Deloitte.

# Agentic AI in financial services: From vision to value

September 2025

# Table of contents

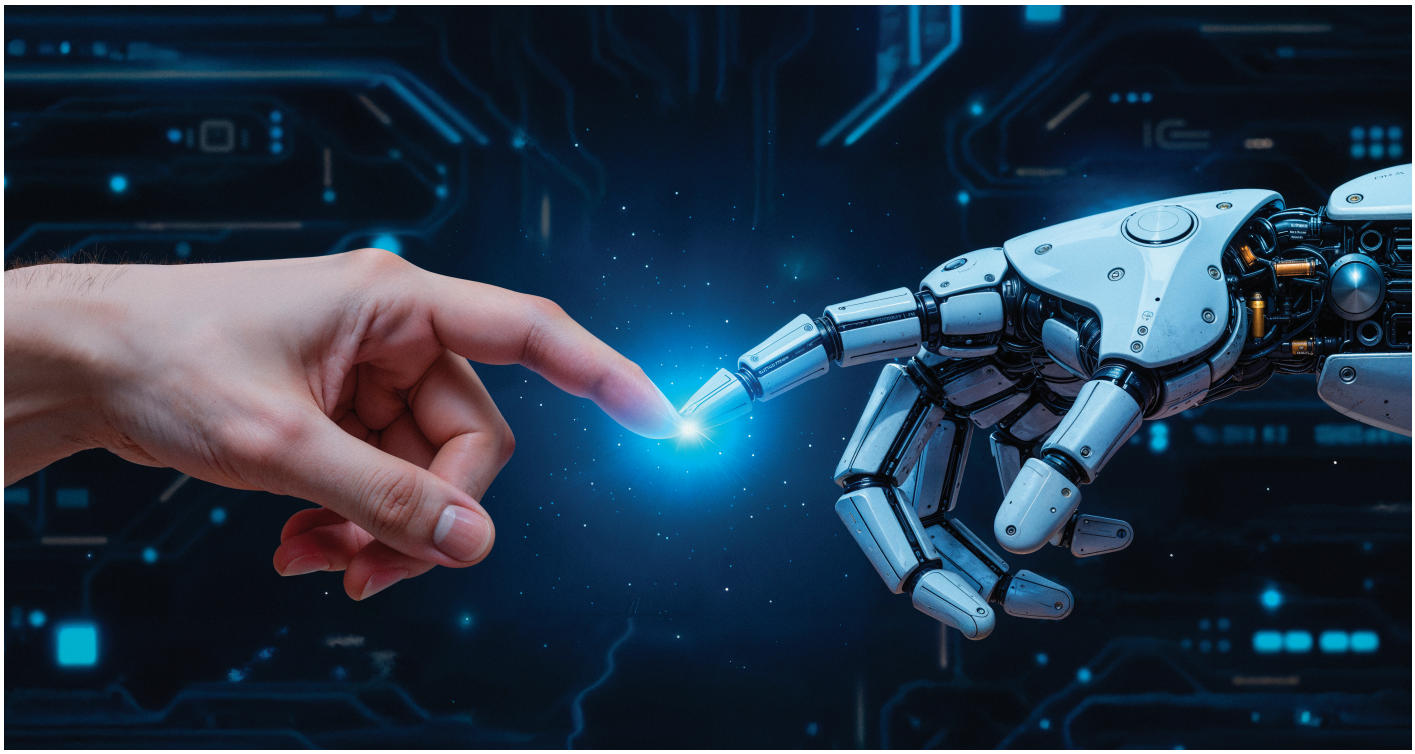

# Overview

Agentic AI represents a significant development in Artificial Intelligence (AI), shifting from reactive tools to autonomous systems that actively plan, decide and execute tasks without substantial human intervention.

At their core, agentic AI systems consist of autonomous AI agents that are **reasoning engines and exhibit agency.** Agentic AI systems are goal-oriented, context-aware and capable of independent decision-making.

These systems represent a natural progression, built on Large Language Models (LLMs) and generative AI (GenAI) capabilities. Traditional AI relied on pre-defined rules and data to implement specific tasks. GenAI built on this foundation with the ability to generate original content. Agentic AI is revolutionising the space with its ability to adapt to changing circumstances independently. These intelligent systems are capable of reasoning, evolving with experience and collaborating seamlessly with humans and other digital agents.

This shift from traditional to agentic AI is being driven by a convergence of technological maturity and business needs. Significant advancements have been made in the domain of LLMs, memory architectures and other areas, making it technically feasible to build systems that operate with greater autonomy and contextual understanding.

The financial services (FS) industry continues to face the pressure to drive down costs, improve sales and delivery efficiencies, streamline processes, respond faster to change, embed automation across workflows and identify and mitigate risks.

Agentic AI helps meet these demands by moving beyond process augmentation to fully autonomous, end-to-end execution across planning, decision-making and task completion with minimal human intervention.

Given the maturity of India's FS industry, agentic AI is likely to have a decisive impact on areas including wealth management, Anti-Money Laundering (AML) and fraud analytics, high volume Straight Through Processing (STPs), customer servicing, customer experience and employee engagement, as well as process reengineering. As a part of this paper, priority use cases for agentic AI have been provided to enable Financial Institutions (FIs) to contextualise and benchmark their own agentic AI journey.

Drawing on Deloitte's AI institute research , our work with leading FS organisations and industry best practices, this paper provides a practical roadmap for **turning agentic AI from concept to competitive advantage.** The purpose of this report is to highlight the core strategic opportunities that FS institutions are capable of unlocking by using AI and agentic AI. It also aims to create a differentiated offering and turn it into a tactical competitive advantage.

