Capgemini

The agentic era is cloud's *moment of truth*

As agentic AI scales, those with the strongest foundations will see its true value.

World Report Series 2027
Cloud in Financial Services





Scan to explore an interactive version of this report

Table of contents

05
Foreword

06
Executive Steering Committee

08
Executive summary

10
Orchestrating agents across operations, technology and customer experience

22
Building a foundational cloud layer for trusted, compliant, and sovereign AI at scale

34
Elevating cloud-AI investments from consumption metrics to business outcomes

42
Conclusion

43
Methodology

44
Partner with Capgemini



Foreword

The financial services industry stands at a pivotal moment in its AI transformation. What began as targeted automation is now reshaping how financial institutions operate, compete, and serve customers. Cloud continues to be at the forefront of a catalyst for this change.

AI agents are moving beyond isolated tasks to coordinate decisions and actions across entire business lines, from underwriting and claims to advisory and compliance. Capgemini's World Cloud Report for Financial Services 2027 frames this shift as the industry's defining challenge and a strategic imperative: building the orchestration and governance foundations that turn scattered AI activity into enterprise-wide value. That value shows up as accelerating revenue growth, reducing operational cost, and differentiated customer experiences.

Our research shows that organizations are rapidly embracing this shift. Based on a survey of 1,100 executives across 14 markets and seven financial services sectors, 80% of firms began by ideating or piloting individual AI agents in 2025, and 59% now report piloting or deploying multi-agent systems in production within limited functions. 10% of firms believe they have reached enterprise-wide multi-agent orchestration capable of supporting autonomous business execution.

This report highlights three priorities for enterprise-wide agentic AI implementation. First, orchestrating AI agents across operations, technology, and customer experience to unlock enterprise-wide value. Second, establishing a foundational cloud layer that serves as the data, governance, and execution authority for trusted and compliant AI. Third, shifting the focus of cloud-AI investments from consumption metrics and experimentation to outcome-based value creation, supported by clear economic and governance models.

The gap between ambition and execution remains significant. Many organizations continue to operate fragmented pilots, while a small group of leaders is institutionalizing agentic AI as a core enterprise operating model. Those that successfully combine orchestration, cloud-governed execution, and value discipline will be best positioned to scale AI responsibly and achieve sustainable competitive advantage.

Ultimately, the next generation of financial services leaders will not be defined by how many AI agents they deploy, but by how effectively they orchestrate autonomous execution, govern it at scale, and convert AI activity into measurable business results.



Kartik Ramakrishnan
CEO of Capgemini's Financial Services Strategic Business Unit
Member of the Group Executive Board

Executive Steering Committee

The Executive Steering Committee for the World Cloud Report for Financial Services 2027 included top executives from across the financial services industry; committee members also represent the Americas, Europe, and Asia Pacific regions to ensure a mix of global perspectives and experiences. We are grateful for their time, experience, and vision in helping to guide our report's content.

Banking firms



Meghan Cavanaugh
USCB: Head of Engineering, Customer Journey Tribes
Barclays



Rafael Cavalcanti
Chief Data and Analytics Officer and CRM Director
Bradesco



Kasimir Hirn
Group CIO, EVP Development & Technologies
OP Pohjola



Sapna Sardana
Global Head of Planning & Stress Testing, and Internal Reporting Technology
Barclays



Hirenkumar Jadav
Head, Cloud CoE
Maybank

Executive Steering Committee

Insurance firms



Shreejit Nair
SVP, Chief Information Officer
Ameritas



Jeff Murr
VP, IT Infrastructure and Operations
Genworth



Robert Pick
Group Deputy CITO
Tokio Marine



Parimal Kumar
SVP, Digital, Services and Experience Tech Division
Erie Insurance



Shashi Reddy
Chief Digital, AI & Technology Officer
PEMCO

Executive summary

The agentic era is redefining how financial services operates

The financial services industry has entered a new stage of AI adoption. What began as isolated agentic AI deployments is rapidly evolving into enterprise-scale agent ecosystems that can coordinate workflows, automate decisions, and execute actions across business, technology, and customer functions.

The shift is already underway. In 2025, firms focused on deploying individual AI agents to automate specific tasks and improve efficiency. Today, leading organizations are building interconnected, multi-agent systems that work together across the enterprise. According to our 2026 survey of 1,100 financial services executives, 59% of firms are already piloting or deploying multi-agent systems, and 10% have achieved enterprise-wide orchestration that enables autonomous business execution at scale.

As agent ecosystems scale, the opportunity extends well beyond productivity gains. Multi-agent systems are beginning to transform core banking and insurance workflows, modernize technology operations, and reshape customer engagement. Financial institutions are focusing on high-volume workflows – such as customer service, compliance, underwriting, and claims management – where coordinated agents can execute tasks with greater speed, contextual awareness, and autonomy. At the same time, customer interactions are becoming increasingly agent-mediated, with 76% of executives expecting AI agents to become a key component of customer engagement within three years.

Despite this momentum, most organizations remain in the pilot stages of adoption: fragmented initiatives, siloed governance, and legacy operating models continue to limit enterprise-wide adoption. Only 5% of financial services firms have reached enterprise-wide agentic AI maturity, based on our assessment of industry players across five key dimensions: orchestration and deployment; governance, trust and autonomy; cloud execution authority; security and identity; and operating discipline and outcome economics.

Enterprise-wide agentic AI maturity will define the next generation of financial services leaders

Our research finds that leaders distinguish themselves not by AI adoption alone, but by their ability to institutionalize execution at scale. Compared with their competitors, leading organizations are significantly more likely to operate enterprise-wide, multi-agent orchestration, to embed governance directly into execution, and to institutionalize outcome accountability for AI investments.

Executive summary

These firms convert agentic AI maturity into enterprise advantage through three mutually reinforcing priorities:

1. Connect the bots

Orchestrating agents across operations, technology and customer experience.

- Leading firms are moving beyond individual use cases toward coordinated agents spanning operations, technology, and customer journeys.
- 44% of executives are prioritizing “agentic meshes” to bridge the gap between siloed pilots and enterprise execution.
- Organizations are also recalibrating their approach to sourcing AI capabilities: building what differentiates them, buying standardized capabilities, and partnering where ecosystem data adds depth and creates additional value.

2. Set the ground rules

Building a foundational cloud layer for trusted, compliant, and sovereign AI at scale

- Cloud is evolving beyond infrastructure into an execution authority for autonomous systems. It acts as the data spine, control plane, and governance fabric that enables trusted execution at scale.
- Regulatory requirements, data residency, and compliance certifications are the leading cloud-selection criteria, cited by 50% of executives, followed by security and identity management (47%) and auditability and traceability (39%).
- Sovereignty is expanding beyond data location to include who controls agent execution. Yet 62% of firms report that their identity and access-management infrastructure lacks the resilience needed for continuous agentic execution.

