Rise of agentic AI: Shaping the future of financial services

October 2025

Table of contents

Foreword

Dear readers,

It is with pleasure that we bring to you the latest edition of our newsletter which explores the use of agentic artificial intelligence (AI) solutions for financial services.

This newsletter focuses on the evolution of AI and agentic AI and how they can be leveraged to enhance the financial ecosystem to make it more efficient.

We hope you find this newsletter to be a useful and insightful read.

For further details or feedback, please write to:

Vivek Belgavi (vivek.belgavi@pwc.com) or
Mihir Gandhi (mihir.gandhi@pwc.com).

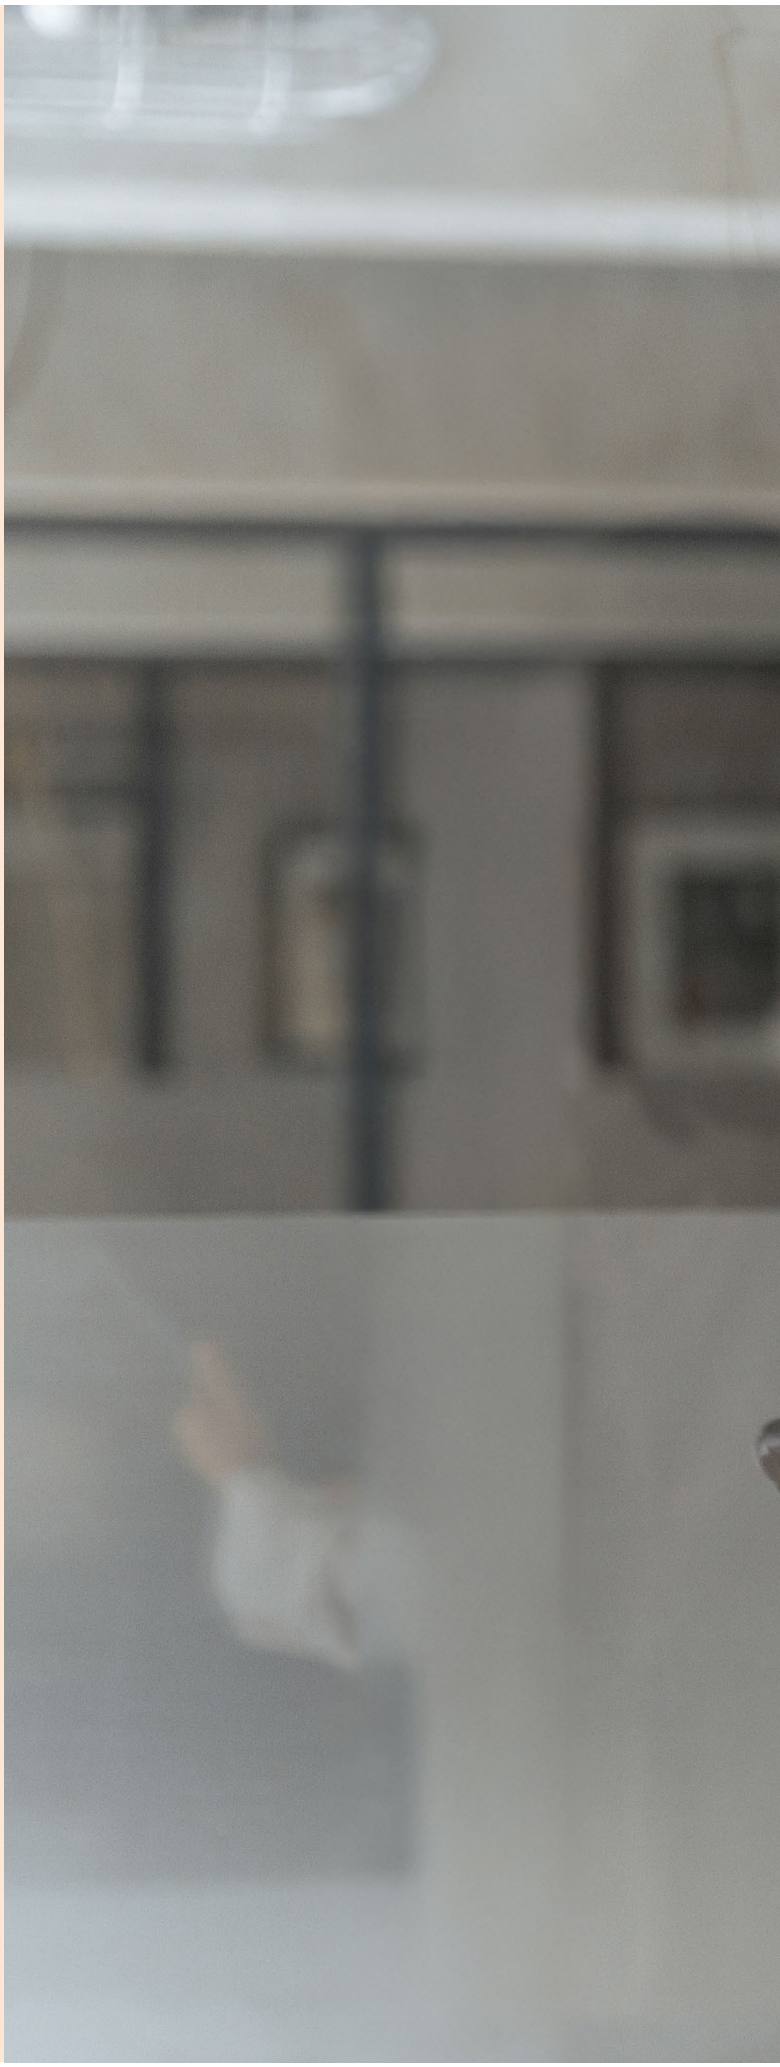

01 Introduction

What is agentic AI?