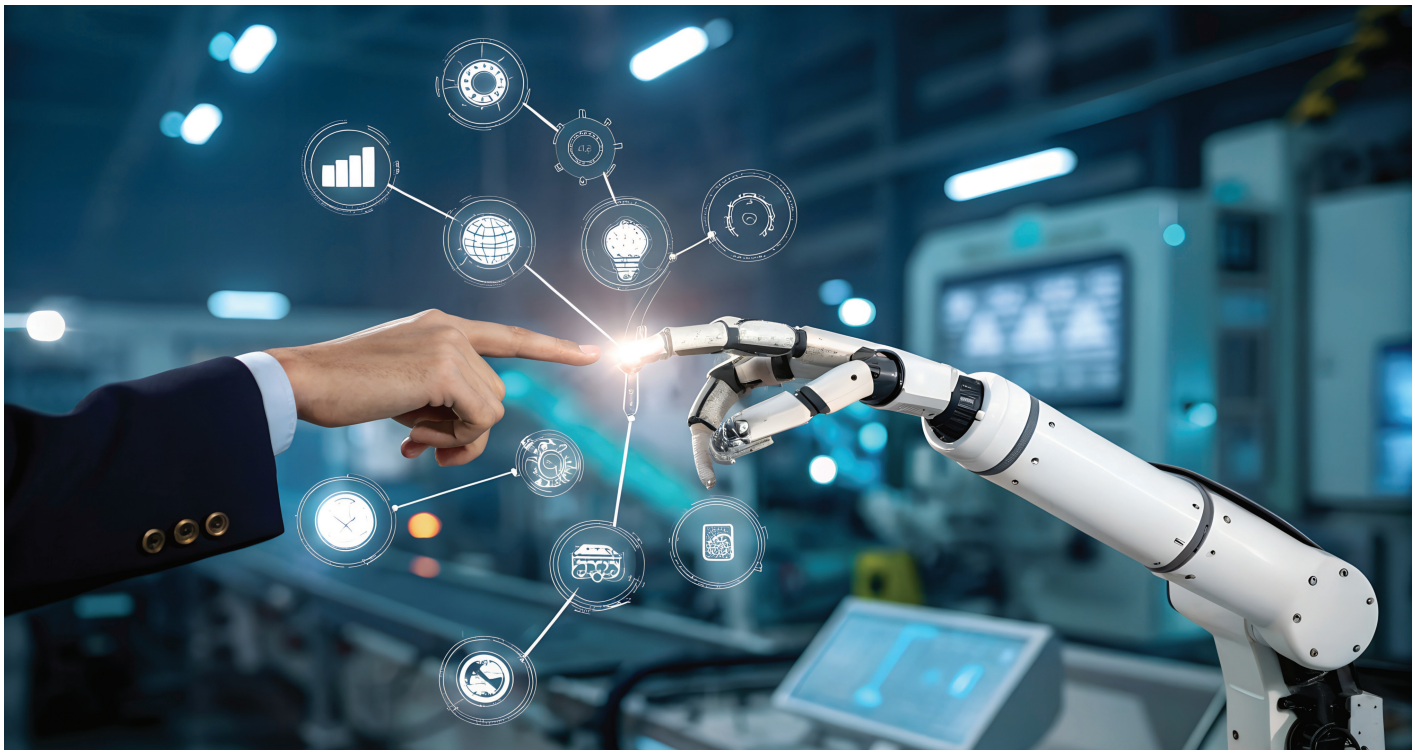

# The state of agentic AI in financial services

The Financial Services (FS) industry is uniquely positioned to lead and benefit from the rapid evolution of AI, in general, and agentic AI, in particular.

Most Financial Institutions (FIs) view agentic AI as technological advancement that is integral to shaping their current operating model. FIs with relatively mature technology stacks, flexible core systems and a reasonable risk appetite are well-positioned to attain **first-mover advantage** and gain a decisive edge over their competitors.

From transaction histories and intelligent customer agents to risk appetites and automated fraud detection, FIs have access to detailed and diverse data sets. Agents' ability to seamlessly engage with and analyse vast amounts of data positions it as a transformative capability for the industry.

Globally, Deloitte predicts that 50 percent of existing companies that have already adopted GenAI will launch agentic AI pilots or proofs of concepts by 2027.[1]

The pace of innovation is increasing, and companies are significantly shifting their priorities regarding investments in AI frameworks that can plan, decide, adapt and work alongside humans and systems.

Therefore, AI agents represent a powerful leap beyond traditional GenAI, equipping systems to not only reason and generate, but also to **act autonomously and collaboratively.** The deployment of agents helps streamline a business' operations by automating front, middle and back-office operations across the enterprise.

A more detailed view on the changing paradigms, autonomy levels, self-adaptive capabilities and other key tenets of AI automation models is defined in Figure 1.

[1] https://www.deloitte.com/us/en/insights/industry/technology/technology-media-and-telecom-predictions.html

**Figure 1. Agentic AI representing the next frontier in automation**

**Advancement in autonomy and reasoning**

| | Traditional RPA Robots | Classic AI models (pre-LLM: rules, ML, NLP) | LLM assistants/copilots | Agentic AI |
|---|---|---|---|---|
| **Primary paradigm** | Scripted UI automation | Pattern recognition models trained for one task | Large-context language generation + use of external functions or APIs | Agents that plan, act, learn and coordinate |
| **Autonomy level** | None (executes fixed steps) | Low (produces scores or text, human embeds action) | Medium (can draft and call APIs but expects user prompt) | High (self-sets subtasks, chooses tools, stops when the goal is met) |
| **Learning/ adaption** | Static | Retrained offline | In-context learning/RAG memory | Continuous feedback loop; fine-tunes heuristics on live signals |
| **Scope of use** | Narrow, repetitive, rules based | Single domain e.g., evaluating credit score | Broad knowledge, multidomain conversation | Cross-workflow orchestration |
| **Decision-making** | None | Outputs probability | Suggests actions | Makes decisions within guard-railed limits; escalates rare cases |
| **Memory across sessions** | None | Model weights only | Short-term conversations | Short-term + long-term memory + shared blackboard across a multiagent environment |
| **Explainability** | High (list of steps) | Medium (dependent on each input or feature) | Medium to low (highly driven by the placement of words in the prompt) | Medium to high (tracks every action and reasoning with logs) |
| **Governance need** | Basic access control and focus on ensuring compliance | Additional controls to ensure ethical use, fairness, transparency, and to manage bias | Added controls such as model risks controls to ensure high accuracy and manage risks of misinformation | Highest level of governance required; must consider agent control room, real-time auditing, action logging, kill switch, and human override |
| **Indicative Use Cases** | • Rule-based Loan Report Generation<br>• Compliance Reporting<br>• Automatically updating Customer Data<br>• Automated Invoice Processing across defined Parameters | • Credit Risk Scoring based on historical data<br>• Pattern recognition for Fraud detection<br>• Sentiment Analysis using NLP (on social media/other feedback mechanisms) | LLM-powered chatbots that can:<br>• Answer customer queries related to offerings<br>• Support employees with summarised data, reports and recommendations | • End to end loan processing including decisions on credit-worthiness, resolution of concerns and final approval<br>• Autonomous fraud detection and response (such as by freezing accounts or contacting customers) |

**Rise in the level and type of risks**

**Source:** Deloitte Center for Financial Services, August 2025

Unlike traditional AI models that rely heavily on human orchestration, agentic AI heralds the possibility of a new heliocentric operating model, one where multiple intelligent agents mimic human collaboration and act as the central execution layer. These agents are managed, and their actions coordinated by supervisory agents, which in turn are supervised and guided by humans. This evolution promises enhanced efficiency, scalability and innovation, allowing FIs to automate complex multi-step workflows by dividing and distributing them into tasks to be executed by specialised agents, mirroring human collaboration and teams.