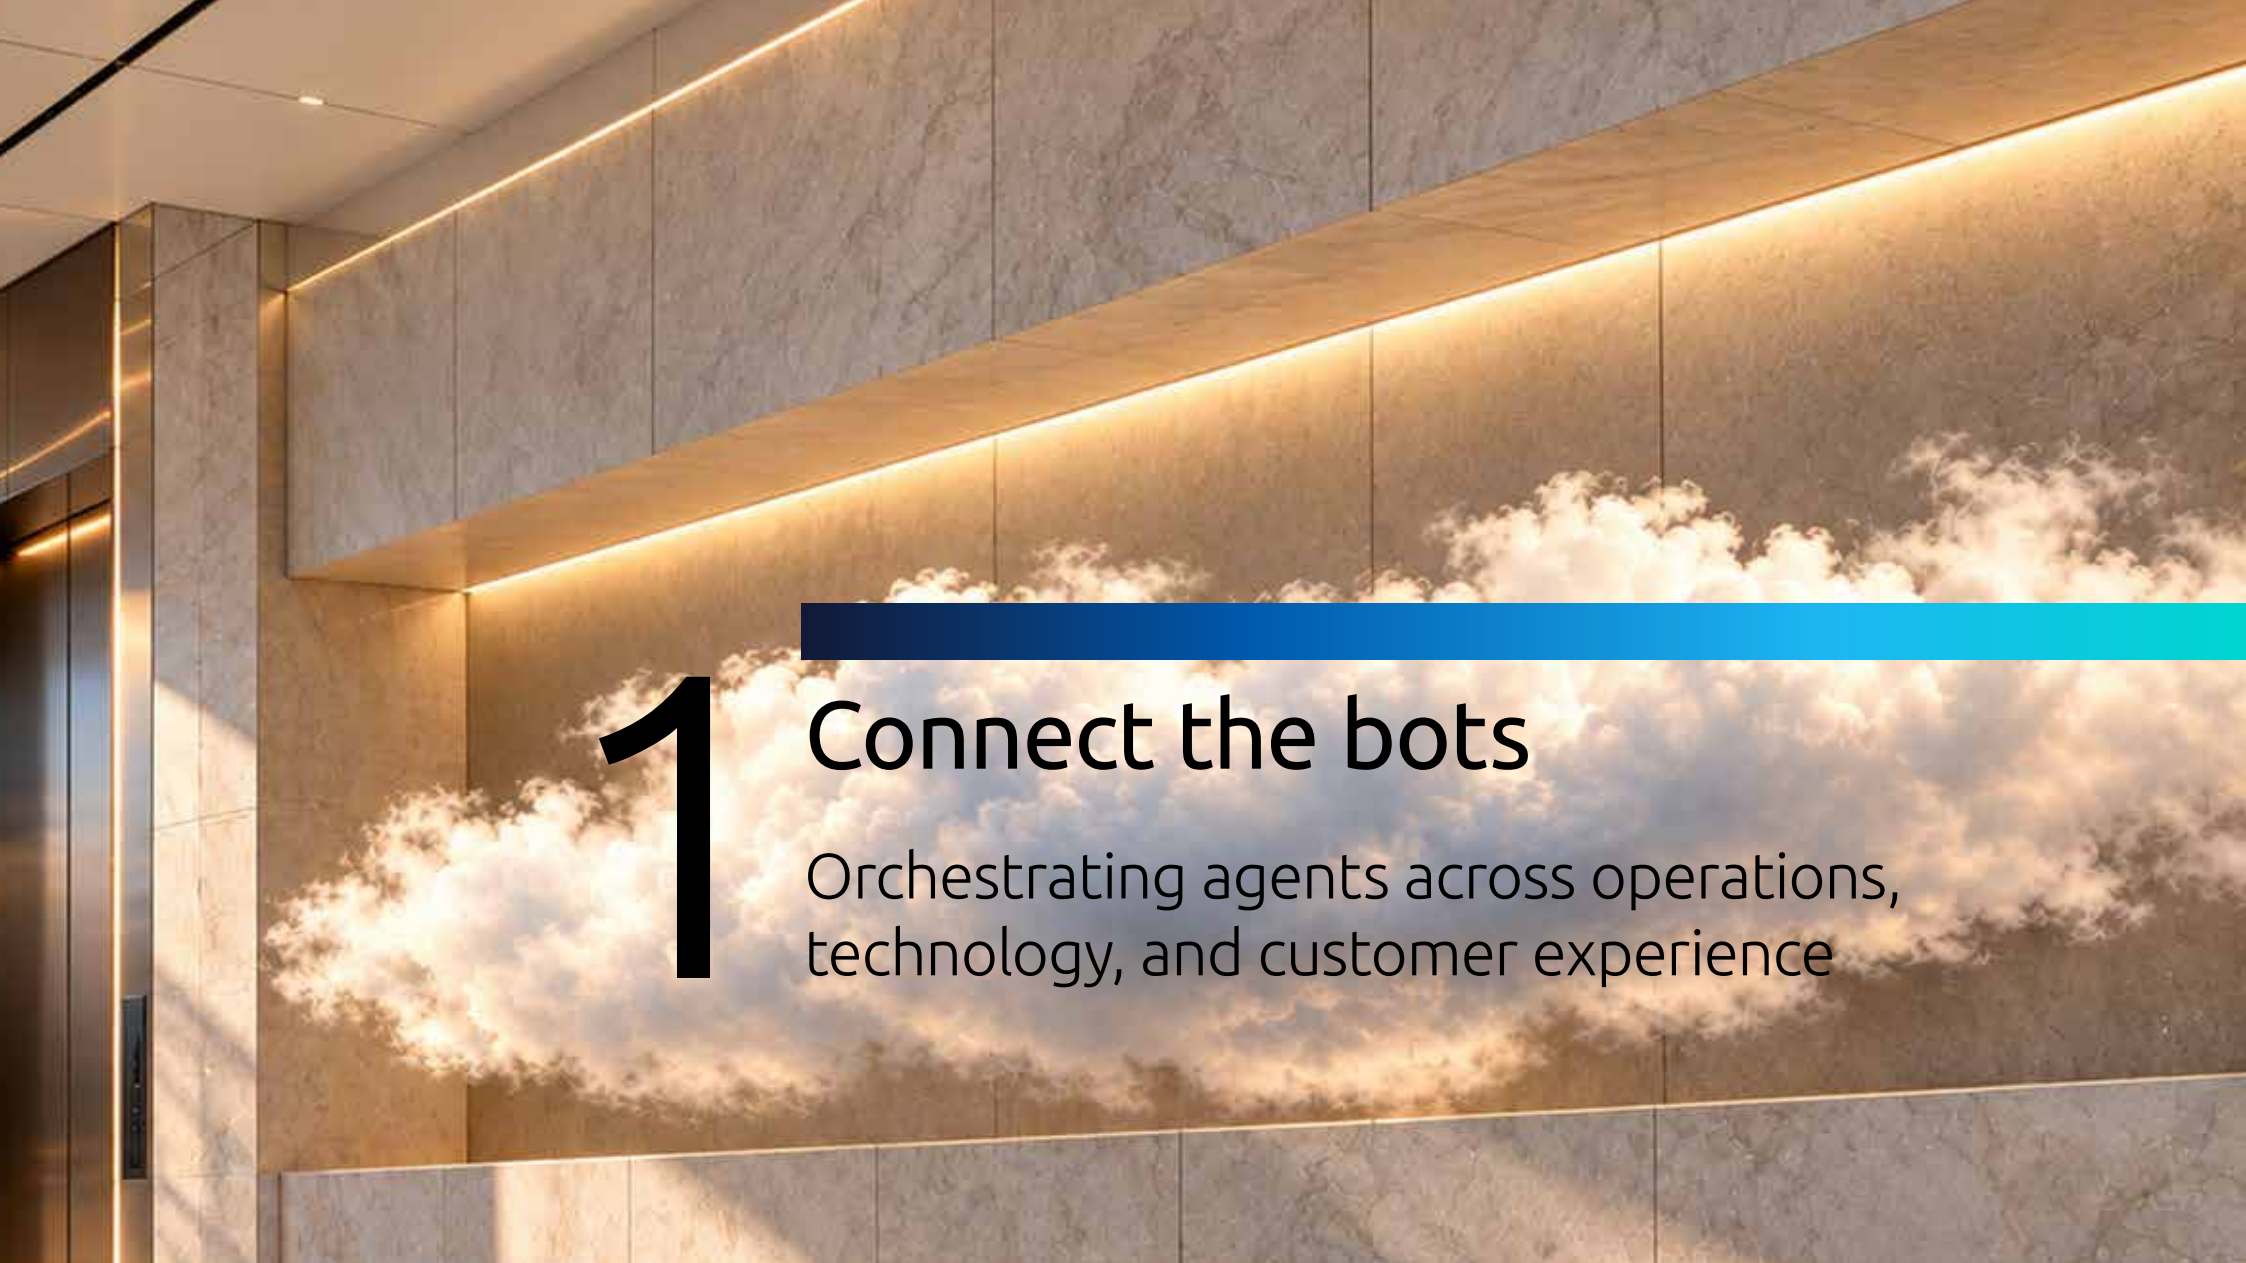
3. Unlock return on intelligence

Elevating cloud-AI investments from consumption metrics to business outcomes.

- As agentic workflows scale, organizations face growing pressure to demonstrate value and measurable returns.
- 48% of firms cannot consistently demonstrate ROI from agentic AI initiatives, while 51% lack mature value gates for governing continued investment.
- Leading organizations are connecting AI spending to measurable outcomes, using token and outcome economics and establishing clear accountability.

The gap between ambition and execution remains the defining challenge of the agentic era, and the next generation of financial services leaders will not be defined by the number of AI agents they deploy. While many firms continue to operate pilot-led initiatives, leaders are institutionalizing enterprise-wide orchestration, cloud-governed execution, security and trust controls, and outcome-driven investment disciplines.

Organizations that successfully combine these capabilities will be best positioned to scale agentic AI responsibly, realize sustainable business value, and establish long-term competitive advantage.



1 Connect the bots

Orchestrating agents across operations, technology, and customer experience

Agentic AI in financial services is at an inflection point. The industry is moving from deploying individual AI agents for isolated tasks to orchestrating connected, multi-agent systems – a shift that redefines how work gets done, where value is created, and how institutions compete. In 2025, financial services organizations began adopting single AI agents, with 80% of firms ideating or piloting them; these firms focused on single use cases and realized incremental efficiency gains and cost savings. Now, firms seek to scale multi-agent systems to enable organizational transformation, which are:

- Facilitating enterprise-wide agentic orchestration
- Transforming the IT function
- Redefining customer engagement.

Enterprise-wide agent ecosystems are the new operating model

Financial services firms are making steady progress toward connected multi-agent ecosystems, and they are unlocking significant business impact. Research conducted to support this report revealed that 59% of firms were piloting or deploying multi-agent systems in production within limited functions. These systems are helping organizations to achieve task-level automation, accelerate workflow execution, and improve operational efficiency. However, only 10% of firms have achieved enterprise-wide multi-agent orchestration capable of supporting autonomous business execution and organization-wide optimization. Orchestration is one of the five dimensions of our Enterprise Agentic AI Maturity Index, which tracks where financial organizations are in their AI journey.

Indeed, the lack of multi-agent orchestration represents the point where many firms struggle to move from pilot activity to enterprise-scale execution. To close this gap, 44% of executives are prioritizing agentic meshes, a decentralized infrastructure layer that connects autonomous agents and enables coordinated, enterprise-wide execution and governance.

10%

of firms have achieved enterprise-wide orchestration for autonomous business execution



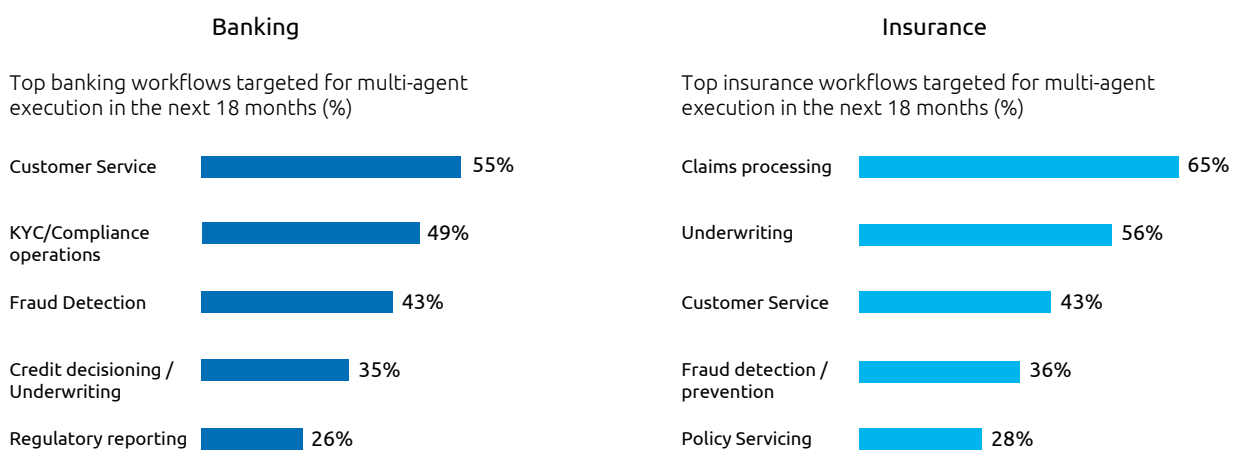
"The future lies beyond individual agents toward reusable capabilities that can be composed, orchestrated, and governed across the enterprise. Value compounds when organizations create common platforms that enable agents to collaborate at scale."

