain) and the dominance of data leakage as the leading AI risk (61 percent) tells a coherent story: financial institutions have deployed AI faster than they have secured it, and they know it.

The traditional cloud security stack — CSPM, CWPP, CNAPP, SIEM — was designed for a world of virtual machines, containers, and APIs. It does not inherently address the novel risk surface created by AI systems that ingest sensitive data through prompts, retrieve it through RAG pipelines, store it in vector databases, and generate outputs that may contain residual training data. The emergence of AI-specific security categories in respondent write-ins — AISPM, DSPM for AI, and shadow AI detection — signals a market that is beginning to address this gap.

For financial services specifically, the AI security integration challenge is compounded by regulatory expectations. Model risk management, traditionally governed by frameworks such as the Federal Reserve's SR 11-7 guidance on model risk management, must now accommodate AI systems that are fundamentally different from the statistical models these frameworks were designed for [29]. The American Bankers Association has called on prudential agencies to update SR 11-7 to address AI-specific applicability, but no substantive update has been issued in 15 years. Large language models cannot be validated through the same approaches used for logistic regression or decision trees. Their behavior is probabilistic, context-dependent, and difficult to characterize exhaustively. Financial regulators worldwide are grappling with how to update model risk management expectations for this new class of system, and the industry is grappling with how to comply with expectations that are still forming.

2. Third-Party Risk in the Age of AI

The positioning of third-party and supply chain risk as the top cloud security concern (55 percent) represents a significant shift from prior surveys. In 2023, while third-party risk was acknowledged, it ranked behind cloud misconfigurations and data privacy as a primary concern. Several factors have elevated it to the top position.

The CrowdStrike incident of July 19, 2024, in which a faulty content update to the Falcon sensor caused an estimated 8.5 million Windows devices to crash worldwide, demonstrated the operational risks of concentrated technology dependencies in financial services. Banking sector losses alone were estimated at $1.15 billion, with JPMorgan Chase, Bank of America, and Wells Fargo among the affected institutions [3]. The incident validated years of regulatory warnings about concentration risk. It was followed by additional

cloud disruptions in 2025: a 15-hour AWS outage in October 2025 affecting over 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 [20]. Between August 2024 and August 2025, AWS, Azure, and Google Cloud experienced more than 100 combined service outages, establishing cloud disruption as a recurring operational reality rather than an exceptional event [20].

The ransomware threat to financial services has also intensified. Sixty-five percent of financial services IT professionals report being hit by ransomware in 2024, above the 59 percent all-industry average, with recovery costs averaging $2.58 million per incident [21]. The November 2023 LockBit attack on ICBC Financial Services — the U.S. subsidiary of the world's largest bank by assets — disrupted U.S. Treasury trade clearing and repo financing in the $26 trillion Treasury market, forcing the bank to transmit settlement details via USB stick by messenger in Manhattan [22]. The Infosys McCamish breach, also via LockBit, exposed data of 6.5 million individuals including Bank of America customers [23]. These incidents underscore that third-party risk is not abstract — it manifests as concrete operational disruption and data exposure at scale.

DORA, which became enforceable across approximately 22,000 EU financial entities on January 17, 2025, codifies third-party ICT risk management requirements that go beyond anything previously mandated. It requires financial entities to maintain a register of all ICT third-party arrangements, conduct thorough risk assessments of critical ICT third-party service providers, ensure contractual provisions for exit strategies and substitutability, and participate in threat-led penetration testing at least every three years on live production systems. In November 2025, the European Supervisory Authorities published their first list of 19 designated Critical ICT Third-Party Providers, including AWS, Google Cloud, Microsoft, Oracle, SAP, and Bloomberg, subjecting these firms to direct EU supervisory oversight with penalties of up to 1 percent of average daily worldwide turnover [4][24]. The regulation also requires financial entities to assess ICT concentration risk before contracting with any provider, and existing contracts must comply by the earlier of next renewal or July 1, 2026 [4].