**Figure 2: Financial Institutions' move towards a helio-centric model with limited human supervision**

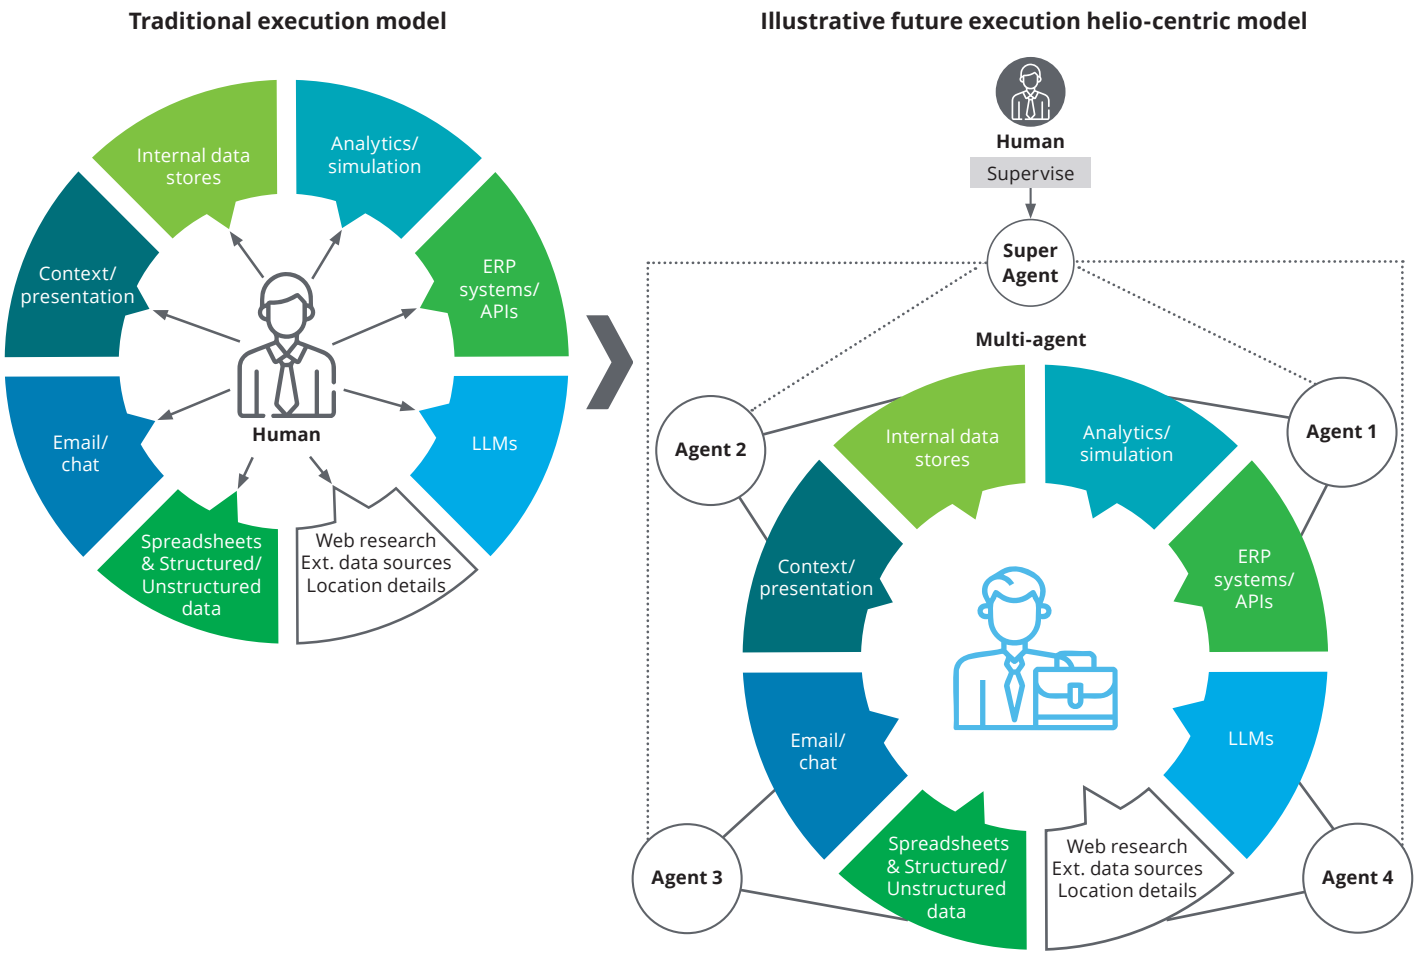


**Source:** Deloitte Center for Financial Services, August 2025

## Key agentic AI value propositions in FS

The FS industry is at the crossroads of a transformative era, where agentic AI is not only maximising operational efficiency but also unlocking innovative use cases across the value chain. With the help of autonomous, goal-driven agents, organisations can enhance customer-facing interactions, streamline back-office processes and create new revenue opportunities. For FIs, agentic AI offers the potential to reimagine offerings across multiple user personas and deliver hyper-personalised experiences throughout the customer lifecycle.

### 1. Multi-agent sales acceleration and customer retention

Agentic AI can enable FIs to identify relevant product offerings, targeted services, complaint patterns and dissatisfied customers. Agents enable targeted resolutions or products to the right customer at the right time. An illustrative flow has been briefly described below.

a. **A Customer retention agent** will track and analyse account-usage patterns to identify disengaged customers. This information and any actionable

insights derived therein may be passed on to a product recommendation agent

b. **A product recommendation agent** will then analyse customer behaviour, financial statements and credit scores to pre-approve them for loans and propose appropriate products.

c. Working in tandem, a **sales engagement agent** will design personalised offers and retention strategies and communicate the same with a communication agent.

d. The **communication agent** will ensure timely outreach through targeted banners, notifications, direct messages or prioritised resolutions.

e. After a customer expresses interest, a **customer support agent** will step in to address queries, provide post-purchase support, review complaint patterns and track dissatisfaction for targeted retention.

f. Finally, a **payment processing agent** will complete the transaction.

Agents collaborating seamlessly across an entire sales journey can enable FIs to accelerate revenue growth through personalised, timely resolutions and product placements, while reducing acquisition costs and strengthening customer lifetime value.

## 2. Automating underwriting

Traditionally, underwriting has been viewed as a complex and time-consuming task, and one prone to human error. With the advent of agentic AI, the process can be completed significantly quicker and with greater accuracy.

a. **Data analyst agents** will analyse relevant financial and non-financial data to flag concerns and arrive at a risk score or classification in real time.

b. Depending on their degree of involvement in the workflow, **advisory agents** may either advise human salespersons on a customer's eligibility and feasible pricing schedules or work with another agent to directly present the offer.

**FIs can deploy agentic AI to empower loan agents with data-backed rationale and recommendations on ranges, including limits, sub-limits and deductibles. Automated underwriting has the potential to save countless person-hours, allowing organisations to drive leaner operations, improving speed-to-market, portfolio quality and cost efficiency, while reducing human bias in lending decisions.**

## 3. Enhanced KYC, next-gen transaction monitoring and investigations

Moving beyond anomaly detection, the unique proposition of agentic AI lies in the potential of **investigative agents** that can autonomously contextualise and escalate suspicious activity.

A **Monitoring agent** identifies unusual transactions, and an **investigation agent** evaluates them against risk thresholds in real time by collating expected financial behaviours, historical records and contextual information from Customer relationship management (CRM). A **compliance agent** determines potential escalation paths and recommends solutions for commonly observed anomalies.

Recent research shows that sophisticated deepfake and spoofing techniques, such as manipulated video streams and virtual-camera injections, can bypass many commercial electronic **Know Your Customer** (eKYC) systems, including those with active liveness detection. AI agents can help solve this by analysing contextual signals, such as how a device is being held or whether unauthorised software is running.

A combination of a **document processing agent** to extract and validate information from uploaded unstructured data, such as IDs and proofs, with a **watchlist screening agent** to continuously monitor against sanctions and fraud databases, can help FIs set up a faster, frictionless onboarding process with lower compliance risks.

**A large Dutch Financial Institution has been using a combination of AI innovations for its KYC and compliance processes, reportedly achieving a 90 percent reduction in onboarding time and cutting staff workload by 30 percent.**

Beyond tracking transactions and KYC, FIs with high-quality data sets can also take the support of agents to create dynamic risk thresholds based on the changing customer profile and automate regulatory disclosures or other periodic reporting exercises. **Such risk management agents** track updates to regulations through integration of multiple data sources, minimising false positives in AML and screening for potential sanctions, ensuring regulatory compliance.

**A large American financial institution is using LLMs to help determine what information clients must report to regulators[2] and how they can improve their business processes.**

[2] https://www.bloomberg.com/news/features/2023-05-31/jpmorgan-s-push-into-finance-ai-has-wall-street-rushing-to-catch-up

**Another world-leading FI is using agentic AI to both reduce fraud[3] and improve portfolio management. Its systems track device activity, customer behaviour and transaction patterns in real time to detect fraud and trigger AML alerts, reducing false positives and accelerating investigations.**

## 4. Operational excellence and employee productivity

Agents can transform internal operations for Fis and also boost employee productivity by automating routine tasks, making internal information easier to access and responding to queries in real-time.

