Reimagining financial services with agentic AI

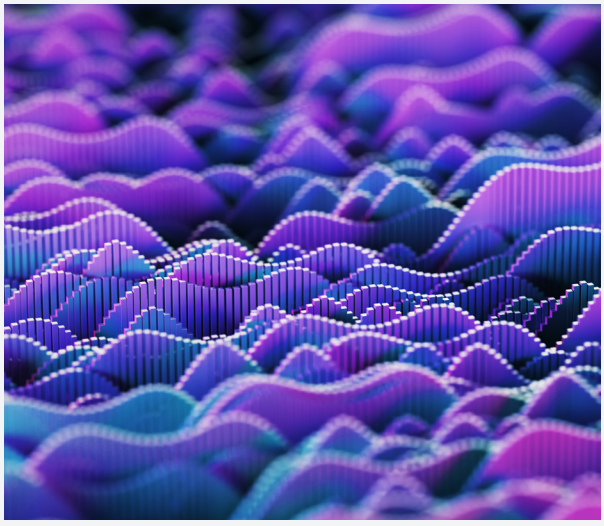


How FS firms can transform use cases and accelerate impact with AI

Capgemini

In brief

- **Agentic AI goes beyond generative chat:** it employs agents that plan, act, and adapt across banking and insurance throughout the cloud modernization lifecycle.
- **Most financial services (FS) firms stall at pilots:** Why? Because they're worried about the ramifications of agent behaviors, they haven't experienced it themselves and need someone who has guided others to implement at scale. Agentic AI reimagines business flows by enabling systems to make autonomous decisions, adapt in real time, and optimize outcomes across complex processes. A proven approach and platform accelerate time to value by aligning AI capabilities with clear business goals, reducing friction from experimentation to deployment.
- **Governance is non-negotiable:** explainability, human-in-the-loop, logging, and security must all be built in from day one to satisfy regulators and auditors. Every agent interaction should be made visible to avoid black box behavior.



Why agentic AI? And why now?

Financial services organizations are under pressure. Costs are up – so is inflation. Regulators are watching. Customers expect speed and personalization. Agentic solutions can think, take decisions, and pick the right path to execute, which brings sophisticated automation to human intensive legacy processes.

Generative AI (GenAI) helps with assistive work – like answering questions and drafting copy – but it's limited. Agentic AI is the new frontier, offering agents that can plan, act, and adapt in real time. Take claims triage, fraud checks, and loan onboarding as few examples. The payoff is cycle-time reduction, higher straight-through processing, and consistent decision-making in a cost-effective manner. While generative AI enhances individual tasks, agentic AI drives end-to-end value by integrating reasoning, memory, and control into enterprise-scale processes.

Most financial services organizations recognize this need. The question is: how? Learn and build everything in-house, or accelerate with a baseline platform and proven partner? Scaling is hard. And the reality on the ground is clear: when it comes to AI adoption, only 26% of companies have the capabilities to move beyond proof of concept and create value at scale.[1] The rest stall in pilots.

Meanwhile, AI is broadly used. As many as 78% of organizations report using AI in at least one function.[2] And yet, end-to-end redesign still lags.

1 Boston Consulting Group "AI Adoption in 2024: 74% of Companies Struggle to Achieve and Scale Value;" October 24, 2024.

2 McKinsey "The state of AI: How organizations are rewiring to capture value;" March 12, 2025.

Why so many stalls? Several reasons:

- Classic project risk plays a role. Only 12% of business transformation projects are successful.[3]
- Newer analyses of large IT programs show that only 31% meet their "success" criteria.[4] Big programs fail far more often and suffer overruns.
- Agentic AI is a complex, multi-dependency program: governance, data, models, controls, and organizational change all at once.

Building bespoke "inside the perimeter" programs can work, but many teams underestimate the lift to industrialize beyond pilots. A pragmatic approach starts with proven partners for bringing agents, runs inside the firm's environment, and layers in domain-specific controls from day one.

Where the financial services industry is seeing benefits

In insurance:

Agentic AI is rapidly transforming the insurance value chain – bringing speed, intelligence, and efficiency to every touchpoint. In claims, it enables faster resolution through autonomous decisioning and real-time data synthesis. In underwriting, agentic systems dynamically assess risk with greater precision, reducing manual effort and improving accuracy. Distribution is becoming more personalized and proactive, as agentic AI hyper-personalizes offerings to customer needs and behaviors.