Rafael Cavalcanti
Chief Data and Analytics Officer and CRM Director, Bradesco

Early movers of enterprise agentic AI report strong returns on investment. For example, BNY launched an enterprise AI platform, Eliza, that enables employees to build task-specific agents using internal data. This AI initiative accelerated the client onboarding rate by 20% and reduced annual legal review time across 3,000+ vendor agreements by 75%.[1, 2]

As enterprises adopt agent ecosystems, they are shifting focus to high-volume, repeatable workflows that cut across functions and systems. Figure 1 highlights priority workflows by sector: banks are prioritizing operational workflows like customer service and compliance operations, while insurers are targeting claims processing and underwriting, where multi-agent systems can support context-rich analysis and complex decisions.

Figure 1:
Top banking and insurance workflows targeted for multi-agent execution in 18 months



Source: Capgemini Research Institute for Financial Services Analysis, 2026; Capgemini Cloud Survey of Financial Services Firms 2026 (N=1,100).



Multi-agent systems are transforming information technology

Orchestration of agentic systems depends on an agent-ready technology estate, requiring modernization of the underlying technology stack. Accordingly, IT and software development lifecycle (SDLC) transformation must progress in tandem.

Multi-agent systems are accelerating tech debt reduction in financial services. According to our executive survey, 80% of executives prioritize tech debt reduction through AI agents. Leading tech debt reduction priorities include:

- **API enablement of core systems:** Agents discover, generate, and maintain APIs around legacy applications and core platforms, making previously isolated systems accessible to modern applications, agent ecosystems, and cloud-native workflows.

80%

of executives prioritize tech debt reduction through AI agents



"One of our key findings is that code generation alone is not enough. Organizations should address the entire software development lifecycle, from initial ideas through production and ongoing operations."

Kasimir Hirn
Group CIO, EVP Development and Technologies, OP Pohjola

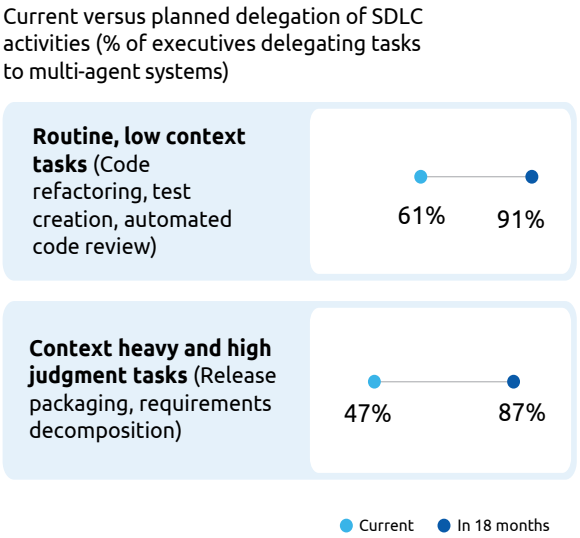
- **Data modernization:** Agents automate data discovery, cleansing, and transformation across legacy environments, improving data quality and accelerating access to trusted enterprise information.
- **Legacy codebase refactoring:** Agents read decades-old code, surface the business rules buried inside it, and restructure it into a maintainable form – a process that would otherwise consume significant engineering time.

One example of tech debt reduction through agentic AI is Morgan Stanley's in-house generative AI platform, DevGen.AI, which analyzed 9 million lines of legacy code, saving 280,000 developer hours.[3]

Beyond tech debt reduction, firms are extending multi-agent systems across the software development lifecycle. As illustrated in Figure 2, this trend is anticipated to grow strongly during the next 18 months:

- Adoption is strongest in routine, lower-context activities such as code refactoring, test creation, and automated code review, where agents can already assume delegated responsibility.
- More context-heavy activities, including release packaging and requirements decomposition, remain less mature today; however, executives expect these use cases to accelerate sharply, too.

Figure 2:
Banks and insurers are expanding multi-agent systems across the SDLC



Source: Capgemini Research Institute for Financial Services Analysis, 2026; Capgemini Cloud Survey of Financial Services Firms 2026 (N=1,100).



AI in action: US bank upgrades mainframe and legacy systems with Agentic and GenAI, reducing costs and efforts by up to 30%

Business challenge

A US financial services provider offering deposit, credit, investment, and wealth management products set out to modernize its mainframe and distributed legacy systems, and a straight lift-and-shift would not suffice. The bank wanted modern, modular architecture to improve agility, shorten time-to-market, and support new products for employees and customers alike. The scope was demanding: more than 700 components and ~1 million lines of decades-old code were tightly coupled to the core banking platform, and any changes had to keep the platform running without service disruption or compromise to data integrity.

Strategy/Implementation

The bank turned to Capgemini for an AI-led approach built on a suite of accelerators. First, the existing systems were analyzed to map how everything connected and to recover the business rules embedded in the legacy code, eliminating redundancy along the way. They then rebuilt those components as modern cloud services, organized around business domains rather than technical boundaries. Automated testing at each stage confirmed that the new services behaved exactly as the originals, allowing the bank to modernize without pausing delivery.

Results

Automation reduced modernization effort and cost by 20-30%, including labor, rework, and testing. Developer productivity rose by a similar margin as engineers spent less time deciphering legacy code. Defect rates and technical debt fell, core banking services ran uninterrupted throughout, and the move to the cloud enabled elastic scaling, high availability, and lower infrastructure costs. This real-world example illustrates how AI and engineering excellence together can solve complex, enterprise-scale challenges, making modernization more efficient and outcome-driven.



Agent-mediated customer journeys are redefining service delivery

Whether agents work inside the enterprise or face the customer, they rest on the same foundation - real-time decision engines, elastic compute, event-driven APIs and machine-verifiable identity. Together, these capabilities let an external agent authenticate, transact, and settle, on behalf of a customer. Firms that build this foundation turn customer-facing agents into a real channel for products and services.

In fact, according to our executive survey, 76% of executives expect AI agents to become a key part of customer engagement within three years; the customer journeys mentioned most frequently as expected to become agent-driven during this period are account servicing (51%), loan origination and underwriting (50%), and product recommendation (47%).

76%

of executives expect AI agents to become a key part of customer engagement within three years

Across the customer lifecycle, customer agents will increasingly act independently by:

1. Detecting financial needs and acting on them without customer prompting
2. Sourcing, comparing, and negotiating offers across multiple institutions on the customer's behalf
3. Executing transactions, managing identity verification and approvals, and modifying product choices in real time

This change will drive fundamental shifts as firms move toward business-to-agent distribution. In turn, this will propel three changes in financial services operations:

- New engagement layers centered around AI agents, shifting differentiation from user interfaces to personalized outcomes
- The movement of trust architecture to AI agents requiring verified agent identity and delegated authority
- A shift in decision-making toward real-time execution, with direct human involvement reserved for when high-stakes judgment is required

These shifts are also redefining how sourcing for agentic AI capabilities works. When agents were optimizing internal workflows only, buying off the shelf was a safe default. But now, the same agents are doing so much more – speaking to customers, acting under firm brand, and making decisions that once required human judgement. Buying capability by default is no longer safe, and building by default is no longer affordable. In the new world, firms no longer need one sourcing strategy for agentic AI – they need a sourcing decision for every individual capability.

Agentic AI is pushing firms toward a mixed sourcing strategy

As agentic AI moves from isolated pilots to enterprise execution, the key sourcing question is no longer how to build AI, but rather what intelligence is worth owning. Financial institutions are converging on a mixed sourcing model because no single approach can simultaneously deliver differentiation, trust, and speed. As a result, sourcing decisions are becoming a strategic governance question rather than a procurement decision.



Proprietary intelligence is increasingly built, commodity capabilities are bought, and ecosystem-driven workflows are partnered. According to our executive survey, today, 36% of firms use a mix of build, buy, and partner approaches; within the next three years, adoption of this mixed model is expected to rise to 69%. Firms are recalibrating their application strategies for agentic solutions across the three levers:

- **Build/generate for differentiation:** Build and own the development of agents that power differentiated workflows such as underwriting logic, risk decisioning, and proprietary customer intelligence, where control, IP, and competitive advantage matter most.
- **Buy for scale:** Procure packaged agentic solutions for standardized, non-differentiating workflows such as employee support, IT service management, and back-office operations, where speed to value outweighs the case for ownership.
- **Partner for ecosystem depth:** Co-develop agents with hyperscalers, integrated software vendors (ISVs), and specialized providers for shared workflows such as fraud detection and regulatory reporting, where ecosystem data and cross-industry signals can strengthen outcomes.

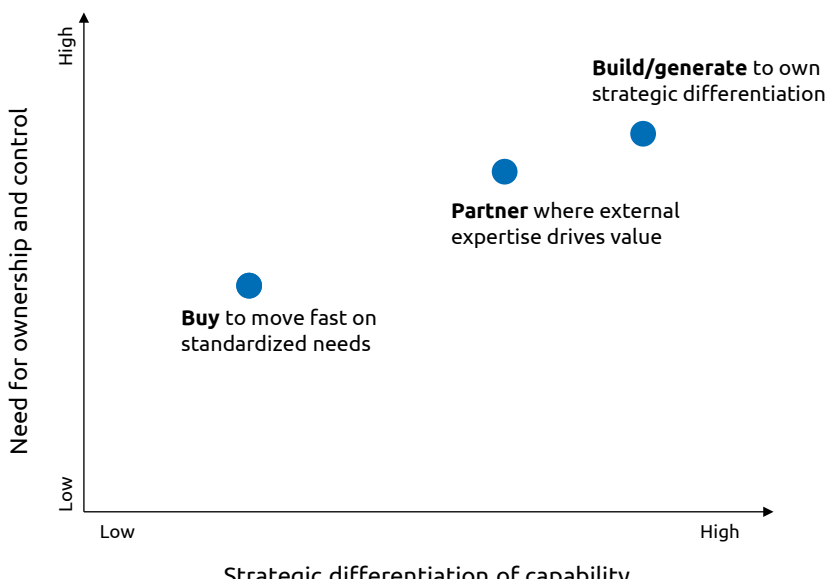


Figure 3 illustrates that strategic differentiation and control requirements are the two defining dimensions that help to determine sourcing model decisions. Workflows requiring high differentiation and tight control, such as proprietary underwriting or risk models, are best suited for in-house builds. Standardized, lower-control workflows favor procurement at scale. Shared workflows that benefit from ecosystem data and cross-industry collaboration sit in the partner zone, where co-development with hyperscalers and ISVs delivers depth that neither building nor buying alone can achieve.

"As agentic AI matures, the build, buy, or partner decision has become more nuanced. Most financial institutions are taking a hybrid approach. Where they land depends on where they see the greatest opportunity to differentiate."