The ability to self-learn and take autonomous action also suggests that the greater the adoption of such models across the enterprise, the more effective they are likely to become.

a. **For instance, an American financial institution uses an agent trained on their internal corporate data. It has boosted staff productivity by combing through relevant internal procedures across HR, Compliance, Finance etc., and providing responses to concerns, complaints, and queries in real time.**

b. The role of AI agents observed across non-FS industries makes them imminently transferable to FS. Agents are being used to author and review code, hastening development efforts to build new platforms and enhance existing ones. **One such tool is the Software Development Life cycle (AI-SDLC). The AI-SDLC embeds Gen AI and agentic capabilities across the software lifecycle, from requirement gathering and code creation to testing and deployment.**

c. Organisations have been using simple chatbots to resolve basic customer service requests, but AI agents now enable customer service automation to evolve into a more robust, contextualised service that can truly mimic human interactions in terms of tonality, proactiveness and sentiment analysis. **Another American Financial Institution's employee facing agent is used by ~90% of its workforce.[4] The agent has the capability to automate common support queries and fetch responses in real time, thereby avoiding manual escalations and reducing calls to the Human-IT desk by more than half. The Bank also offers a customer-facing version of the chatbot. Here, the chatbot acts as a 24/7, on-demand customer assistant that enhances customer experience and boosts brand loyalty. The Agent has transformed how customers engage with the bank, with over 2 billion registered interactions, and an average of ~2 million interactions per day.[5]**

d. Across global finance, investment banks and ecosystem vendors are rolling out AI-driven solutions that empower employees to make faster, data-driven decisions. **Notable examples include an American financial institution's LLM. Trained on 50 Bn+ parameters[6] across sentiment analysis, named entity recognition, news classification, and question answering, this model is tailored to the FS domain. Another example comes from a reputed financial research firm, whose flagship investment analysis tool leverages advanced algorithms to assess a startup's growth prospects and accordingly create exit predictions.**

Agentic AI has the potential to redefine core banking processes across growth (sales), resilience (risk/compliance) and trust (customer experience). For CXOs, the imperative is clear:

1. Capture quick wins through onboarding and underwriting pilots.
2. Build enterprise-wide orchestration capabilities that link agents across business units.
3. Shape governance and risk frameworks early to ensure trust, transparency and regulatory alignment.

[3] https://www.jpmorgan.com/insights/payments/payments-optimization/ai-payments-efficiency-fraud-reduction
[4] https://newsroom.bankofamerica.com/content/newsroom/press-releases/2025/04/ai-adoption-by-bofa-s-global-workforce-improves-productivity--cl.html
[5] https://newsroom.bankofamerica.com/content/newsroom/press-releases/2024/04/bofa-s-erica-surpasses-2-billion-interactions--helping-42-millio.html
[6] https://www.bloomberg.com/company/press/bloomberggpt-50-billion-parameter-llm-tuned-finance/

A view of some of Deloitte's recent work related to agentic AI in Financial Services can be found in Figure 3.

**Figure 3: Deloitte's work with financial services clients**

### Market sensing and reporting for treasury at a large Indian Bank

- **What:** Automated financial reporting by Integrating multi-source financial data, real time insights, and contextual news processing
- **How:** Use of agents for news summarization, de-duplication, categorization, validation, and integration of expert insights
- **Why:** Reduced manual efforts, improved operational efficiency and better financial reporting

### GL Scrutiny for a leading bank

- **What:** Automation and review of general ledger entries to detect anomalies and errors
- **How:** Use of a multi-agentic workflow which Involves agents for data extraction, validation, anomaly detection, exception management
- **Why:** Increased accuracy, reduced manual effort

### Tax reconciliation for a leading Indian bank

- **What:** Reconciliation of tax records received from various internal sources and external Govt tax portals
- **How:** Using a series of agents for data extraction, validation, orchestration, fuzzy watch and data consolidation
- **Why:** Improved reconciliation, enhanced visibility, and timely resolution

### Customer feedback assist tool for a leading Canadian bank

- **What:** Automation of customer feedback and review analysis
- **How:** Use of GPT-4 turbo for deriving insights from unstructured customer reviews of cilent's app from various app hosting platforms
- **Why:** Ensure consistency in the analysis by reducing human biases while reducing human efforts by automating the tag assignments

### Customer dispute resolution for a leading payments company

- **What:** Automate and intelligently manage customer payment disputes through end-to-end AI-driven case handling
- **How:** Leveraging Agentic AI to analyse dispute data, engage with customers via chatbots, and coordinate with backend teams to resolve issues swiftly
- **Why:** Reduced resolution time, improved customer satisfaction

### CAM creation at one of India's largest private banks

- **What:** Agentic AI solution to automate the process of CAM creation and update
- **How:** Deploying multi-agent solution including drafting agent, data extraction & summarization, identification agent, deviations approval agent
- **Why:** Efficient and quicker CAM update process

### Automated bank vetting process for leading private bank

- **What:** Aim to reduce turnaround time & human intervention in the bank guarantee wetting process
- **How:** Use agentic AI bots to validate, modify details, and generate updated BG documents
- **Why:** An autonomous, efficient process with minimal human intervention for BG text

### RM nudge engine for a leading corporate bank

- **What:** Agentic nudges and notifications to RMs for better engagement with clients and increased
- **How:** Muti-agent solution focusing on prioritization of clients, next best actions, scheduling, and optimization routes/channels
  **Why:** Focused targeting of customers with more focus on cross-selling and up-selling of products, reduced rejection rate

## Case in point: Intelligent decision making by an AI system for loan processing

Banks can transform the customer underwriting process by using a combination of GenAI, agentic AI and other advanced digital technologies. With these tools, banks can streamline their operations and enhance the fairness, efficiency and security of the underwriting process.

**Figure 4: Intelligent AI assisted loan processing Journey map**

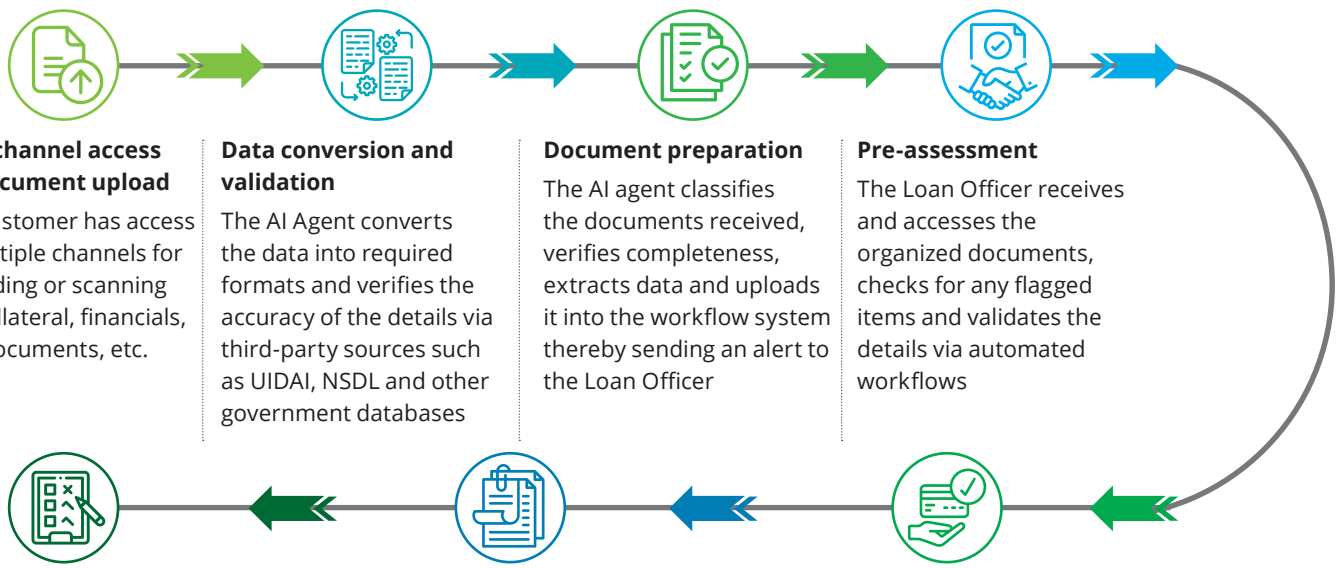

**Multichannel access for document upload**

The customer has access to multiple channels for uploading or scanning the collateral, financials, KYC documents, etc.