The AI dimension adds a new layer of third-party risk. When a financial institution uses a cloud-hosted AI service for credit scoring, customer interaction, or fraud detection, it is dependent on a third party not only for infrastructure availability but also for the behavior, accuracy, and security of the AI model itself. Model updates, training data changes, capability modifications, and security vulnerabilities in the AI service all represent risks that the financial institution may have limited ability to detect or control. The 10 percent of respondents who cited third-party or supply chain risk from AI providers as a top AI risk may understate the actual exposure, as many organizations may not yet have visibility into the full chain of AI dependencies in their environment.

3. The Coming Wave of Agentic Finance

The survey data on AI agents paints a picture of an industry moving toward a fundamentally new operating model. With 62 percent already deploying agents, 85 percent anticipating autonomous AI payments, and 38 percent of agent-using organizations granting conditional or high autonomy, financial services is approaching a threshold where AI systems are not merely tools that humans use but actors that operate within financial systems on behalf of humans and institutions.

This transition raises questions that existing governance frameworks were not designed to answer. Existing law provides a clear answer to the most fundamental question: regulatory liability remains with the regulated institution. Delegation of a function to an AI agent does not transfer liability to the AI provider, the model developer, or the agent itself. What is not yet clear is the operational governance framework that must accompany this liability. How should agent identity be established, verified, and audited? What authorization model governs an AI agent that acts on behalf of a customer who is not present for the transaction? How should an institution's risk appetite be translated into machine-enforceable guardrails? Supervisory authorities will want to understand whether institutions deploying high-autonomy AI agents have subjected those systems to appropriate model risk management, whether boards have explicitly approved the risk appetite for autonomous AI decision-making, and whether the institutions have notified their supervisors as may be required under applicable material change notification requirements.

CSA's research on agentic AI identity and access management offers a framework for addressing some of these questions. The core principle is that agent identity must be treated as a first-class concept in IAM systems, with agents receiving verifiable credentials, time-bounded permissions, and auditable action logs that are distinct from but linked to the human or organizational principal on whose behalf they act [6]. The MAESTRO framework provides a multi-layer security model for agentic AI systems that addresses agent identity, tool authorization, inter-agent communication, and human oversight at each level of an agentic architecture [30].

The payment networks' response to agentic commerce — Visa's Intelligent Commerce initiative and similar programs — suggests that the infrastructure for AI-mediated transactions is being built now, even as the security and governance frameworks lag behind. The 65 percent of respondents who said agentic payments will require a new authorization model are correct, and the financial services industry has a narrow window to develop that model before the technology outpaces the controls.

4. Regulatory Convergence and Persistent Divergence

The global regulatory environment for cloud and AI in financial services is characterized by a paradox: there is broad consensus on the principles (third-party risk management, operational resilience, AI transparency, data protection) but persistent divergence on the specifics of implementation. Financial institutions

operating across jurisdictions face a patchwork of requirements that are directionally aligned but tactically incompatible.

In Europe, DORA and the EU AI Act form the most comprehensive regulatory framework for digital resilience and AI governance in the world. DORA's requirements for ICT risk management, incident reporting, and third-party oversight are unprecedented in scope and specificity. The EU AI Act entered into force on August 1, 2024, with a phased implementation: prohibited AI practices applied from February 2, 2025; general-purpose AI model obligations from August 2, 2025; and the high-risk AI system obligations — which classify financial AI applications including credit scoring and insurance risk assessment — take full effect on August 2, 2026. Financial institutions reading this report in 2026 face an imminent compliance deadline. High-risk obligations include conformity assessments, human oversight, fundamental rights impact assessments, and detailed technical documentation, with penalties of up to EUR 35 million or 7 percent of global annual turnover [28]. National financial supervisory authorities are designated as market surveillance authorities for high-risk AI deployed by regulated institutions, meaning the same regulators who examine banks for safety and soundness will also assess their AI systems for EU AI Act compliance.

In the United States, the regulatory approach remains principles-based and distributed across multiple agencies. The Federal Reserve's SR 11-7 guidance on model risk management, while predating modern AI, provides the foundation for supervisory expectations on AI models. The OCC, FDIC, and Federal Reserve have issued interagency guidance on third-party risk management, and several agencies have signaled interest in AI-specific guidance without yet promulgating binding rules. The NIST AI Risk Management Framework provides a voluntary structure that many U.S. financial institutions have adopted (39 percent in this survey) as a de facto governance standard.