Rianne van Veldhuizen, Vice President, Financial Services, AWS

Figure 3:
Differentiation and control requirements determine the optimal sourcing model for agentic solutions



Source: Capgemini Research Institute for Financial Services Analysis, 2026.

Building today also increasingly means generating, as generative development tools become the primary mechanism for custom agent creation, compressing timelines and shifting engineering effort from coding to orchestration. Across every application sourcing strategy, however, governance remains non-negotiable, requiring a cloud native approach for auditable execution at scale. Even though AI generation aims to solve the cost implications of building in-house, high token consumption and output reliability remain key concerns.

As organizations move toward autonomous operations and a mixed sourcing approach, a reimagined, foundational cloud layer becomes ever more critical.



Capgemini Research Institute 2026





2 Set the ground rules

Building a foundational cloud layer for trusted, compliant, and sovereign AI at scale

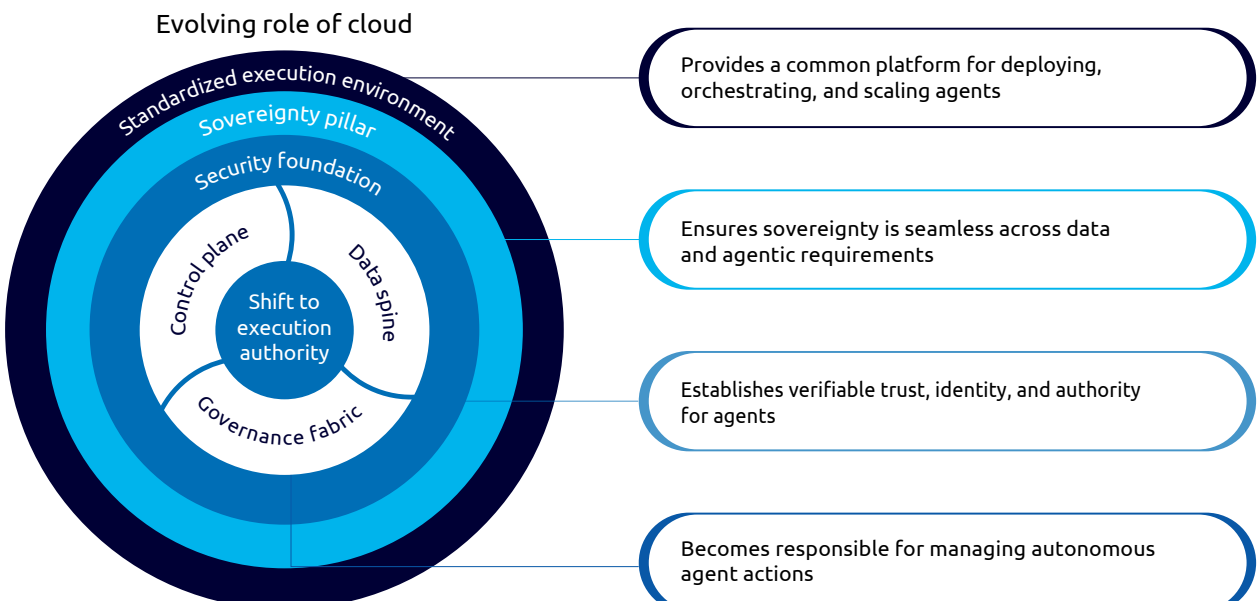
Pilots show what agents can do. Enterprise deployment demands clarity about what they are allowed to do: which data they can use, whose authority they carry, and how their decisions are recorded. In financial services, where autonomous actions may affect payments, portfolios and customers, those rules cannot be added after agents are deployed.

Cloud is becoming the common execution layer for applying those rules across systems. This chapter examines how governed data access, real-time controls and traceability support trusted action. Why agent identities require a new security model, and how sovereignty now reaches beyond data location to control over decisions and execution. Together, these foundations allow institutions to extend autonomy across the enterprise while retaining critical oversight and accountability.

Cloud is becoming the execution authority for autonomous financial services

Our research shows that financial institutions increasingly view cloud through three execution domains: how agents access data, how agent actions are controlled, and how autonomous decisions are governed. These requirements emerge as organizations scale agentic AI beyond pilots into enterprise-wide execution. Scaling multi-agent systems requires cloud to become the trusted execution foundation. As shown in Figure 4, its role expands from

Figure 4:
Cloud now plays execution authority, sovereignty pillar, and security foundation roles in financial services



Source: Capgemini Research Institute for Financial Services Analysis, 2026



"Cloud has become the practical enabler of AI experimentation and scaling. Ready access to models, infrastructure, and services reduces barriers to innovation while creating new imperatives around economics, governance, and long-term operating strategy."

Shashi Reddy
Chief Digital, AI & Technology Officer
PEMCO

being an innovation enabler to governing, securing, and orchestrating autonomous agent actions – an execution authority for autonomous financial services which is one of the dimensions in our Enterprise Agentic AI Maturity Index.

As autonomous execution grows, cloud capabilities further evolve and support become more layered:

- **Data spine:** The layer through which agents access data at runtime. Here, cloud provides governed, context-rich enterprise data for AI agents and enforces secure, purpose-bound, and consistent data access across agents. This in turn enables AI agents to deliver more accurate outcomes like risk and investment decisions.
- **Control plane:** The layer that manages agent identity, access, and policy enforcement. Here, cloud establishes the guardrails within which agents operate, enforcing real-time policies, approval thresholds, and execution boundaries before actions are taken. This allows firms to automate higher-value financial workflows while maintaining operational and risk controls.
- **Governance fabric:** The layer that makes autonomous execution traceable. Here, cloud tracks agent actions and decision lineage across multi-agent workflows and enables explainability, auditability, and regulatory compliance. This builds confidence in autonomous decisions across customer, risk, and compliance functions.

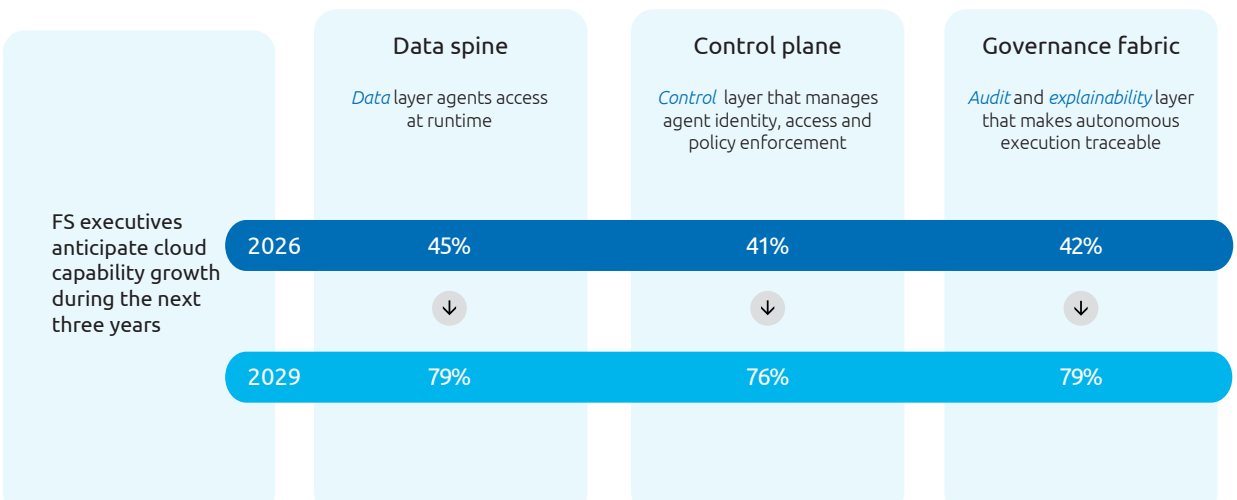


Across each of these layers, the executives we surveyed expect cloud execution capabilities to rise significantly in three years; Figure 5 highlights significant anticipated change across all three layers.

To illustrate this evolution, let's look at how cloud could operate across simple banking or insurance workflows.

As cloud evolves into an execution authority, agents gain direct access to data, systems, and decision workflows. Ensuring that these actions are trusted, authorized, and continuously governed becomes paramount. As a result, security evolves from protecting users and systems to establishing verifiable trust, identity, and delegated authority for autonomous agents.

Figure 5:
Cloud execution capabilities across data, control, and governance are expected to double by 2029



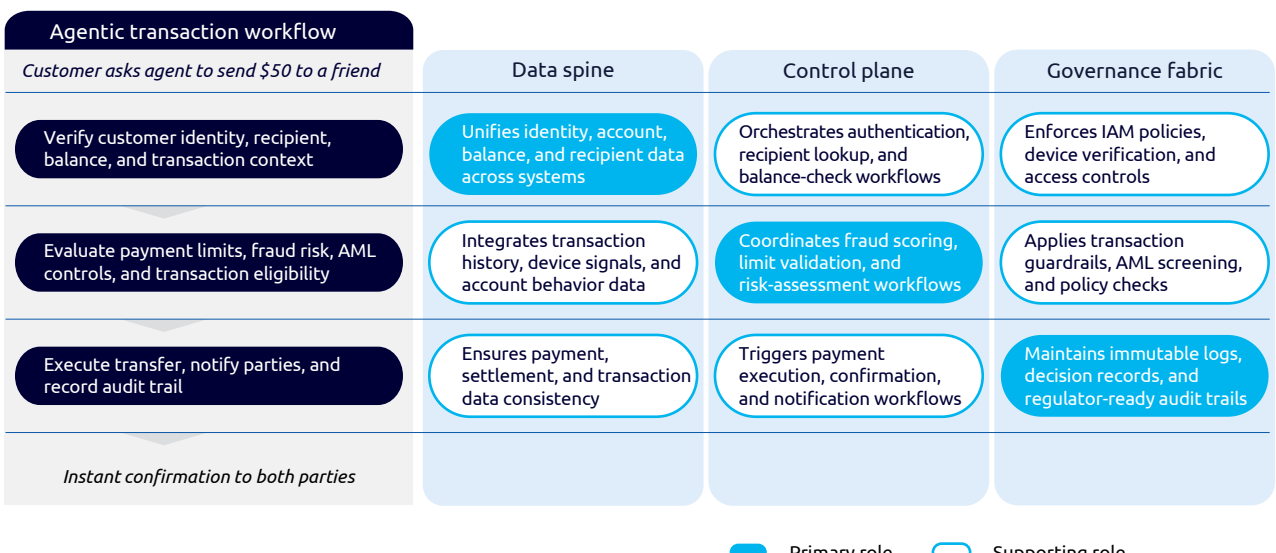
Source: Capgemini Research Institute for Financial Services Analysis, 2026; Capgemini Cloud Survey of Financial Services Firms 2026 (N=1,100)

Banking: Making a money transfer

Figure 6 illustrates how cloud capabilities support an agent-mediated banking payment workflow. Consider a customer instructing an agent to send USD 50 to a friend:

- The process begins with customer authentication, recipient verification, and account validation. At this stage, cloud brings together identity, account balance, and recipient data from multiple systems, enabling secure access and establishing the context required for the transaction.
- The workflow then moves to payment validation, where it assesses fraud risk, AML requirements, transaction limits, and payment agility. Cloud coordinates these decisioning processes, integrating behavioral signals, risk models, and policy controls to determine whether the transaction can proceed.
- Once approved, the payment is executed, notifications are issued, and the transaction is recorded. Cloud ensures seamless processing across systems while maintaining a complete audit trail.
- By preserving data consistency, automating fulfillment workflows, and supporting regulatory recordkeeping, cloud enables reliable and compliant transaction execution at scale.