**Data conversion and validation**

The AI Agent converts the data into required formats and verifies the accuracy of the details via third-party sources such as UIDAI, NSDL and other government databases

**Document preparation**

The AI agent classifies the documents received, verifies completeness, extracts data and uploads it into the workflow system thereby sending an alert to the Loan Officer

**Pre-assessment**

The Loan Officer receives and accesses the organized documents, checks for any flagged items and validates the details via automated workflows

**Credit assessment**

The credit officer pulls the financial and other data using conversational AI to analyze trends and history to estimate probability of defaults or other risks

**Memo creation**

The AI Agent generates a credit memo and further makes a credit decision and generates the loan terms automatically based on the credit assessment

**Contract generation**

The AI Agent independently generates the contract based on the loan terms customized to the applicant profile ensuring the loan structure is aligned with both the borrower's and the bank's criteria

The use of AI can offer the following benefits to organisations strategically leveraging it for various use cases:

**Faster decision making**

The AI agents can instantly summarise and analyse the data and documents and make quick decisions, improving overall turnaround time.

**Proactive risk management**

AI-based predictive analytical agents can foresee potential defaults or credit risks based on customer behaviour patterns, enabling proactive decision-making.

**Improved customer experience**

AI-powered chatbots and agents provide instant responses and personalised interactions, improving the customer journey from application to approval.

**Enhanced efficiency and scalability**

Agentic AI and automation allow FIs to scale their underwriting processes to handle a larger volume of applications, without sacrificing quality and saving on time and effort.

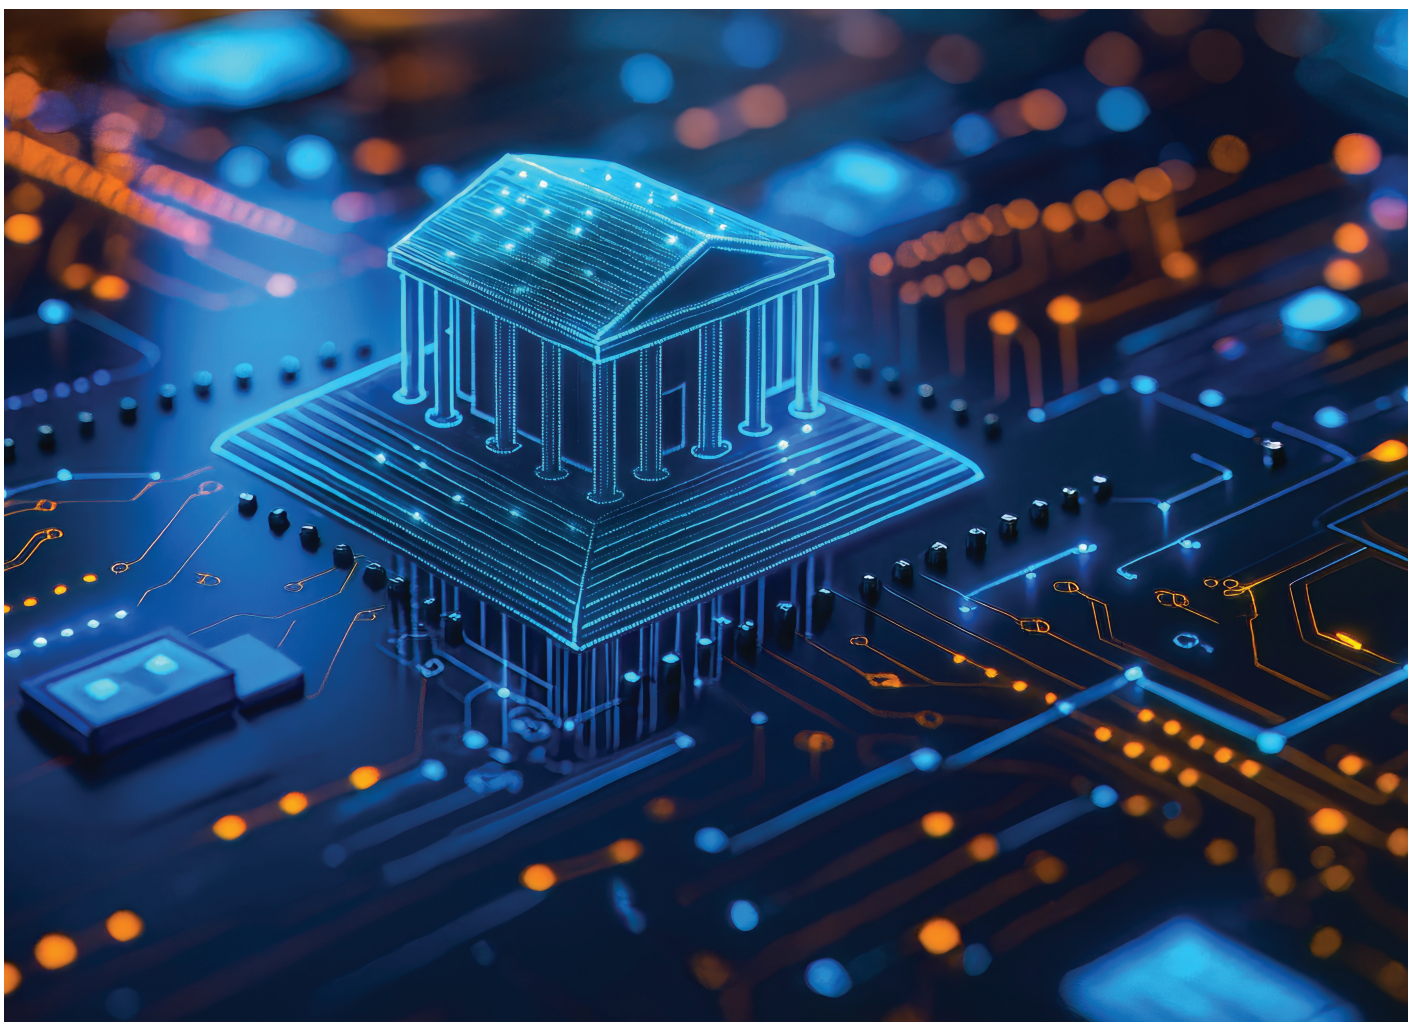

# Adopting agentic AI: From proof of concept to scale commercialisation

Indian FIs are set to reap the demographic dividend (higher digital literacy), digital dividend (in terms of accelerated digitization driven by initiatives such as UPI-enabled payment infrastructure, UIDAI, etc.) and data dividend (in terms of ability to triangulate customer insights). In this context, agentic AI represents a transformational opportunity. A strategic and well-structured approach will enable the agentic AI journey to be more deliberate and value-enhancing.

As Indian FS institutions embrace the transformative potential of Agentic AI, they need to evaluate the following four dimensions in their journey:

## 1. Which business processes / use cases are best suited for agentic AI adoption and what are the likely risks and outcomes?

FIs make a key decision when settling on the scope and complexity of agents that are to be deployed.

The simplest form of deployment may be referred to as an ***overlay,*** a scenario where there is minimal change in the business process. Instead, an agent is introduced to bolster the existing process by providing efficiency and additional value with minimal disruption.

A second approach is one where FIs build software applications specifically for Agentic AI. These applications are hyper-targeted, often focusing on individual workflows within offerings and are ***agentic by design.*** FIs may incorporate these applications within existing processes, facilitating targeted changes without an overhaul of processes.