And in servicing, intelligent agents are enhancing responsiveness and accuracy, driving better customer experiences, while reducing operational costs. Together, these advancements are helping insurers unlock differentiated value, accelerate time to market, and build more resilient, adaptive businesses.

In banking:

Banks are moving towards agentic AI platforms to drive intelligent automation, personalization, and operational agility. In retail banking, agents are enabling hyper-personalized customer engagement and real-time financial guidance, improving satisfaction and loyalty. Wealth management teams are leveraging agentic AI to deliver tailored portfolio strategies and proactive insights, enhancing advisor productivity and client outcomes. In investment management, agentic systems accelerate research, optimize risk models, and streamline compliance, allowing firms to respond faster to market shifts. Meanwhile, in cards and payments, agentic AI enhances fraud detection, automates dispute resolution, and enables dynamic credit decisioning – delivering safer, faster, and more adaptive services. Together, these innovations are enabling banks to reduce operational costs and streamline speed to market across the enterprise.



3 Bain "The Three Common Transformation Talent Mistakes and How to Avoid Them;" April 14, 2024.

4 OpenCommons "CHAOS Report on IT Project Outcomes;" accessed October 1, 2025.

Agentic AI: driving acceleration and efficiency in the cloud modernization lifecycle

In legacy modernization:

Agentic AI is becoming a catalyst for legacy modernization by bridging the gap between outdated systems and future-ready architectures. By deploying intelligent agents that autonomously assess, transform, and validate legacy and mainframe systems, organizations can reduce dependency on scarce SMEs while improving accuracy and functional consistency. These agents bring contextual awareness and dynamic orchestration to legacy environments – streamlining data flows, optimizing processes, and enhancing interoperability across silos. This approach doesn't just reduce technical debt: it also unlocks faster innovation cycles, allowing banks and insurers to modernize incrementally while delivering safer, smarter, and more adaptive services to clients.

In application modernization:

Application modernization is being radically accelerated by the integration of agentic AI, which brings intelligence and velocity to every phase of the software development lifecycle (SDLC). These capabilities extend across the SDLC, from requirements analysis and code generation to testing, integration, and deployment. Agentic AI accelerates development cycles by dynamically identifying optimization opportunities, orchestrating workflows, and ensuring continuous validation across environments. One Capgemini customer, a US bank, reported a 20–30% increase in developer throughput with an agentic approach, as developers spent less time on code comprehension, validation, and documentation, and more time on value-added tasks – enabling more features to be delivered per sprint. This not only streamline modernization efforts, but also empowers firms to build adaptive, resilient applications more quickly – unlocking differentiated value and reducing time to market.

In cloud management:

Intelligent agents are taking the reins – autonomously monitoring workloads, assisting production support teams, predicting resource needs and failure, and orchestrating cloud environments across hybrid and multi-cloud landscapes with precision and speed. These systems don't just react – they anticipate and fix, continuously optimizing for performance, cost, and compliance in real time. The result? Financial institutions are scaling faster, slashing operational overheads, and unlocking high-performing, secure cloud ecosystems that fuel innovation. With agentic AI, cloud management becomes not just smarter – but radically more agile, resilient, and future-ready.

Governance you can take to audit

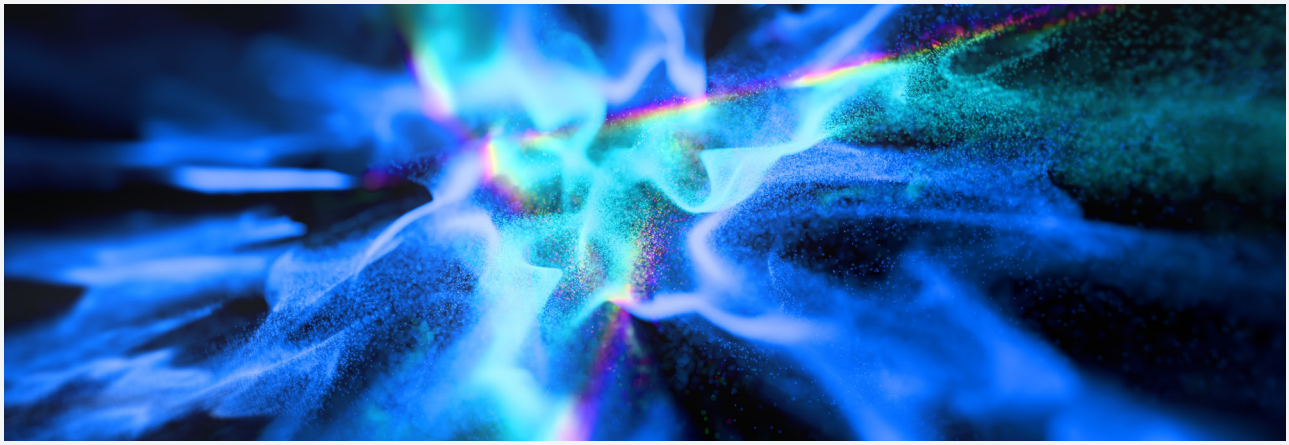
For financial institutions, it's not just about driving efficiencies: regulators are signaling tighter oversight too. The EU AI Act is staged: prohibitions and literacy obligations apply from February 2, 2025. Governance rules and GPAI obligations in August 2025. Full application by August 2, 2026, with extended timelines for certain high-risk systems into 2027.[5] FS firms with EU exposure must plan now.