Figure 6:
The role of cloud has evolved across an agentic banking workflow



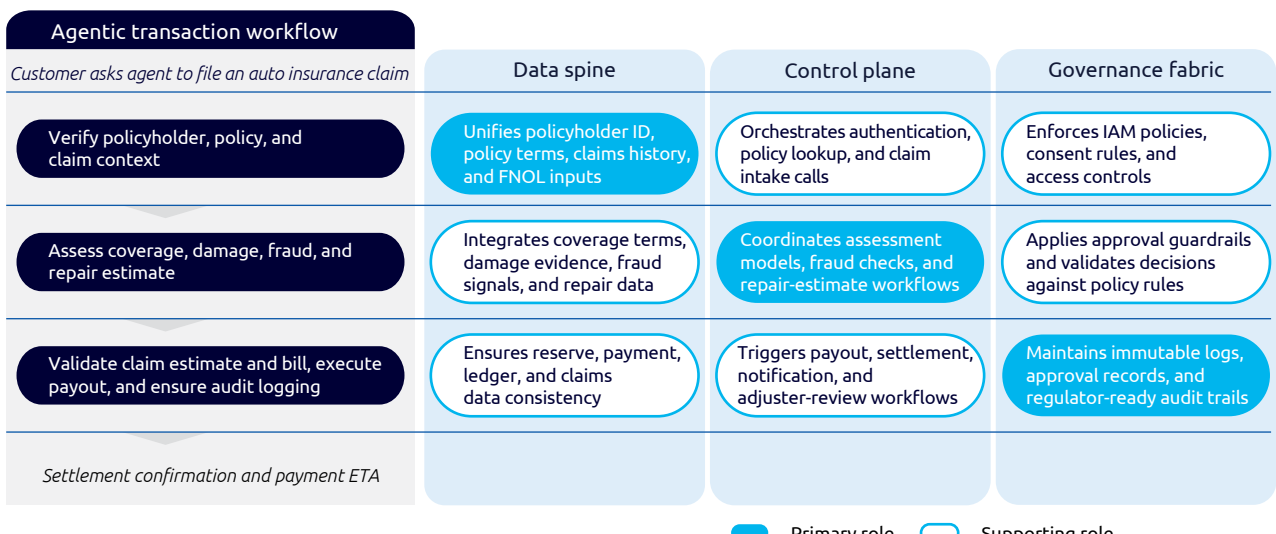
Source: Capgemini Research Institute for Financial Services Analysis, 2026.

Insurance: Filing an auto insurance claim

Figure 7 illustrates how cloud capabilities support an agent-mediated insurance claims workflow. Consider a customer asking an agent to file an auto insurance claim:

- The process begins with policyholder verification, policy validation, and claim intake. Cloud brings together policyholder identity, policy terms, claims history, and first notice of loss data from across systems, enabling the agent to establish claim context while ensuring secure access and consent compliance.
- The workflow then progresses to claim assessment, where coverage eligibility, damage estimates, fraud indicators, and settlement considerations are evaluated. Cloud coordinates these processes, integrating damage evidence, historical claims data, and fraud signals, enabling agents to assess the claim against policy and approval criteria.
- Once the claim is approved, the process moves to settlement, payout, and recordkeeping. Cloud supports seamless execution across claims, finance, and customer servicing systems while maintaining data consistency throughout the workflow.
- By automating notifications and preserving a complete record of decisions, approvals, and transactions, it enables efficient claims processing and regulatory-ready auditability.

Figure 7:
The evolving roles of cloud across an agentic insurance workflow



Source: Capgemini Research Institute for Financial Services Analysis, 2026.

"Cloud technology has shifted from an IT utility into a true growth engine, but as financial institutions move toward agentic AI, the scorecard fundamentally changes. A standalone model can answer questions, but an agentic platform delivers actions. Unlocking that ROI requires surrounding the model with persistent memory, enterprise context, and a robust action harness so multi-agent workflows can operate securely within regulatory guardrails."

Elizabeth Bourgoin, Director, Banking, Google Cloud North America

Autonomous agents require a new security model

Agents are starting to execute trades, approve payments and manage portfolios on their own. Identity and access frameworks were built for people and cannot govern systems acting at machine speed. Financial services needs a new security model built for agents, to authenticate, authorize, and govern what they do. This is also a critical dimension of our Enterprise Agentic AI Maturity Index.

50%

of executives expect AI agents to become an enterprise workforce by 2029

- AI agents are gaining economic authority across financial services: in fact, 50% of the executives we surveyed said they expect AI agents to become an enterprise workforce by 2029.
- During roughly the same period, end-user driven, agentic commerce transaction value is projected to rise from USD 8 billion in 2026 to USD 1.5 trillion in 2030, highlighting its growing importance.[4]
- As agents gain economic authority, traditional identity and access management (IAM) systems are insufficient to govern autonomous execution. In fact, only 18% of organizations across industries are confident that their current IAM can manage AI agent identities effectively.[5] Furthermore, 62% of financial services firms said their IAM stack lacks the resilience needed to support uninterrupted operations required for continuous agentic execution.[6]

To overcome the flaws of traditional IAM, new cloud-centric security approaches are emerging to govern autonomous agent execution. These fall into two categories:

- **Formal standards in development:** Workload Identity in Multi-System Environments (WIMSE) verify agent workloads across hybrid and multi-cloud environments and establish cryptographic trust between cloud workloads.[7]

"The challenge is not building intelligent agents, but ensuring they operate reliably at scale. Resilience, recovery mechanisms, oversight, and operational discipline become increasingly important as autonomous systems assume greater responsibility."

Richard Harmon, VP & Global Head of Financial Services, Red Hat



- **IAM vendor-proposed frameworks:**
  - Agent Identity Management System (AIMS) by Aembit, creates persistent identities for AI agents across cloud services and enforces delegated authorization with auditability. AIMS standardizes AI agents as distinct security principles rather than as a standard human user or a shared service account.[8]
  - Cross App Access (XAA) by Okta, enables secure agent access across SaaS and cloud platforms, and applies zero-trust controls to cross-cloud interactions.[9]

While these formal and IAM vendor-proposed standards mature, financial institutions are not waiting: they are piloting adjacent trust frameworks that tackle the same problem – how to authenticate, authorize, and audit the actions of autonomous agents. Two early examples illustrate the direction of travel:

- In 2026, HSBC piloted end-to-end agentic B2B transactions, enabling AI agents to autonomously execute procurement and payments in collaboration with Mastercard. The pilot integrated Mastercard Agent Pay with HSBC's payments infrastructure to enable secure, governed agentic transactions across buyers, procurement platforms, and suppliers.[10]
- In 2025, Visa and Massachusetts-based Akamai Technologies announced a collaboration to deploy Visa's Trusted Agent Protocol (TAP) across Akamai Cloud – enabling authentication, identity verification, and fraud controls for AI agents participating in agentic commerce. The initiative helps merchants verify both AI agents and the consumers they represent, creating a trusted framework for autonomous transactions.[11]

New trust frameworks for agents raise a broader question: Who controls execution once agents begin acting autonomously?

In the agentic era, sovereignty has evolved from where data sits to who controls agent execution

As purpose-built trust frameworks emerge to govern autonomous agent execution, focus is shifting from verifying agent identity to determining who sets the rules, authority boundaries, and accountability for agent actions.



"As AI scales, sovereignty becomes less about where data resides and more about how it is controlled, monitored, protected, and governed throughout processing, movement, and decision-making across distributed environments."

Hirenkumar Jadav
Head, Cloud CoE,
Maybank

Sovereignty is being redefined for financial services: it's become a strategic priority.

- In fact, sovereign cloud spending across industries in the EU is projected to grow 235% between 2025 and 2027 – from USD 6.9 billion to USD 23.1 billion.[12]
- Globally, spend on sovereign infrastructure during the same period is expected to grow at a rate of 87% - from USD 59.3 billion in 2025 to USD 110.6 billion in 2027.[12]

While most financial institutions are meeting sovereignty requirements through the cloud offerings of established hyperscalers, a segment of market players is supporting the development of independent European infrastructure to reduce reliance on any single provider, foreign or domestic. As one example, in 2026, Paris-based Mistral.ai secured USD 830 million from a consortium of seven banks (Bpifrance, BNP Paribas, Crédit Agricole CIB, HSBC, La Banque Postale, MUFG, and Natixis CIB) to build sovereign AI infrastructure that keeps compute, data, and execution within Europe.[13]

Traditionally, sovereignty in financial services centered on data residency – focusing on where firms store data, which jurisdictions govern it, and who accesses it. In the agentic era, however, sovereignty spans two layers, with cloud sovereignty forming the foundation and AI sovereignty sitting on top:

- **Cloud sovereignty** governs where data resides and which third parties can access it. This foundational layer addresses data location, access, and third-party exposure through data residency controls and jurisdictional oversight.
- **AI sovereignty** determines who authorizes, governs, and holds accountability for autonomous agent decisions, and who controls the models and compute powering them. This layer extends sovereignty into agent authority, execution paths, decision logic, and model and compute control, addressing whether models run on owned, licensed, or foreign-controlled infrastructure.

This evolution is clearly visible across the regulatory landscape, with key shifts defining sovereignty in the agentic era apparent across the EU (GDPR, NIS2, DORA, EU AI Act), Asia-Pacific (MAS, APRA), and North America (OSFI, CLOUD Act).