In Asia Pacific, the Monetary Authority of Singapore published a consultation paper on AI Risk Management Guidelines in November 2025, establishing supervisory expectations for AI governance including board oversight, AI inventories, risk materiality assessments, and lifecycle controls spanning data management, fairness, transparency, and third-party risk [36]. Australia's APRA implemented CPS 230 on Operational Risk Management effective July 2025, which requires comprehensive service provider management policies and business continuity planning that extend to cloud and AI governance [31]. The Basel Committee published its Principles for the Sound Management of Third-Party Risk in December 2025, establishing 12 high-level principles that create a common international baseline for managing cloud, AI, and other technology dependencies [37].

The survey data reflects this complexity. Organizations average three or more frameworks, and the "Other" responses reveal additional regulatory requirements from DORA, the EU AI Act, PCI DSS, SOX, and various national authorities. The compliance burden is not merely administrative — it consumes security resources that could otherwise be directed toward improving actual security outcomes.

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

The survey data repeatedly surfaces concerns — data leakage at 61 percent, RAG exfiltration at 27 percent, credential exposure at 19 percent — that point to a technical security gap in the AI data pipeline that traditional cloud security tools do not address. Understanding this gap requires examining the specific architectural components of enterprise AI deployments and the novel attack surfaces they introduce.

The most common enterprise AI architecture in financial services involves a retrieval-augmented generation (RAG) pipeline in which user queries are converted to vector embeddings, matched against a vector database containing enterprise documents, and the retrieved content is assembled into a prompt that is sent to a large language model for completion. Each stage introduces distinct security risks. At the ingestion layer, malicious or poisoned content in source documents can persist in the vector store and be retrieved in response to future queries — a form of persistent indirect prompt injection. At the retrieval layer, vector similarity search does not inherently respect document-level access controls; a user with access to the RAG system may receive content from documents they are not authorized to view, because the retrieval mechanism operates on mathematical similarity rather than authorization boundaries. At the generation layer, the model may incorporate sensitive data from retrieved documents into its output, expose patterns learned during training, or be manipulated through adversarial prompts embedded in the retrieved content.

Vector databases themselves represent a new and largely unsecured component of the AI data infrastructure. Most commercial vector stores lack the row-level security, field-level encryption, and audit logging capabilities that financial institutions expect from traditional databases. Embedding vectors can, under certain conditions, be reverse-engineered to recover information about the source text, creating an information leakage channel that does not exist in conventional data architectures.

For organizations deploying AI agents, the Model Context Protocol (MCP) and similar tool-use architectures introduce additional security concerns. The relatively low survey citation rate for tool and protocol security risks (9 percent) likely reflects nascent awareness rather than low actual risk. As AI agents increasingly use tools to interact with financial systems — executing API calls, querying databases, triggering workflows — the security of tool authorization, input validation, and execution sandboxing becomes critical. Early research has identified tool poisoning (malicious tool descriptions that cause agents to misuse tools), excessive tool permissions (agents granted broad API access when they need narrow scopes), and injection via tool descriptions as significant risks in agentic architectures.

The technical controls needed to address these risks are emerging but not yet widely deployed in financial services. Input and output guardrails that scan prompts for sensitive data before submission and filter completions for PII, credentials, and proprietary information represent the first line of defense. Access control enforcement at the retrieval layer, so that RAG systems respect document-level permissions rather than relying solely on user-level authentication, is essential for regulated data. Model behavioral baselines with monitoring for drift, anomalous outputs, and potential extraction attempts provide ongoing assurance. Purpose-built AI security platforms — including AI firewalls, LLM gateways, and AI security posture

management tools — are now commercially available and beginning to see adoption at financial institutions that recognize the limitations of traditional cloud security stacks for AI workloads. Cloud providers have also begun offering native AI security controls, including Amazon Bedrock Guardrails, Azure AI Content Safety, and Google Cloud's model monitoring capabilities.