Finally, in line with quantum scaling, FIs may elect to introduce agentic AI alongside an overhaul of existing processes. Rather than mere deployment of agents, the overhaul approach would require an assessment of existing capabilities, an identification and prioritisation of opportunities for agentic AI introduction, and finally, a ***redesign of processes*** that organically includes agents as integral parts of the workflow.[7]

[7] https://www.deloitte.com/us/en/insights/industry/financial-services/agentic-ai-banking.html

A few considerations that shape the agentic approach for FIs are as follows:

**Figure 5: Key considerations in deciding the approach to agentic AI**

| | Consideration | Potential Approach (smart overlay, agentic by design or process redesign) |
|---|---|---|
| **Process Clarity** | **How clearly defined and standardized is the process for this task?**<br>Clear, partially defined, ambiguous | **Clear:** Smart overlay (quick results), or agentic by design for incremental modernization<br>**Partially defined:** Agentic by design recommended, but smart overlay may still be viable initially<br>**Ambiguous:** Consider process redesign to fundamentally reshape workflow |
| **System Openness** | **How easily can your core systems connect and communicate via APSs or event streams?**<br>Highly connected, moderately connected, disconnected | **Highly connected:** Appropriate for smart overlay or rapid deployment of agentic by design<br>**Moderately connected:** Prioritize agentic by design for flexible integration<br>**Disconnected:** Begin full-scale process redesign or comprehensive agentic by design replacement |
| **Elastic Compute** | **How flexibly can you scale cloud infrastructure during high-demand periods?**<br>Fully elastic, somewhat elastic, rigid | **Fully elastic:** Ideally suited for process redesign or comprehensive agentic by design<br>**Somewhat elastic:** Start incrementally with agentic by design<br>**Rigid:** Begin with smart overlay; progressively enhance infrastructure |
| **Degree of Risks** | **What is the severity of the potential impact from a wrong decision by an AI agent?**<br>Mission critical, moderate impact, low stakes | **Mission critical:** Prioritize agentic by design or strategic process redesign for strong governance<br>**Moderate impact:** Agentic by design preferred; implement smart overlay initially<br>**Low stakes:** Safely start with smart overlay to gain immediate benefits |
| **Speed to Value** | **How quickly does the agent need to demonstrate measurable benefits?**<br>Immediately (within months), within a year, medium to long term (one to two or more years) | **Immediate:** Leverage smart overlay<br>**Within a year:** Thoughtfully adopt agentic by design, evolving gradually<br>**Medium to long term:** Ideal for holistic process redesign and deep agentic by design |
| **Data readiness** | **How mature is your data management and governance?**<br>Robust and ready, developing, early stage | **Robust and ready:** Ideal for process redesign and comprehensive agentic by design<br>**Developing:** Start with agentic by design selectively<br>**Early Stage:** Begin with smart overlay to progressively build capabilities |

| | Consideration | Potential Approach (smart overlay, agentic by design or process redesign) |
|---|---|---|
| **Talent Strength** | **What is your current strength in AI talent?**<br>Ready, growing, limited | **Ready:** Ideal for process redesign and agentic by design<br>**Growing:** Start with agentic by design selectively<br>**Limited:** Begin with smart overley to progressively build capabilities |
| **Replacement vs Integration cost** | **Considering money and time, how cost-effective is integration with legacy systems versus replacing them?**<br>Cheaper to integrate, balanced cost, cheaper to replace | **Cheaper to integrate:** Start with smart overlay, then phase in agentic by design over time<br>**Balanced cost:** Use agentic by design selectively<br>**Cheaper to replace:** Immediately pursue agentic by design, consider process redesign where justified |
| **Regulatory Considerations** | **How mature are your regulatory and compliance automation frameworks?**<br>Highly automated, partially automated, mostly manual | **Highly automated:** Strong position for agentic by design and eventually process redesign<br>**Partially automated:** Proceed carefully with agentic by design<br>**Mostly manual:** Begin conservatively with smart overlay to gradually mature compliance frameworks |

**Note:** API: Application Programming Interface
**Source:** Deloitte Center for Financial Services analysis

## 2. How can agentic AI initiatives be designed to complement existing infrastructure, digital, governance and talent-readiness

Beyond technology initiatives, agentic AI needs to be viewed as a business transformation opportunity. To enable successful adoption of agentic AI, FIs need to look at it through a people, process and technology lens:

- **People (culture and talent):** A top-down direction to uplift firmwide AI literacy will enable different lines of businesses to adopt relevant agentic use cases across key business priorities. Companies need to develop in-house abilities in data science, engineering and ML operations to build and scale AI. Additionally, a clear set of roles and escalations need to be defined for deployed agents.
- **Process:** To enable seamless adoption of agents, companies need to identify key workflows (ERP, CRM, compliance, ops systems etc.) that agentic AI can plug into. Banks need to set up a sandbox or Centre of Excellence (CoE) for piloting emerging AI technologies. Each pilot must be measured against agreed-upon business and technology-related Key Performance Indicators (KPIs).
- **Technology:**
  - **Data readiness:** Companies need to clearly define data standards, lineage, consent and retention policies, which would be the foundation for agentic AI implementation readiness. Pipelines to transform unstructured data, along with data quality monitoring to prevent bad inputs, are essential for adequate outputs to be generated.
  - **Infrastructure:** To ensure dynamic accessibility, cloud adoption and scalable deployment, pipelines are critical. This, clubbed with a flexible core, API-driven microservice architecture, low-code orchestration layers, and digital platforms capable of supporting autonomous, swappable agents across the entire journey, will enable true adoption of AI across core banking processes.

## 3. What is the best approach to scale, incremental, quantum or hybrid? What is the best way to deploy agentic AI: buy/partner, build or hybrid?

The sheer volume and scope of agentic AI may, at first, seem overwhelming to FIs. To avoid resource misallocation and ballooning costs, they must adopt a personalised approach suited to their specific needs, business objectives and technological maturity.

## Deploying agentic AI

A detailed view contrasting three broad strategic approaches to agentic AI deployment is defined in Figure 6.

**Figure 6: The Build vs Buy/Partner vs Hybrid approach**

| Criteria | Build | Partner | Hybrid |
|---|---|---|---|
| **Strategic value** | When Agentic AI is core-to your IP, product or long-term strategy | When Agentic AI is important but requires high customisation and domain-specific capabilities | When Agentic AI is not core, off-the-shelf solutions will suffice |
| **Data protection** | Data security is a significant ethical concern | Data can be shared with external partners | Only internal data can be shared |
| **Talent Availability** | Strong in-house data/AI teams available | Some expertise but seek external augmentation | Limited AI/tech expertise internally |
| **Time-to-market pressure** | High upfront investment; better long-term ROI at scale | Shared cost/risk model with scalable roadmap | Instant deployment and quick time-to-value are top priorities |
| **Cost and Scalability** | Longer lead time, full control | Collaborative model that allows for co-investment, minimising risk and maximising value | Prefer solutions with set budgets and implementation schedules |

**Source:** The business imperative for Agentic AI, July 2025

Each approach has its own set of risks and benefits, and none of the deployment strategies may be viewed as objectively superior to the others. Therefore, the "right" approach for an organisation is the one that makes the most business sense based on its specific organisational context.

## Scaling up agentic AI

Agentic AI strategies may begin by focusing on initial deployment, but in the medium to long term, this shifts towards effective scaling of offerings. As business demands grow, so do the opportunities to accrue benefits from the deployment of additional agents. The potential for increased earnings and efficiency is also accompanied by a higher strain on technology architecture and more complex compliance scenarios. These considerations must be measured by FIs as they seek to settle on an appropriate approach. FIs must maintain a fine balance between innovation, risk, resource allocation and compliance.

## Incremental

Incremental scale-up refers to an approach that emphasizes gradual expansion of agentic AI in the organisation. This allows organisations to adopt a careful approach, keeping considerations around safety, security, compliance and efficacy at the forefront. In addition to simplifying regulatory compliance, this will also facilitate gradual staff adoption and training, reducing resistance and errors, particularly in mission-critical processes such as underwriting or fraud detection. A potential drawback may be that in a rapidly evolving market, an FI may risk losing ground to competitors using quicker scale-up approaches.

## Big bang approach

This approach focuses on large-scale and simultaneous adoption of Agentic AI across the enterprise. It provides an instant edge over competitors, pushing the organisation to the forefront of innovation and the benefits derived therein. However, the approach also comes with drawbacks. First, executing a change of this significance at such a large scale will exert unprecedented pressure on existing technology infrastructure. Additionally, at a time when the regulatory environment is dynamic and regulators are focused on ensuring robust governance, the complexity of and organisational unfamiliarity with the changed state of agentic AI would present significant compliance challenges.