Regulatory changes manifest across three critical facets:

1. Reporting requirements are expanding

Where firms previously needed to maintain evidence of data location, vendors, and ICT dependencies, the agentic era demands that real-time traceability, oversight, and accountability for agent decisions are also maintained. This moves reporting from periodic, compliance-driven

documentation to continuous, decision-level auditability across autonomous workflows.

2. Primary ownership and accountability is broadening

Historically, accountability for technology governance resided primarily within technology, security, and risk functions. In the agentic era, accountability expands beyond these functions to include business leadership and the board, which now share responsibility for agent actions and their outcomes. This reflects the reality that agent-driven decisions increasingly affect core business outcomes, customer trust, and regulatory standing – requiring enterprise-wide governance rather than purely technical oversight.

3. Points of control are shifting

Traditional sovereignty controls focused on data location, access, and third-party exposure. In the agentic era, controls must extend to agent authority, execution paths, and decision logic. Organizations must now govern not only where data resides but also how agents act on that data, what decisions they are authorized to make, which execution paths they follow, and what logic underpins their autonomous actions.

Sovereignty matters for financial services firms today because:

- Core industry processes like credit, underwriting, and AML fall within resilience and risk frameworks of key regulations such as DORA and ICT requirements.
- Penalties compound across regions. For instance, the EU's Artificial Intelligence Act penalties can reach EUR 35 million (USD 41 million) or 7% of global turnover. Additional DORA and NIS2 liabilities increase overall regulatory exposure.[14]

To handle these requirements, leading cloud providers are acting today. For example, Amazon Web Services launched AWS European Sovereign Cloud in January 2026 to meet Europe's sovereignty requirements. This enables data residency, operational control, and governance independence with full AWS capabilities.[15] Similarly, Microsoft introduced new sovereign cloud services in November 2025 – expanding data residency, governance, and operational sovereignty controls across public and private cloud environments.[16]



"AI agents will raise expectations for service velocity, quality, and outcomes, making cloud central because its data models, infrastructure access, and operational consistency enable faster insight generation across service workflows at scale."

Jeff Murr
VP, IT Infrastructure and Operations, Genworth

With sovereignty requirements expanding from data governance to execution governance, the question of where and how these controls are enforced becomes critical, increasingly pointing to cloud platforms as the natural enforcement layer for sovereign agentic execution.

Governed agentic AI is consolidating cloud platform selection around hyperscalers

Cloud's evolving capabilities across data spine, control plane, and governance fabric layers are driving platform decisions. Participants in our executive survey identified three top sourcing criteria for cloud platforms:

1. Regulatory, data residency, and compliance certifications (50%)
2. Security, identity, and access management (47%)
3. Auditability and traceability at scale (39%)

"The next challenge in agentic AI is understanding how agents behave once they're deployed, and evaluating that continuously. Effective governance requires observability, appropriate access controls, meaningful human oversight, and evidence that makes agent behavior easy to understand."

Hemanth Sundararaj, Solutions Sales Leader-WW Financial Services, Microsoft, USA

Across these parameters, hyperscalers offer integrated capabilities that organizations value, such as global compliance certifications, sovereign cloud, and data residency. Further, they have built-in IAM, zero-trust, and confidential computing across layers. Hyperscalers also provide centralized audit logs with end-to-end traceability across cloud services.

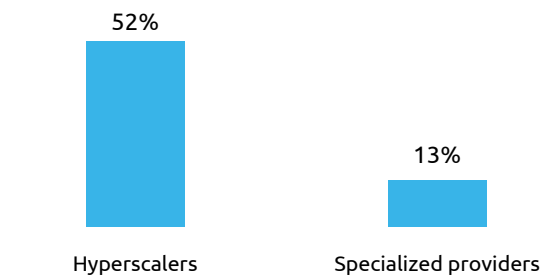
These capabilities make hyperscalers the preferred cloud sourcing partners for financial services organizations in 2026, much more so than other specialized providers. In fact, executive survey results in the bar chart in Figure 8 show hyperscalers hold a four to one lead.

Once hyperscalers provide a common regulatory, security, and auditability baseline, differentiation shifts from infrastructure to outcomes. Organizations will therefore measure value based on business outcomes rather than compute consumption.

50%

of executives prioritize regulatory, data residency, and compliance certifications for cloud platform selection

Figure 8:
Financial services firms prefer hyperscalers as cloud sourcing partners



Source: Capgemini Research Institute for Financial Services Analysis, 2026; Capgemini Cloud Survey of Financial Services Firms 2026 (N=1,100).





3 Unlock return on intelligence

Elevating cloud-AI investments from consumption metrics to business outcomes

Agent mediation adds new competitive dimensions across the customer journey. Customer agent-led product shortlisting works in tandem with the existing discovery model, as agents pull rates, eligibility criteria, and terms from multiple financial providers to shortlist products. Moreover, evidence-weighted agent decisioning adds depth to perception-led choice, as agents prioritize fit, price, SLA adherence, service history, and outcome reliability over brand, interface, or marketing influence. Finally, agent-enabled execution complements existing buying and service channels, as agents connect through APIs to core systems to complete purchases, manage changes, and provide ongoing support.

Agent mediation can deliver measurable improvements across business and customer outcomes, including a:

- 85% reduction in fraud false positives, leading to lower fraud and risk costs[17]
- 70% reduction in loan processing time, driving faster customer decisions[17]
- 65% reduction in manual KYC workload, resulting in reduced operational burden.[17]

As agents handle increasingly complex customer journeys, token consumption rises sharply, placing new pressure on AI economics.

Funding decisions now require stronger proof of AI value and governance of business outcomes

Declining token costs are offset by growing token consumption across agent workflows. From January 2025 to April 2026, there was a 93% reduction in inference costs per million tokens.[18] Despite declining token costs, more complex agentic workflows drove a 497% increase in enterprise token spend, offsetting many of the efficiency gains created by lower inference costs.[18]

This rising token spend exposes gaps in value attribution and governance. 48% of financial services firms cannot consistently demonstrate return on investment from their agentic AI initiatives, while 51% of firms lack mature value gates to govern continued AI funding. What's needed is outcome economics, which measures an organization's ability to translate AI activity into demonstrable business value and increasingly separates leaders from the rest of the market.

48%

of firms cannot consistently demonstrate ROI from agentic AI initiatives



"The most important value from AI may not come from expense reduction but from speed, quality, productivity, and the capacity to deliver more outcomes. Growth and customer impact are becoming more meaningful measures."

Robert Pick
Group Deputy CITO,
Tokio Marine

As a result, AI funding decisions are shifting from efficiency gains toward business outcomes. According to our survey, the number of executives prioritizing business outcomes as the primary AI funding criterion is expected to increase by 141% between 2026 and 2028, while the number prioritizing operational efficiency is expected to decline by 42%. Additionally, nearly 63% of the executives are either piloting or using intelligent routing at scale to optimize AI cost and performance, while another 27% are planning to implement it in the next 12 months.

63%

of executives are either piloting or using intelligent routing at scale to optimize AI cost and performance



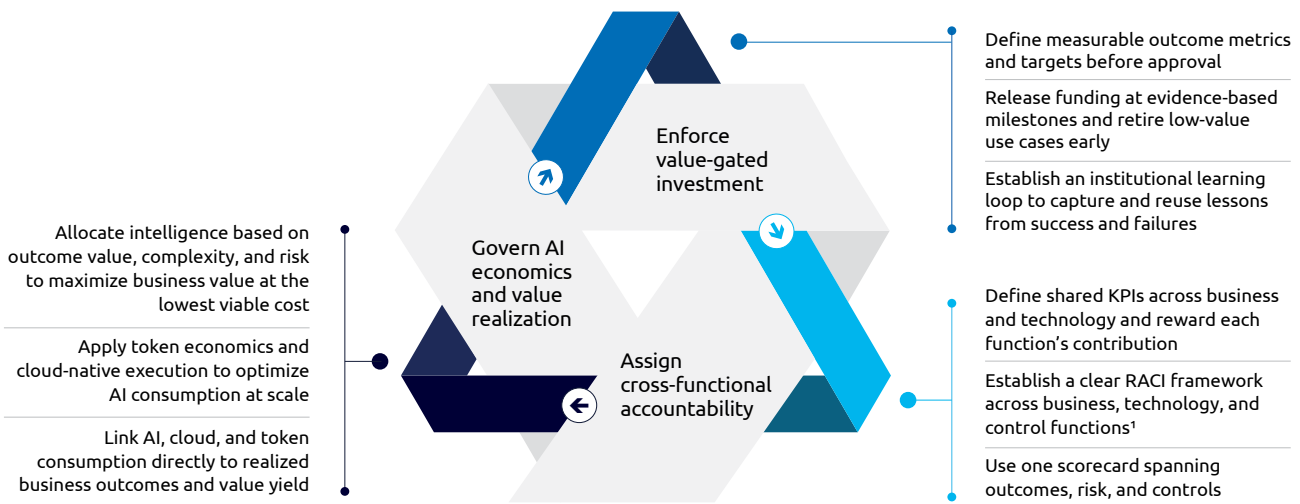
Value economics requires a disciplined execution and governance strategy

Closing the gap between AI consumption and realized value requires disciplined execution from funding to architecture. To obtain and measure value from their AI investments, Figure 9 illustrates how firms should prioritize three interconnected pillars:

- **Enforce value-gated investment:**

  AI initiatives will increasingly become business transformation programs rather than just technology experiments. To manage this transition efficiently, firms should not approve funding based solely on technical feasibility or model performance, but instead should establish explicit outcome metrics and release funding in stages as those outcomes are demonstrated. For example, an underwriting AI solution may advance from pilot to enterprise deployment only after proving measurable improvements in straight-through processing rates, underwriting productivity, or loss ratio performance. This approach helps organizations to concentrate investment on initiatives with proven business impact and retire projects that fail to demonstrate value.

Figure 9:
The value economics playbook is defined around three interconnected pillars



Source: Capgemini Research Institute for Financial Services Analysis, 2026.
Note: 1. RACI - Responsible, Accountable, Consulted, and Informed



"The measure of AI isn't how much of it an organization deploys, but whether it improves meaningful business outcomes. That means looking at improvements in areas such as speed, quality, productivity, customer experience and reliability, while balancing value, risk and cost."