Confidential computing represents a particularly promising approach for financial services AI workloads. Hardware-attested trusted execution environments such as Intel TDX, AMD SEV-SNP, and AWS Nitro Enclaves enable AI inference on sensitive financial data while providing cryptographic assurance that neither the cloud provider nor other tenants can observe the data or model in use. For institutions where the 54 percent data privacy concern is the primary barrier to AI adoption, confidential computing offers a technical path to deploying AI on regulated data without exposing that data to the infrastructure provider.

Additionally, the paper would be incomplete without noting the distinction between data leakage as a confidentiality concern and output integrity as a reliability and compliance concern. The survey foregrounds confidentiality risks, but in practice, some of the most damaging AI incidents in financial services involve incorrect AI-generated outputs that are acted upon without verification — hallucinated compliance determinations, inaccurate risk assessments, or credit decisions that inadvertently violate fair lending requirements. Both the confidentiality and integrity dimensions of AI output require distinct controls and monitoring approaches.

6. Systemic Risk: Beyond the Individual Institution

The analysis throughout this report has operated primarily at the firm level — what risks does an individual institution face from AI adoption? But the survey findings also carry systemic implications that warrant attention from supervisory authorities and financial stability bodies.

If multiple financial institutions use the same foundation models for similar functions without significant customization, their AI outputs could become correlated. In such a scenario, a systematic bias, hallucination pattern, or failure in a widely-used model might produce simultaneous errors across institutions. This is a form of model risk that SR 11-7's institution-level framework was not designed to address, and that the Financial Stability Board has flagged as a financial stability concern in its October 2025 report on AI adoption, which identifies third-party dependencies and service provider concentration as key vulnerabilities [40].

The concentration risk analysis extends beyond cloud providers to AI model providers. Three hyperscalers control approximately 63 percent of global cloud infrastructure, but the concentration in foundation model providers is even more acute. If the AI capabilities of the global banking system depend on a handful of model providers, the compromise or failure of one becomes a systemic event. The 19 Critical ICT Third-

Party Providers designated under DORA do not yet include AI model providers as a distinct category, though supervisory authorities may need to consider whether this oversight framework should extend to firms whose models underpin critical financial functions [24].

AI-driven trading and risk management systems that learn from the same market data may amplify market movements by taking similar positions simultaneously — the AI analogue of quantitative strategy crowding, but potentially at much greater scale and speed. Similarly, institutions using similar AI models for credit scoring may simultaneously tighten or loosen credit to the same borrower segments, amplifying the credit cycle in ways that individual model validation would not detect.

The 5 percent of responding organizations that have granted high autonomy to AI agents for critical actions represents a systemic risk finding, not merely a firm-level governance observation. Autonomous AI agents operating across interconnected financial networks could propagate errors or malicious actions at machine speed, faster than human circuit-breakers can intervene. This dimension of AI risk requires macro-prudential analysis that extends beyond the scope of this survey but that the survey data underscores as urgently needed.

7. From Cloud Security to AI-Cloud Security

The final trend is perhaps the most fundamental: the emergence of AI-cloud security as a unified discipline. The traditional separation between cloud security (protecting infrastructure, data, and applications in cloud environments) and AI security (protecting AI models, training data, and AI-driven decisions) is breaking down as AI becomes deeply embedded in cloud-hosted financial services.

This convergence is evident in the survey data. The top cloud security risk (third-party/supply chain at 55 percent) and the top AI security risk (data leakage via AI at 61 percent) are not separate problems but manifestations of the same underlying challenge: financial institutions are processing sensitive data through increasingly complex chains of technology providers and AI services, and the controls designed for simpler architectures are insufficient.

The response requires an integrated approach that treats AI workloads as first-class citizens in cloud security architecture. The CSA AI Controls Matrix (AICM) represents one such approach, extending the Cloud Controls Matrix with controls specific to AI data governance, model lifecycle management, AI-specific identity and access management, and AI operational resilience [5]. The 24 percent adoption rate of AICM in this survey, while still minority, suggests growing recognition that AI and cloud security must be governed together rather than in separate silos.

Opportunities and Recommendations

The survey findings point to several actionable opportunities for financial institutions, cloud service providers, regulators, and the broader security community.