In the US, agencies continue to emphasize existing risk-management frameworks and third-party governance. That means firms can expect scrutiny on vendor controls, explainability, and model risk. The US Department of the Treasury has urged more coordination across regulators, with the Office of the Comptroller of the Currency (OCC), Federal Reserve (FED), and Federal Deposit Insurance Corporation (FDIC) continuing to reinforce third-party risk expectations.[6]

For FS firms, explainability isn't optional. Banking and insurance decisions must stand up to model risk governance, audit, and customer review.

Governance must-haves

- Approved LLMs: tools that deliver unbiased outcomes.
- Expert oversight: human-in-the-loop at defined decision gates.
- Agent logging and observability: full event trails, prompts, model versions, and data lineage.
- Control structures: role-based permissions, policy constraints, and kill switches.
- Bias, drift, and error checks: pre-deployment tests and ongoing monitoring.
- Compliance-by-design: regulatory checkpoints embedded in workflows.
- Security inside your perimeter: agents are run within the firm's infrastructure, no data leaves the enterprise, and only pre-approved models are used.
- Third-party risk controls: contractual SLAs on audit access, incident response, data use, and model updates.
- Well-defined strategy to adopt tools, platforms, and agents.



5 European Commission "AI Act;" accessed October 1, 2025.

6 US Department of the Treasury "Artificial Intelligence in Financial Services;" December 2024.

What to look for in an agentic partner or platform

Many firms opt to build their own platforms, which can offer greater security and control. That's valid – but beware of the scaling tax. The most successful programs combine inside-the-perimeter deployment with pre-built, domain-specific agents, and a playbook that's been battle-tested across banks and insurers.

Use this checklist:

1. **Human-centric by design:** Agents augment people. Clear handoffs. Friendly UX. No anthropomorphism.

   Ask: *Where are the human approval points? How are users trained?*

2. **Responsible by default:** Controls, explainability, bias tests, and privacy are built-in, not bolted-on.

   Ask: *Show me the audit log and the kill switch.*

3. **Clarity and simplicity:** Short time-to-first-value. Minimal jargon. Clean dashboards.

   Ask: *How quickly can we get to a production-grade PoC in one use case?*

4. **Proven outcomes, not promises:** Case studies with quantified impact, across metrics like TAT, Straight Through Processing (STP), fraud precision and recall, and CSAT scores.

   Ask: *What baselines did you start from, and what moved?*

5. **Adaptable architecture:** Cloud- and model-agnostic. Works inside your environment. Modular agents that plug into admin systems, CRM, and claims cores.