Parimal Kumar
SVP, Digital, Services and Experience Tech Division, Erie Insurance

- **Assign cross-functional accountability:**

  Scaling agentic AI requires ownership that extends beyond technology teams. Leading firms define shared accountability across business leaders, technology teams, risk functions, and control groups to ensure AI outcomes remain aligned with enterprise objectives. For example, while technology teams may be responsible for system performance and orchestration, business leaders remain accountable for productivity, customer experience, or revenue outcomes, and risk teams oversee compliance and control effectiveness. Clearly defined accountability reduces the risk of isolated AI deployments that achieve technical success but fail to deliver measurable business value.

- **Govern AI economics and value realization:**

  As agentic AI adoption accelerates, organizations need governance frameworks that link AI consumption directly to business outcomes. Leading firms are creating mechanisms to measure and optimize the relationship between intelligence consumption (such as model usage, token consumption, and cloud execution costs) and the business value generated. Firms can track metrics such as cost per customer interaction, cost per underwriting decision, or cost per resolved service request alongside outcomes such as revenue growth, productivity gains, and risk reduction. This enables them to allocate intelligence resources to the highest-value use cases and continuously optimize the return on AI investments.





Move from agentic-AI experimentation to enterprise-wide leadership

Financial services firms have made meaningful progress in adopting agentic AI, but only a small group has translated experimentation into true enterprise-scale execution. Closing that gap requires that the three mutually reinforcing priorities discussed throughout this report – orchestration, cloud-governed execution, and outcome accountability – advance together. Capgemini's 'Enterprise Agentic AI Maturity Index,' newly developed for this report, tracks these priorities across five key dimensions:

1. Orchestration and deployment
2. Governance, trust and autonomy
3. Cloud execution authority
4. Security and identity
5. Operating discipline and outcome economics



"The greatest challenge is not proving agentic AI in controlled environments but operationalizing it at enterprise scale. As complexity increases, success depends on data readiness, integration maturity, governance, and the ability to deploy trusted capabilities consistently."

Shreejit Nair
SVP, Chief Information Officer, Ameritas

Significant effort and achievement across all five dimensions is required to move up the maturity ranking for enterprise agentic AI. Our assessment of financial services organizations today, summarized in Figure 10, finds that only 5% of firms demonstrate enterprise-grade maturity – these are the leaders on the industry's multi-agent ecosystem journey.

The gap between the leaders and the rest of the group shows up consistently across all three priorities:

- **Orchestrate agents across operations, technology, and customer experience**

  More than half (51%) of leaders operate enterprise-wide, multi-agent orchestration as opposed to only 8% of the rest of the firms within the other three levels. In order to move up in the index rankings, organizations should extend agent systems into business operations and the software development lifecycle to accelerate modernization and reduce technical debt, driving enterprise agility.

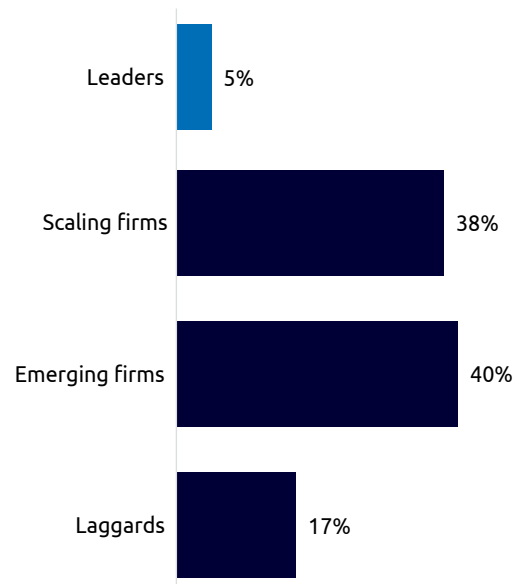
- **Building a foundational cloud layer for trusted, compliant, and sovereign AI at scale**

  A full 94% of leaders embed enterprise governance through automated enforcement or policy-as-code, compared with 43% of the rest of the organizations. Firms should embed governance, accountability, and oversight across data, operations, and model layers to enable scaled autonomy with enterprise-grade trust.

- **Elevate cloud-AI investments from consumption metrics to business outcomes**

  83% of leaders institutionalize or optimize outcome accountability, as compared with only 25% of the rest of firms. The opportunity for other organizations lies in enforcing value- and outcome-gated investment decisions while applying token and outcome economics to balance consumption, cost, and business value.

Figure 10:
The Enterprise Agentic AI Maturity Index



Source: Capgemini Research Institute for Financial Services Analysis, 2026.

Organizations that successfully combine these capabilities will be best positioned to convert agentic AI from a productivity initiative into a source of sustainable competitive advantage. In the agentic era, leadership will increasingly be defined by the ability to orchestrate autonomous execution, govern it with trust and accountability, and continuously convert AI activity into business outcomes.



Conclusion

The agentic AI era marks a fundamental shift in how financial services organizations operate, compete, and create value.

While adoption is accelerating, moving from isolated pilots to enterprise-scale execution remains a significant challenge. Success will require more than technology investment; it demands coordinated progress across orchestration, cloud architecture, security, governance, and business accountability.

Three priorities will define this transition: orchestrating agents across business and technology functions, establishing trusted cloud foundations for secure and compliant autonomous execution, and ensuring AI investments deliver measurable business outcomes.

Organizations that advance these capabilities together can move beyond incremental productivity gains toward broader enterprise transformation. Competitive differentiation will depend not on how many agents organizations deploy, but on how effectively they orchestrate, govern, and translate autonomous intelligence into trusted execution and demonstrable business value.



Methodology

The World Cloud Report for Financial Services 2027 draws data from two primary sources – the Global Financial Services Executive Survey, administered during June and July of 2026, and the Global Executive Interviews conducted in June through September of 2026. This primary research includes insights from 14 markets: Australia, Belgium, Brazil, Canada, France, Germany, Japan, Luxembourg, the Netherlands, Singapore, Spain, the United Arab Emirates, the United Kingdom, and the United States.

2026 Global Financial Services Executive Survey

We polled 1,100 leaders of financial services firms (C-suite executives, heads of Cloud, and more) about the evolution of AI in financial services and the changing role of cloud platforms. Respondents represented seven financial services domains – retail banking, wealth management, payments, corporate and investment banking, P&C insurance, life insurance, and health insurance – with balanced participation from each. These respondents also represented all of the markets listed above, covering all three regions of the globe – the Americas, Europe, and Asia-Pacific.

2026 Global Executive Interviews

The report also includes insights from focused interviews completed with cloud experts holding leadership positions across financial services and cloud ecosystem firms around the globe. These interviews included senior executives from banks, insurance companies, and hyperscalers.

2026 Agentic AI Maturity Index

The Enterprise Agentic AI Maturity Index assesses the maturity of financial services firms in scaling agentic AI from pilots to enterprise-wide execution. The index is based on responses from the 2026 Global Financial Services Executive Survey and evaluates organizations across five dimensions: orchestration and deployment, governance, trust and autonomy, cloud execution authority, security and identity, operating discipline and outcome economics. They were assigned weighted scores across these five dimensions and were classified as Laggards, Emerging Firms, Scaling Firms, or Leaders

Leaders (80-100)

Leaders have institutionalized agentic AI as an enterprise capability. They orchestrate agents across business functions, embed governance directly into execution, establish cloud as a trusted execution authority, secure agent identities and actions at scale, and consistently link AI investments to measurable business outcomes.

Scaling Firms (60-79)

Scaling Firms have moved beyond experimentation and are actively expanding agentic AI across multiple functions. They demonstrate strength in selected dimensions but have not yet institutionalized capabilities consistently across the enterprise.

Emerging Firms (40-59)

Emerging Firms have established foundational agentic AI capabilities but remain largely focused on pilots and localized deployments. Adoption is growing, but orchestration, governance, and accountability mechanisms are not yet mature enough to support enterprise-wide execution.

Laggards (0-39)

Laggards remain in the early stages of agentic AI adoption. Efforts are fragmented and experimental, with limited operational deployment and few foundational capabilities required for scaled autonomous execution.

Partner with Capgemini

Our FS-specific cloud-enabled agentic AI platform is designed to accelerate enterprise transformation: A single, governed, platform to build, modernize, manage and run applications, while reimagining financial services business processes. One governed platform to build, modernize, manage and run applications, while reimagining Financial Services business processes.

Custom Application Development

Custom Apps Development is a comprehensive capability to design, build, modernize and industrialize cloud-native applications, full-stack digital solutions, APIs, integration services, agentic applications and mainframe-connected solutions. It combines next-generation AI-enabled SDLC, Platform Engineering, API Engineering, DevSecOps and legacy modernization to rebuild applications, decompose monoliths, expose mainframe business capabilities through APIs and progressively transition workloads to modern architectures.

At Capgemini, we connect these capabilities with the client's architecture, engineering toolchain, security controls and cloud platforms, while OrchestrateAI brings domain assets, certified patterns, multi-agent quality controls and automated traceability to accelerate delivery and reduce rework.

Custom Application Management

Custom Application Management is a comprehensive capability to enhance, modernize, govern and continuously improve brownfield application estates across cloud, legacy and hybrid environments. It combines a unified application Control Plane, AI-Ops, SRE, brownfield modernization, application maintenance, security remediation and release management to provide end-to-end visibility and control across applications, dependencies, service health, change and technical debt.

At Capgemini, we connect application knowledge, engineering toolchains, ITSM, observability and cloud platforms to create a continuously updated application knowledge twin, enabling impact-aware changes, predictive maintenance, risk-based releases and progressive modernization aligned with business criticality, resilience and cost.

Custom Application Run

Custom Application Run is a comprehensive capability to operate, govern and continuously optimize business-critical applications across cloud, legacy and hybrid environments. It combines Knowledge Engineering, a unified operations Control Plane, AI-Ops, SRE, Auto-Heal services, Agentic Governance and Token Cost Management to connect application telemetry, service dependencies, operational knowledge, business impact and technology consumption.

At Capgemini, we bring together ITSM, observability, cloud, security and FinOps platforms to enable predictive issue detection, contextual root-cause analysis, governed autonomous remediation and continuous optimization of reliability, performance, cloud cost and AI consumption.



Ask the experts



Ravinder Khokhar
FS-Global Head of Cloud
ravinder.khokhar@capgemini.com

Ravi leads Capgemini's worldwide financial services cloud initiatives. He is a domain-focused technology leader with a proven track record of successfully establishing long-tenured client partnerships and delivering results across financial services.