## Hybrid

A hybrid approach tries to find an attainable middle ground between the incremental and big bang approaches. This approach can systematically balance innovation with organisational risk, facilitating quick adoption across priority areas while ensuring FIs do not over-extend their capabilities. A considered approach would also allow for greater flexibility to react quickly in case of market or regulatory shifts. However, a hybrid approach is not without its cons. It significantly complicates strategy and resource allocation efforts. Teams are required to maintain legacy structures, while also learning and incorporating new technologies into their everyday workflows. This acts as a drain on resources with both rising maintenance costs and overstretched teams.

# 4. What is the governance end game?

The ideal governance model for effective agentic AI implementation would be one that balances innovation and trust within the boundaries of regulatory requirements and is auditor-ready. Policies for explainability, bias, security, model approval, and release management are to be embedded in relevant systems before scaling up with clear escalation paths. AI ethics guidelines, along with proactive regulatory compliance, are essential.

**Policy and ethics:** Establishment of standards for fairness, explainability, acceptable use and documented decision policies per use case.

**Human-in-Command architecture:** Even with autonomous abilities, adequate human intervention is needed for high-stakes actions, such as credit approvals and large payments. Protocols must be established to determine when agents should halt and escalate to humans for sign-off.

**Regulatory alignment and compliance by design:** Organisations should ensure that the agents to be embedded must automatically comply with RBI's digital lending norms, SEBI's algo-trading guardrails and data localisation requirements. They must also have the ability to dynamically adjust processes based on evolving RBI/SEBI/IRDAI circulars (e.g. on data privacy, customer consent and algorithmic bias). Every decision by an AI agent must be traceable, explainable and reversible, ensuring regulators can conduct post-facto reviews.

**Future-state end game:** The ideal future-ready state would be of a self-regulating ecosystem in which AI agents can monitor each other to guard against compliance breaches.

**Regulator-ready dashboards:** Real-time supervisory dashboards accessible to RBI/SEBI to monitor AI decisioning at the systemic level.

## Positioning Agentic AI across the Banking Value Chain

Implementing agentic AI solutions is likely to accrue benefit across an enterprise. Deloitte has identified four key pillars of the banking value chain that may be used to categorise use cases and their intended business objectives. These are:

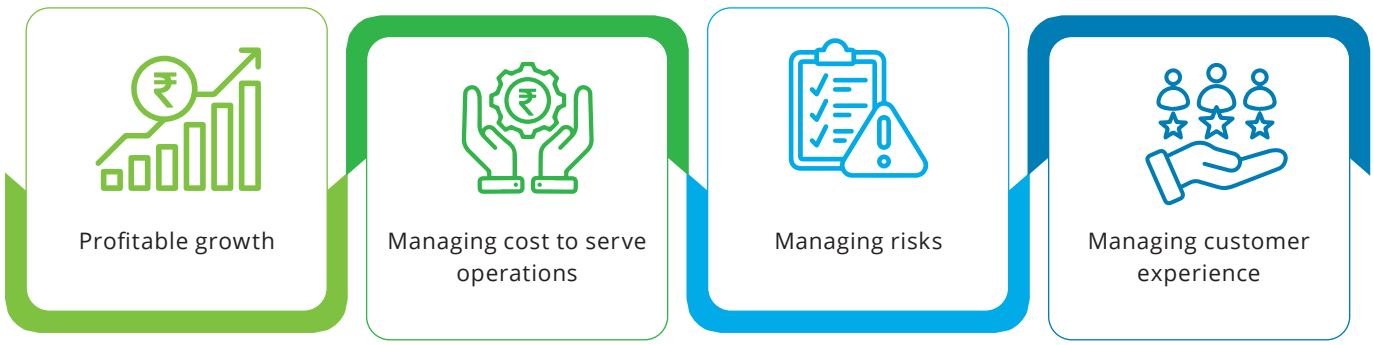


**Figure 7: Key agentic AI use cases across the banking value chain**

| Profitable growth | Managing cost to serve operations | Managing risks | Managing customer experience |
|---|---|---|---|
| Agent-powered Loan Approval ● | Chatbot Agents – Customer support, website, IVR, in-app ● ● | Predictive Modelling and Risk ● | Chatbot Agents ● ● |
| Automated Credit Underwriting through Alternative Data ● ● | Customer-facing and inward-looking Virtual Assistant Agent ● ● | Profiling Agent ● | Customer-facing Virtual Assistant Agent ● ● |
| Financial Advisory Agent ● ● | Multi-Agent Client Onboarding Workflows ● ● | Transaction Monitoring Agent ● | Dispute Management Agents ● ● |
| Optimized Tax Planner ● ● | Payments Processing Agents ● ● | Agent-led Automated KYC ● ● | Multi-Agent Client Onboarding Workflows ● ● |
| Agent recommended product innovation ● ● | Agent Generated Financial Reporting ● | Automated Compliance and Reporting Agents ● | |
| Trade Planning and Execution Agents ● | Process Automation Agents ● | | |
| Agent-driven Trade Finance ● | Software Development Assistance Agents ● | | |
| | Contract/Document Summary and Analysis Agents ● | | |
| | Agent-run IT Helpdesks ● | | |

● Retail
● Corporate
● Back-office

A comprehensive view of this categorisation of agentic AI use cases can be found in Figure 7.

# Conclusion

Agentic AI has rapidly evolved from a technology enabler to a **strategic imperative** for FIs worldwide. Today, leading FS players view AI and agentic AI not merely as an IT upgrade, but as a transformational force reshaping core business functions. Early mover FIs are using agentic AI to move beyond narrow process automation use cases and building intelligence into the core of their operations.

The benefits of agentic AI are likely to be multi-dimensional: from optimised cost-to-income ratios and productivity gains to accelerated go-to-market (GTM) strategies and enhanced customer engagement. With robust governance and strategic deployment, agentic AI can revolutionise FI operations and drive operational excellence and business growth. In the next two to three years, agentic AI is likely to drastically optimise the cost-to-income ratios of banks, NBFCs, insurers and wealth management firms. Depending on the use cases, it is also likely to have an impact on revenues and cross sales.

FIs that invest in Agentic AI are likely to see an increase in employee productivity with the support of AI agents across back-office, mid-office and customer-facing roles. Sales and marketing in FIs will benefit from faster turnaround times, driven by more effective content creation and the ability of agentic technologies to respond swiftly to market conditions and in routine sales and servicing activities.

To unlock the full potential of agentic AI capabilities, FIs must place AI at the centre of their talent agenda. There is a need to create organisation-wide agentic AI fluency through internal capability development initiatives. Capability enhancements will help employees to utilise AI tools effectively, aiding productivity gains. As the FS Industry inevitably pivots towards integration with agents, fostering an AI-first way of thinking will provide FIs with a competitive edge in the medium to long-term.

Given early detection and collection efforts and corrective actions, undertaken by agentic technologies; there is a long-term potential for reduced slippages and delinquencies, creating highly robust balance sheets and giving AI-first organisations a competitive edge.

Despite its potential, agentic AI continues to struggle with preconceived biases and hallucinations. At each stage of evolution, these risks will have to be mitigated, including robust data testing as well as regular data audits, feedback loops and testing of AI outputs.

Agentic AI systems are likely to **accelerate from proof of concepts to commercialisation in the next 12 to 18 months.** The journey for each FI is likely to be different as the agentic AI use case really depends more on the maturity of the FIs on data, cloud, AI infrastructure and capabilities and less on the technology of Agentic AI.

AI-first organisations will set the benchmark for speed, scale and customer centricity in Financial Services. Those who invest in agentic AI not as a technology upgrade, but as a catalyst for business model reinvention by blending human supervision with autonomous intelligence to deliver measurable enterprise value, will emerge as winners.