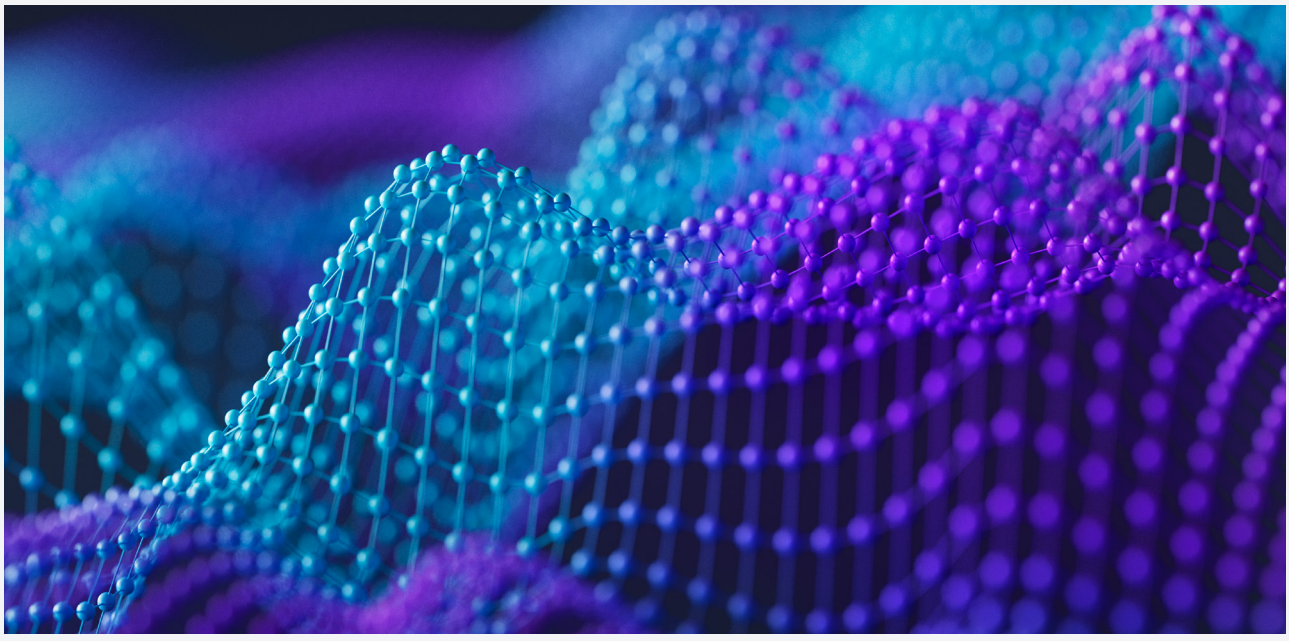
   Ask: *Does this model and platform align with our overall cloud/AI strategy?*

6. **Tone and trust:** Calm confidence beats hype. Clear risk disclosures. Pragmatic roadmaps.

   Ask: *How do you measure and report risk, cost, and value? How often?*

7. **Flexibility:** Agentic ecosystem is maturing at a lightning speed with advancements taking place every day rather than every week.

   Ask: *Am I tied to one approach or solution?*



Practical FAQs leaders will ask

1. **Will our data leave our perimeter?**
   It doesn't have to. Many programs deploy agents inside the firm's environment and restrict models to those already approved by risk and security.

2. **Can we swap models later?**
   Prioritize cloud- and model-agnostic orchestration. You should be able to swap or fine-tune without rewriting the whole stack.