Agentic AI refers to AI systems that can autonomously perform tasks involving multiple steps while making context-based decisions with minimal manual oversight. These systems combine multiple large language models (LLMs) and utilise their reasoning abilities with multiple servers which enables them to execute tasks independently through frameworks such as model context protocol (MCP). This allows agentic AI systems to work as self-directed 'digital workers'.

How it works:

Foundation in LLMs: Agentic AI uses LLMs as its cognitive engine for tasks like analysing unstructured data and generating reports. However, unlike standard LLM applications that create and analyse content, agentic AI adds an additional layer of decision-making.

Retrieval using vector databases: Vector databases allow data storage as highly multi-dimensional data in mathematical form using semantic similarity to store and retrieve data. Agentic AI requires access to relevant and context-rich information, which becomes possible using semantic search.

Action via MCP: MCP allows AI agents to access and exchange information with external tools and servers. AI agents can identify and call configured tools dynamically to pass contextual information and collect the response.

Multi-agent collaboration: Complex workflows are managed by teams of multiple specialised AI agents that interact and work together.

Figure 1 illustrates how agentic AI can be used for a bank which wants to collect and process loan applications.

Figure 1: Illustration of multi-agent orchestration for loan processing

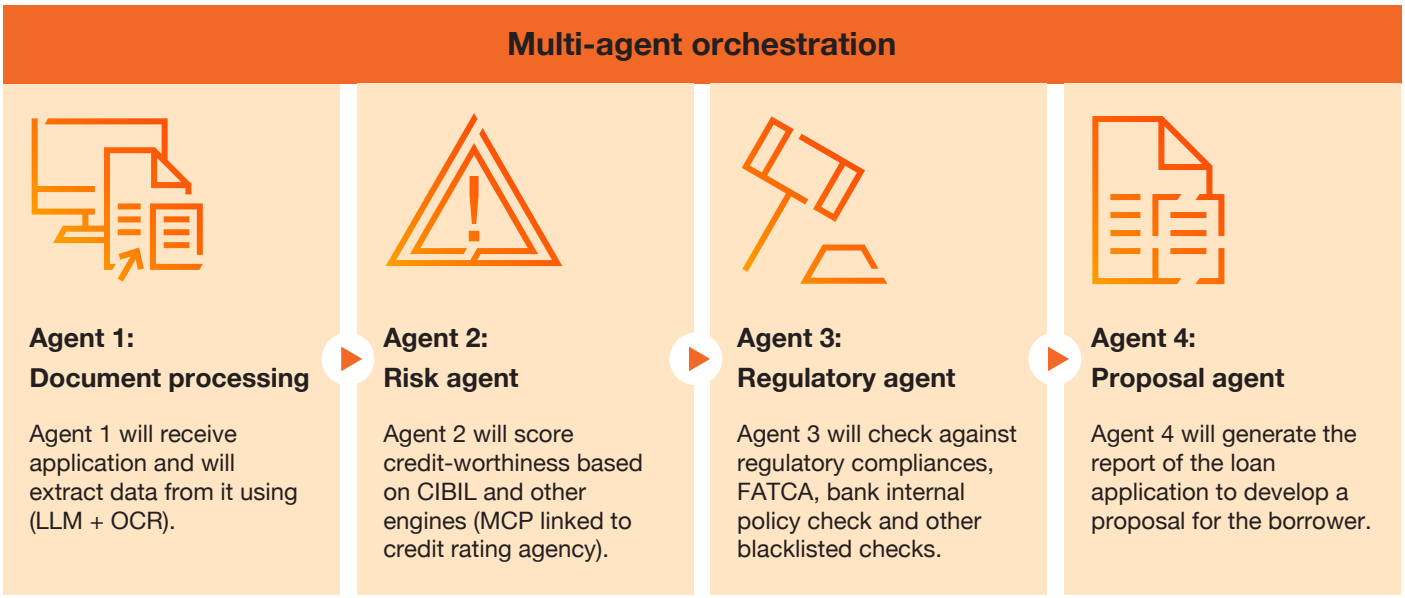

OCR: Optical character recognition, FATCA: Foreign Account Tax Compliance Act, CIBIL: Credit Information Bureau (India) Limited

Source: PwC analysis

Comparison: Generative AI and agentic AI

Agentic AI is like a new frontier to GenAI capabilities, enabling decision-making and execution capabilities.