# Sources

**World Economic Forum.** *Artificial Intelligence in Financial Services 2025. Jan. 2025.* https://reports.weforum.org/docs/WEF_Artificial_Intelligence_in_Financial_Services_2025.pdf

**Deloitte.** *State of Generative AI in the Enterprise: Fourth Quarter 2024.* Dec. 2024. https://www.deloitte.com/content/dam/assets-zone3/us/en/docs/services/consulting/2024/us-state-of-gen-ai-q4.pdf

**Citi.** "Agentic AI: The Future of Intelligent Financial Systems." *Citi Global Perspectives & Solutions,* 2024. https://www.citigroup.com/global/insights/agentic-ai

**Salesforce.** "AI in Banking." *Salesforce India, 2024.* https://www.salesforce.com/in/financial-services/ai-in-banking/

**Citi.** *Agentic AI: The Future of Intelligent Financial Systems.* Citi GPS: Global Perspectives & Solutions, 2024. https://www.citigroup.com/rcs/citigpa/storage/public/GPS%20Report_Agentic%20AI.pdf

**Deloitte.** *Generative AI Use Cases Across Industries:* May 2025. May 2025. https://www.deloitte.com/content/dam/assets-zone3/us/en/docs/services/consulting/2025/us-ai-institute-gen-ai-use-cases-2025-05-08.pdf

**Finastra.** "Agentic AI Assistants in Banking." Finastra, 2024. https://www.finastra.com/viewpoints/articles/agentic-ai-assistants-in-banking

**BNY Mellon.** "How Generative AI Is Transforming Trade Finance." *BNY Insights,* 2024. https://www.bny.com/corporate/global/en/insights/trade-finance-digital-transformation.html

**Deloitte India.** "India Rides the Agentic AI Wave." *Deloitte Press Room,* 2025. https://www.deloitte.com/in/en/about/press-room/india-rides-the-agentic-ai-wave.html

**Stanford Institute** for Human-Centered Artificial Intelligence. *AI Index Report 2025.* 2025, https://hai.stanford.edu/assets/files/hai_ai_index_report_2025.pdf.

**Deloitte China.** "DeepSeek Industry Applications." *Deloitte AI Institute,* 2025, http://deloitte.com/cn/zh/issues/digital/deloitte-ai-institute/deepseek-industry-applications.html.

**Gartner.** "Gartner Predicts Over 40 Percent of Agentic AI Projects Will Be Cancelled by End of 2027." *Gartner Newsroom,* 25 June 2025, https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027.

**Gartner.** *Agentic AI Strategic Planning Guide.* 2025, https://www.gartner.com/document-reader/document/5629291?ref=pubsite.

**"RBI** Announces Free AI Committee to Develop AI Framework." *The Economic Times,* 30 July 2025, https://economictimes.indiatimes.com/news/economy/policy/rbi-announces-free-ai-committee-to-develop-ai-framework/articleshow/116684195.cms?from=mdr.

**Bank of America.** "AI Adoption by BofA's Global Workforce Improves Productivity." *Bank of America Newsroom,* 23 Apr. 2025, https://newsroom.bankofamerica.com/content/newsroom/press-releases/2025/04/ai-adoption-by-bofa-s-global-workforce-improves-productivity--cl.html.

**Bank of America.** "BofA's Erica Surpasses 2 Billion Interactions, Helping 42 Million Clients." *Bank of America Newsroom*, 13 Apr. 2024, https://newsroom.bankofamerica.com/content/newsroom/press-releases/2024/04/bofa-s-erica-surpasses-2-billion-interactions--helping-42-millio.html.

**Deloitte Digital.** *Agent Advantage for Salesforce Financial Services.* 23 Apr. 2025, https://www.deloittedigital.com/content/dam/digital/global/documents/accelerator-20250423-agent-advantage-for-salesforce-financial-services.pdf.

**Haptik.** "AI Agents: Use Cases in Banking." *Haptik Blog, 2025,* https://www.haptik.ai/blog/ai-agents-use-cases-in-banking.

**Deloitte.** "Generative AI in Investment Banking." *Deloitte Insights, 2025,* https://www.deloitte.com/us/en/insights/industry/financial-services/generative-ai-in-investment-banking.html.

**Deloitte.** "AI Consulting Services." *Deloitte Consulting, 2025,* https://www.deloitte.com/us/en/services/consulting/services/ai-consulting.html.

**Digitalisation World:** https://digitalisationworld.com/news/69250/majority-of-cios-and-ctos-view-agentic-ai-as-core-to-future-it-operations

# Glossary

| | |
|---|---|
| AI | Artificial Intelligence |
| AML | Anti Money Laundering |
| API | Application Programming Interface |
| BG | Bank Guarantee |
| BoFA | Bank of America |
| CoE | Centre of Excellence |
| CRM | Customer Relationship Management |
| eKYC | Electronic Know Your Customer |
| ERP | Enterprise Resource Planning |
| FI | Financial Institutions |
| GenAI | Generative Artificial Intelligence |
| GPT | Generative Pre-trained Transformer |
| GTM | Go-to-Market |
| IT | Information Technology |
| IP | Intellectual Property |
| IRDAI | Insurance Regulatory and Development Authority of India |
| KYC | Know Your Customer |
| LLM | Large Language Models |
| ML | Machine Learning |
| NLP | Natural Language Processing |
| ML | Machine Learning |
| RBI | Reserve Bank of India |
| RPA | Robotic Process Automation |
| RAG | Retrieval-Augmented Generation |
| ROI | Return on Investment |
| SDLC | Software Development Lifecycle |
| SEBI | Securities and Exchange Board of India |
| UI | User Interface |
| UIDAI | Unique Identification Authority of India |
| VC | Venture Capital |

# Connect with us

**Monish Shah**
Partner
Financial Services Industry Leader,
Technology & Transformation
Deloitte South Asia
monishshah@deloitte.com

**Ashvin Vellody**
Partner
Technology & Transformation
Deloitte South Asia
ashvinv@deloitte.com

**Sanjay Kumar Gupta**
Partner
Financial Services
Technology & Transformation
Deloitte South Asia
sanjaygupta@deloitte.com

**Shweta Shetty**
Partner
Financial Services
Technology & Transformation
Deloitte South Asia
sshweta@deloitte.com

# Contributors

**Vijay Mani**
**Moumita Sarkar**
**Nilesh Jajoo**
**Narayanan Raman**
**Priyank Gandhi**
**Sooday Jhaveri**
**Aneesh Fernandes**
**Aryaman Sood**
**Vritti Bang**

# Acknowledgements

**Dinesh V**
**Monica Bhattacharya**
**Aashumi Sheth**
**Chaitanya Kapadia**
**Param Sharma**
**Yashika Darbari**
**Shahzada Allam**
**Sampreeti Sen**
**Shreya Rajopadhye**

Deloitte.

Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited ("DTTL"), its global network of member firms, and their related entities (collectively, the "Deloitte organization"). DTTL (also referred to as "Deloitte Global") and each of its member firms and related entities are legally separate and independent entities, which cannot obligate or bind each other in respect of third parties. DTTL and each DTTL member firm and related entity is liable only for its own acts and omissions, and not those of each other. DTTL does not provide services to clients. Please see www.deloitte.com/about to learn more.

Deloitte Asia Pacific Limited is a company limited by guarantee and a member firm of DTTL. Members of Deloitte Asia Pacific Limited and their related entities, each of which is a separate and independent legal entity, provide services from more than 100 cities across the region, including Auckland, Bangkok, Beijing, Bengaluru, Hanoi, Hong Kong, Jakarta, Kuala Lumpur, Manila, Melbourne, Mumbai, New Delhi, Osaka, Seoul, Shanghai, Singapore, Sydney, Taipei and Tokyo.

This communication contains general information only, and none of DTTL, its global network of member firms or their related entities is, by means of this communication, rendering professional advice or services. Before making any decision or taking any action that may affect your finances or your business, you should consult a qualified professional adviser.

No representations, warranties or undertakings (express or implied) are given as to the accuracy or completeness of the information in this communication, and none of DTTL, its member firms, related entities, employees or agents shall be liable or responsible for any loss or damage whatsoever arising directly or indirectly in connection with any person relying on this communication.

© 2025 Deloitte Touche Tohmatsu India LLP. Member of Deloitte Touche Tohmatsu Limited