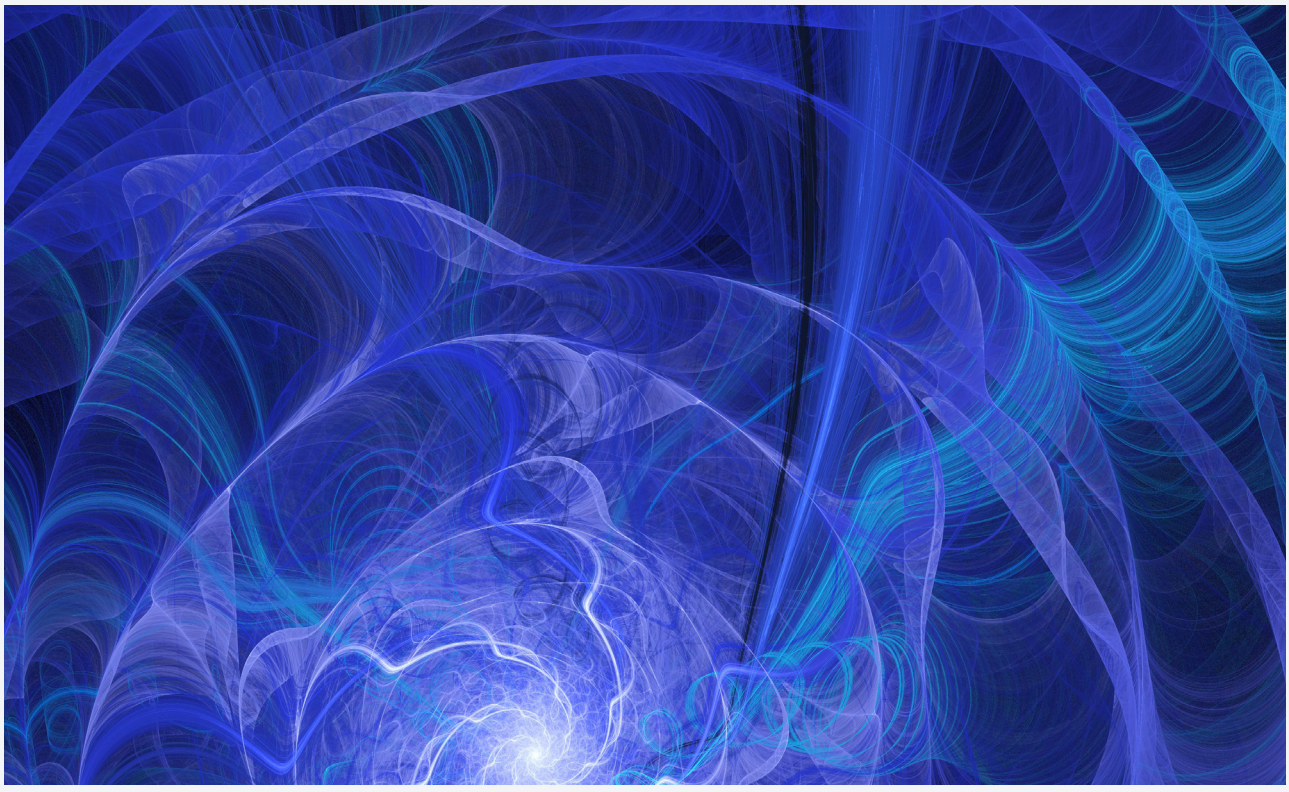
3. **What about explainability?**
   Require decision logs that tie inputs, prompts, model versions, and outcomes together. Make explanations readable for customers, not just auditors.

4. **How do we avoid bias and drift?**
   Run pre-deployment tests, then monitor post-deployment with drift and error checks. Escalate to a human when confidence drops below a threshold or when decisions impact vulnerable customers.

5. **What is our testing approach?**
   Design test cases while documenting the agentic requirements. Monitor using tools to baseline the performance and measure drift, if any, while it progresses across environments.

6. **How do we keep humans in control?**
   Design explicit approval gates, have more involvement initially, and reduce reliance on approval gates as the system matures and learns. Use a visible kill switch with clear roles and permissions.

7. **What's a realistic ROI?**
   Don't overpromise. Anchor on operational outcomes – like the TAT and STP, the number of reworks needed and FTE hours reallocated, the complaint rate, fraud metrics, and false positives. Start with one business process, prove it, then scale.





In conclusion: Proof beats promise

Agentic AI isn't magic – it's disciplined engineering and change management. The winners aren't the firms with the flashiest demos. They're the firms that have the domain-specific knowledge of workflows, deploy with strong guardrails, prove impact, and scale responsibly. In a market where technology is changing every day, that methodical approach is the difference between another pilot and measurable business success.

Ready to transform your financial services enterprise? Let's start with a strategic agentic AI plan tailored to your business.

Authors:



Kamlesh Pandey

Platform Engineering & Agentic AI Leader, Capgemini

Kamlesh Pandey is a Senior Director at Capgemini, specializing in Cloud Platform Engineering and agentic AI solutions. He is leading Capgemini's agentic platform build specific to financial services leveraging cloud, AI, and data.



Leah Klempner

Cloud & AI Acceleration, Capgemini

Leah Klempner is a Senior Consultant at Capgemini, specializing in cloud and AI acceleration. She plays a key role in leading Capgemini's agentic AI platform capabilities and go-to-market, to drive enterprise transformation in financial services.

About Capgemini

Capgemini is an AI-powered global business and technology transformation partner, delivering tangible business value. We imagine the future of organisations and make it real with AI, technology and people. With our strong heritage of nearly 60 years, we are a responsible and diverse group of 420,000 team members in more than 50 countries. We deliver end-to-end services and solutions with our deep industry expertise and strong partner ecosystem, leveraging our capabilities across strategy, technology, design, engineering and business operations. The Group reported 2024 global revenues of €22.1 billion.

Make it *real*.
www.capgemini.com

Capgemini

© Capgemini, 2025. All rights reserved.