Parminder Dhillon
FS-Global Head Cloud Operations and Acceleration
parminder.dhillon@capgemini.com

Parminder leads global Cloud Operations, Acceleration, and FinOps for Financial Services. With 15+ years in enterprise IT, she architects governance, management, and operations for multi-cloud environments, specializing in cost optimization, automated governance, platform operations, and financial efficiency to drive measurable business outcomes



Ramandeep Singh
FS-Global Platform engineering head and cloud solutions lead
ramandeep.singh@capgemini.com

Raman leads the Financial Services Global Platform Engineering & Cloud Solutions capability unit. He has 20+ years of IT experience, focusing on large enterprise application design and development. He is a leader focused on business growth, customer delight, and people.



Mamatha S
FS- Global Head AWS COE
mamatha.shanmugam@capgemini.com

Mamatha leads the Global AWS Center of Excellence for Financial Services and enables solutions across banking, capital markets, and insurance.



Vivek Desai
FS-Global Head Microsoft COE
vivek.desai@capgemini.com

Vivek leads the Global Microsoft Center of Excellence for Financial Services. He is responsible for hands-on solution architecture and strategy, and engineering at scale for complex cloud transformation initiatives across the banking, payments, capital markets, and insurance sectors.



Daniel Wolff
FS-Global Head Google Cloud COE
daniel.wolff@capgemini.com

Daniel leads the Global Google Cloud Platform Center of Excellence. He is a senior technology executive known for accelerating sustainable company growth within financial services. He focuses on unlocking new opportunities for improved competitive advantage, cost containment, process re-engineering, and business agility.

Ask the experts



Ajay Walgude
FS-Global Head IBM/Red Hat COE
ajay.walgude@capgemini.com

Ajay leads the Global IBM and Red Hat Hybrid CoE for FS. With over 20 years of experience in financial services, he is also the founder of Turiya Lab, an initiative that builds technology labs focused on physics, chemistry, and astronomy across schools in India.



Luca Russignan
Global Head of Capgemini Research Institute for Financial Services
luca.russignan@capgemini.com

Luca leads Capgemini's global portfolio of thought leadership in financial services. Luca provides strategic insights for financial services to shape industry dialogue on transformation. Each year, he partners with 80+ C-suite leaders at global financial institutions, drawing on research across multiple markets to identify emerging patterns and translate them into strategic guidance.



Vivek Singh
Head of Banking of Capgemini Research Institute for Financial Services
vivek-kumar.singh@capgemini.com

Vivek leads the banking, wealth management, FinTech, and payments sectors within the Capgemini Research Institute for Financial Services. He has over 12 years of experience in digital, consulting, and business strategy. Based in New York, Vivek is a tech enthusiast who tracks industry disruptions, thought leadership programs, and business development.

Key contacts

Gen AI

Sagar More (UK)
sagar.d.more@capgemini.com

Kieran Maltz
Kieran.maltz@capgemini.com

Manjula Samuel (India)
manjula.samuel@capgemini.com

Agentic AI

Ramandeep Singh (FS)
ramandeep.singh@capgemini.com

Kamlesh Pandey (Banking)
Kamlesh.pandey@capgemini.com

Deepak Sharma (Insurance)
deepak.h.sharma@capgemini.com

APAC

Rishi Garg
rishi.garg@capgemini.com

Andrew Markwell
andrew.markwell@capgemini.com

Continental Europe

Christian Klaus Schenk
christian-klaus.schenk@capgemini.com

Laurent Lefrancois
laurent.lefrancois@capgemini.com

Tudor Uricec
tudor.uricec@capgemini.com

AWS

Alex Manta (UK)
alex.manta@capgemini.com

Chris Covington (US)
chris.covington@capgemini.com

Mamatha S (India)
mamatha.shanmugam@capgemini.com

Ashwini Thambi (India)
ashwini.thambi@capgemini.com

Microsoft

Pavel Dokov (UK)
pavel.dokov@capgemini.com

Sameer Yande (US)
sameer.yande@capgemini.com

Vinay Bhatia (US)
vinay.a.bhatia@capgemini.com

Swaminathan Santhanam (India)
swaminathan.santhanam@capgemini.com

IBM

Shyamasundar Mallobha (India)
shyamasundar.mallobha@capgemini.com

Subhash Balan (India)
subash.balan@capgemini.com

Google Cloud

Naveen Kadala (India)
naveen-kadala.manikoth@capgemini.com

Acknowledgments

We extend special thanks to all the financial services and technology/cloud-ecosystem firms and individuals who participated in our executive interviews and surveys.

The following firms agreed to be publicly named:

FS firms: Ameritas, Barclays, Bradesco, Erie Insurance, Genworth, Maybank, OP Pohjola, PEMCO, Tokio Marine

Technology firms: Amazon Web Services, Google Cloud, IBM, Microsoft, Red Hat

Survey partners: INJ Partners

Teams and individuals

Capgemini Financial Services Cloud team: Ajay Walgude, Daniel Wolff, Parminder Dhillon, Ramandeep Singh, Ravinder Khokhar, and Vivek Desai.

Capgemini Research Institute for Financial Services: Luca Russignan, Vivek Singh, Pranav Shivram, and Chayan Bandyopadhyay for their overall leadership for this year's report. Aysha Haider, Anish Sukumaran, Himanshu Mishra, and Prerna Remy for in-depth market analysis, research, compilation, and drafting of the findings. Tamara Berry for editorial contributions and content leadership. Dinesh Dhandapani Dhesigan for graphical interpretation and design.

Marketing team: Meghala Nair, Jyoti Goyal, Sunanda Ganguly, David Merrill, Rob Hurst, Neha George, Fahd Pasha, Manasi Sakpal, and Manisha Singh for their overall marketing support for the report.

Creative Services Team: Sushmitha Kunaparaju, Pravin Kimbahune, Priyanshi Srivastava and Sushmita Singh for report production.

Capgemini Research Institute for Financial Services Team



Aysha Haider
Project Manager
aysha.batul-haider@capgemini.com

Project Manager for Financial services thought leadership with over 8 years' experience in cross-sector research and consulting.

Endnotes

1. *Pymnts, "BNY's AI Strategy Signals a New Era of Platform Banking" April 16,2026.*
2. *OpenAI, "BNY builds "AI for everyone, everywhere" with OpenAI;" December 12,2025.*
3. *WSJ,* "How Morgan Stanley Tackled One of Coding's Toughest Problems;" June 3, 2025.
4. *Human Security,* "The 2026 State of AI Traffic & Cyberthreat Benchmark Report;" August 20, 2026.
5. *CSA "Securing Autonomous AI Agents" February 4, 2026.*
6. *EMA Research* "EMA Research Finds Legacy IAM Systems Are Slowing AI Adoption in Financial Sector" May 19, 2026
7. *IETF,* "Workload Identity in a Multi System Environment (WIMSE) Architecture;" June 7, 2026.
8. *AEMBIT,* "AIMS: A Model for AI Agent Identity;" May 1, 2026.
9. *OKTA,* "Cross App Access: Securing AI agent and app-to-app connections;" November 25, 2025.
10. *HSBC,* "HSBC expands B2B payments capabilities in Singapore;" June 9, 2026.
11. *VISA,* "Akamai and Visa Collaborate to Build Trust in Agentic Commerce;" December 17, 2025.
12. *Gartner,* "Gartner Says Worldwide Sovereign Cloud IaaS Spending Will Total $80 Billion in 2026;" February 9, 2026.
13. *CNBC,* "Mistral secures $830 million in debt financing to fund AI data center;" March 30, 2026.
14. *Holistic AI,* "Penalties of the EU AI Act: The High Cost of Non-Compliance;" January 23, 2025.
15. *AWS,* "Opening the AWS European Sovereign Cloud;" January 14, 2026.
16. *Microsoft Azure,* "Microsoft strengthens sovereign cloud capabilities with new services;" January 15, 2026.
17. *Kore AI,* "Agentic AI in financial services and insurance sector;" February 20, 2026.
18. *AlphaSense,* "ATB Cormark Capital Markets Reports;" August 27, 2026.

About us

Capgemini is the business transformation partner for enterprises in the age of AI. We help organizations imagine and build an intelligent, sustainable future, combining AI, technology and human ingenuity to transform how they operate, innovate and grow. With unique end-to-end capabilities spanning strategy, technology, engineering and intelligent operations, we bring together deep industry expertise and market-leading capabilities in AI, cloud and data to turn ambition into measurable business outcomes at scale. Supported by a robust ecosystem of partners and nearly 60 years of expertise, Capgemini is a responsible and diverse global organization of over 410,000 team members in more than 50 countries. The Group reported 2025 revenues of €22.5 billion.
Make it real | *www.capgemini.com*

Capgemini RESEARCH INSTITUTE

Financial Services World Report Series

The Capgemini Research Institute for Financial Services is an in-house think tank focused on digital innovation and technology issues impacting global banks, wealth management firms, and insurers. The institute annually publishes its signature Financial Services World Reports, which draw upon primary research voice-of-the-customer surveys, CXO interviews, and partnerships with technology companies and academia. These data-driven perspectives explore how financial institutions can meet emerging business challenges with transformative thinking enabled by technology and data.
To find out more or to subscribe to receive reports as they launch, visit *worldreports.capgemini.com*

Disclaimer

The information contained herein is general in nature and is not intended and should not be construed as professional advice or opinion provided by Capgemini America, Inc. (Capgemini) to the user. Capgemini assumes no liability for errors or omissions, or use of this material. This document is provided for informational purposes only; it is meant solely to provide helpful information to the user. This document does not purport to be a complete statement of the approaches or steps necessary for a business to address or solve any particular matter or to accomplish any particular business goal. The user also is cautioned that this material may not be applicable to, or suitable for, the user's specific circumstances or needs and may require consideration of additional factors if any action is to be contemplated. The text of this document was originally written in English. Translation to languages other than English is provided as a convenience to our users. Capgemini disclaims any responsibility for translation inaccuracies. The information provided herein is on an as-is basis. Capgemini disclaims any and all representations and warranties of any kind.



Read more of our research
https://www.capgemini.com/insights/research-library/world-cloud-report/

For more information, please contact:

financialservices@capgemini.com

For press inquiries, please contact:

Fahd Pasha

Capgemini Financial Services

Tel.: +1 647 860 3777

fahd.pasha@capgemini.com

Capgemini

Make it *real.*