For Financial Institutions

Financial institutions should prioritize establishing unified AI-cloud security governance that integrates AI risk management into existing cloud security programs rather than treating it as a separate initiative. The data classification gap identified by 26 percent of respondents as a barrier to security implementation is a prerequisite for effective AI governance — organizations cannot control what AI systems access if they do not know where their sensitive data resides and how it is classified.

Identity and access management for AI agents and non-human identities should be a near-term priority for organizations deploying AI agents. With 62 percent of organizations deploying AI agents and non-human identity risk ranking as a top-five cloud security concern, financial institutions need to implement agent-specific IAM capabilities including verifiable agent credentials, scoped permissions, behavioral monitoring, and automated credential rotation.

Organizations should conduct AI-specific threat modeling that goes beyond traditional application security. The OWASP Top 10 for AI Applications, MITRE ATLAS framework, and CSA's MAESTRO framework provide structured starting points [30][32]. In practical terms, financial institutions should prioritize four immediate technical investments: (1) deploy input/output guardrails that scan prompts for sensitive data before submission and filter completions for PII, credentials, and proprietary information; (2) implement access control enforcement at the RAG retrieval layer so that AI systems respect document-level permissions, not just user-level authentication; (3) conduct tool authorization reviews for all MCP-connected tools to ensure agents operate under least-privilege principles with scoped, time-bounded permissions; and (4) establish model behavioral baselines and monitor for drift, anomalous outputs, and potential extraction attempts.

Exit and contingency planning for critical AI services should be integrated into third-party risk management programs. The survey data shows that 37 percent of organizations have introduced exit or contingency plans for cloud providers, but this planning needs to extend to AI model providers, AI platform services, and the data pipelines that feed them. Fourth-party risk — the chain of dependencies behind an AI service provider, including the foundation model vendor, training data sources, and evaluation services — requires explicit assessment.

AI security governance should be formalized with clear organizational ownership. At minimum, this requires: a comprehensive inventory of all AI systems and models in use across the organization (including shadow AI); a risk classification scheme that maps AI use cases to appropriate autonomy levels and control

requirements; board-reportable metrics for AI security posture including incident rates, guardrail trigger rates, and compliance status; and a defined RACI model for AI security spanning the CISO, Chief Data Officer, model risk management, and business unit AI teams.

For Cloud Service Providers

Cloud service providers serving financial institutions should invest in transparency around AI service changes, model updates, and security events. The 2023 survey highlighted that CISOs were frustrated by less than 48 hours' notice of significant changes from cloud providers. For AI services, where model behavior can change with each update, the transparency requirement is even more acute.

Providers should develop financial-services-specific AI governance tools that support regulatory requirements for model explainability, audit trails, and data lineage. The 54 percent of respondents who cite data privacy and security concerns as the top barrier to AI adoption represent a market that will pay for demonstrable security assurance.

For Regulators

Regulatory clarity on AI governance expectations would benefit both supervised entities and supervisory authorities. Thirty-six percent of respondents cite unclear regulatory requirements as a barrier to AI adoption, creating compliance uncertainty that can either slow beneficial adoption or, conversely, lead institutions to deploy AI systems without adequate governance on the basis that specific requirements do not yet exist. Supervisory authorities may wish to consider whether existing frameworks such as SR 11-7 require supplementary guidance to address AI/ML systems, and whether AI-specific examination procedures should be developed to ensure consistent supervisory assessment.

Several specific actions would advance supervisory effectiveness. First, supervisory authorities should develop AI-specific examination procedures and invest in training examiners in AI/ML model validation, as the current gap between the models being deployed and examiners' ability to assess them creates supervisory risk. Second, regulators should consider requiring financial institutions to maintain inventories of AI systems deployed in material functions, analogous to DORA's register of ICT third-party arrangements. Third, international supervisory coordination on AI risk — through the FSB, Basel Committee, or bilateral agreements — should be prioritized to reduce the jurisdictional fragmentation that the survey data reveals as a real compliance burden. Fourth, the concentration risk framework developed for critical ICT third-party providers under DORA could be extended to AI model providers whose foundation models underpin critical financial functions across multiple institutions. Fifth, supervisory stress testing and scenario analysis should incorporate AI failure scenarios, including correlated model failures across institutions using the same foundation models.

Appendix: About CSA's Financial Services Program

The Cloud Security Alliance's Financial Services Working Group serves as the primary venue for industry collaboration on cloud and AI security standards, regulatory engagement, and best practices development for financial services. Several CSA initiatives are relevant to the findings in this report:

The **AI Controls Matrix (AICM)**, adopted by 24 percent of survey respondents, extends the Cloud Controls Matrix to address AI-specific governance, risk, and compliance requirements. The **MAESTRO framework** for multi-agent security provides architectural guidance for organizations deploying AI agents. The **STAR program** enables cloud service providers to demonstrate security assurance, with alignment work underway for DORA requirements and PCI DSS v4.0. CSA's research on **agentic AI identity and access management** addresses the agent credential and authorization challenges identified throughout this report.

Priority areas for future work identified through this survey include AI agent identity and authorization models, AI-specific incident detection and response frameworks, cross-jurisdictional regulatory compliance mapping, and the development of financial-services-specific AI security training.

For more information, visit: https://cloudsecurityalliance.org/research/working-groups/financial-services/

Conclusion

The 2026 State of Cloud and AI for Financial Services survey captures an industry in the midst of a profound transformation. Cloud computing has become infrastructure — universal, essential, and no longer novel. The conversation has moved to artificial intelligence, and specifically to the question of how much autonomy financial institutions are willing to delegate to AI systems and whether their security and governance frameworks can support that delegation.

The data tells a story of an industry that is simultaneously optimistic and anxious. Ninety-eight percent are using AI, 62 percent have deployed AI agents, and 85 percent expect AI to autonomously conduct financial transactions. At the same time, 55 percent identify third-party risk as their greatest cloud security threat, 61 percent worry about data leakage through AI, 20 percent have already experienced AI-related security incidents, and 45 percent cite budget constraints as the primary barrier to improving security.

The resolution of this tension will define the financial services security landscape for the remainder of the decade. Organizations that invest in integrated AI-cloud security governance, workforce development, and the emerging frameworks designed to address agentic AI risk will be positioned to capture the productivity, efficiency, and competitive advantages that AI offers. Those that deploy AI without proportionate investment in security and governance face elevated risk of regulatory scrutiny, operational incidents, and erosion of customer trust.

The findings in this report should serve as both a benchmark and a roadmap. The benchmark is clear: the industry has adopted cloud, is rapidly adopting AI, and has identified but not yet resolved the security challenges at their intersection. The roadmap points toward four capabilities that financial institutions should prioritize in the near term: AI-specific security monitoring and observability tools that can detect prompt injection, data exfiltration, and model misuse in real time; identity and access management systems that treat AI agents as first-class principals with scoped, time-bounded, and auditable permissions; data classification and labeling infrastructure that enables policy enforcement at the AI ingestion layer; and guardrail technologies that enforce organizational policy on AI inputs and outputs without creating unacceptable latency for production workloads. The financial system's stability, the privacy of customer data, and the integrity of financial markets depend on making these investments before the next generation of autonomous AI capabilities outpaces the controls designed to govern them.

Demographics

The 2026 survey collected 340 responses from professionals across the global financial services industry. The demographic profile of respondents is summarized below.

Category	Segment	Percentage
Region	Americas	57%
	EMEA (Europe, Middle East, Africa)	48%
	APAC (Asia Pacific)	39%
Organization Size	<100 employees	17%
	101–1,000 employees	22%
	1,001–5,000 employees	19%

Category	Segment	Percentage
	5,001–10,000 employees	12%
	10,001+ employees	29%
Job Role	C-level or Executive	12%
	Director	26%
	Manager	38%
	Staff	22%
	Other	3%
Industry Segment	Banking / Credit Union	37%
	Fintech	19%
	Insurance	4%
	Investment Management	4%
	Professional Association	4%
	Cloud Service Provider	6%
	Regulatory Agency	2%
	Other	17%
Involvement in Cloud/AI Security	Yes	89%
	No	11%

The regional distribution reflects a more balanced global sample than the 2023 survey (52 percent Americas, 28 percent EMEA, 20 percent APAC), with increased participation from EMEA and APAC regions. The multi-select nature of the regional question reflects the multinational operations of many respondents.

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.