Figure 2: GenAI vs agentic AI – a comparison

	GenAI	Agentic AI
Main function	Creates new content — text, images, code or simulations — based on learned data patterns.	Autonomously plan, make decisions and take actions to achieve desired outcomes.
Trigger mechanism	Prompt-based, requires human inputs	Operates on goals, environmental signals or triggers
Decision capability	Provides output-based on patterns but lacks decision-making and validation capabilities.	Performs action based on context and autonomous decision making
Workflow and collaboration	Singlular model for content generation within workflow	Performs workflow orchestration for multi-step processes and collaborates with multiple agents for execution
Maturity	Emerging yet rapidly adopted in customer-facing areas	Early-stage adoption in decision-heavy processes. ~33% of enterprise applications are expected to integrate Agentic AI by 2028
Integration complexity	Moderate – can be embedded in existing tools	High – requires orchestration across systems
BFSI use cases	Chatbots, document generation, customer profiling, etc.	Credit underwriting, call-centre voice support, debt recovery, compliance monitoring, etc.

Source: PwC analysis

Core principles of agentic AI

Agentic AI systems have four key characteristics that allow them to function as autonomous, intelligent digital workers in financial services and beyond.

Figure 3: Core principles of agentic AI

Autonomy

- Designed to carry out complex workflows independently, minimising manual intervention
- Utilises pre-programmed logic combined with machine learning (ML)
- Enables real-time decisions without constant human oversight

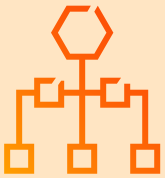

Autonomy

- Centrally orchestrates actions using diverse tools and skills
- Transforms complex tasks into manageable sub-tasks
- Determines the most efficient plan to reach the goal

Goal-oriented collaboration

- Optimises actions across a team of agents to achieve specific outcomes
- Coordinates end-to-end workflows for better efficiency
- Focuses on objectives to ensure collaboration success

Contextual decision-making

- Operates across multiple data streams, dynamically adapting to changes
- Responds to current context, new information, or unexpected events like adapting to market volatility or regulatory shifts

Continuous learning

- Improves decision models from every interaction
- Incorporates human feedback (e.g. overriding a trade transaction) to refine processes
- Supports reinforcement learning from human feedback (RLHF) for ongoing enhancement

Source: PwC analysis

02
Adoption of agentic AI

With rapid growth of digital transformation in the banking and financial services industry, agentic AI is also gaining traction with many companies experimenting with integrating the technology with their solutions in this space. In the Indian market, financial institutions are focusing on breaking down their operations into smaller use cases and then use agentic AI-led solutions rather than adopting the technology for enterprise-wide solutions.

2.1 Key milestones and breakthroughs

In recent years, AI has evolved from being rule-based systems that required heavy human input to becoming autonomous, agentic AI systems. Several critical developments have accelerated the rise of agentic AI in financial services as illustrated in Figure 4.

Figure 4: The evolution of AI (2020 – 2025)

2020	Growth of LLMs	Decision-making	Workflow orchestration and collaborative agents	Emergence of MCP	2025
	Being the foundation of agentic AI, LLMs experienced significant advancement between 2020 and 2023, with the adoption of vector databases, enhancing their ability to understand, research, generate and summarise unstructured data.	LLMs gained decision-making capabilities that enabled them to select next step in the path of execution without human intervention in its selection.	This phase is marked a significant progress for execution of complex multi-step processes. This enables AI systems to take decisions autonomously to orchestrate the workflow for multiple AI agents assigned for specific sub-task, thereby increasing the efficiency of system.	MCP is recent development which enhances agentic system to autonomously communicate with external systems/ customer relationship management (CRM)/databases via configured tools to get the task done. It allows the AI system to identify, select and use tools autonomously.	
	Example: AI bots can understand user's loan request, interact to collect completed application forms and collateral related documents, and summarise required data elements in structured format for processing.	**Example:** Based on data collected from users, the AI system can make decision related to which loans (personal/agricultural/home loan) are to be processed and initiate the processes accordingly.	**Example:** When a user loan request comes, it assigns different processes like customer profiling, document verification, risk analysis, proposal drafting to different AI agents that work in collaboration based on interactive communication between them.	**Example:** During the loan process, to get the risk profiling done, an AI system can be configured to autonomously seek the CIBIL score from CIBIL's system or verify PAN details from National Securities Depository Limited (NSDL) by passing the required information to the respective system.	

Source: PwC analysis

2.2 Emerging themes in banking and financial services institutions (BFSI)

Agentic AI is evolving from point solutions to enterprise-wide capabilities across the front, mid and back office — shifting from isolated pilots to integrated, goal-driven systems that plan, act and learn with built-in guardrails.

There are a variety of emerging themes across categories of solutions including those catering to the front, mid and back office functions of financial institutions. Vibe coding functionality could be considered as a horizontal one that caters to multiple functionalities.

While agentic AI presents a wide array of opportunities, Figure 5 provides a non-exhaustive list of use cases based on the complexity and maturity of its adoption.

Figure 5: Maturity and complexity of AI's use cases

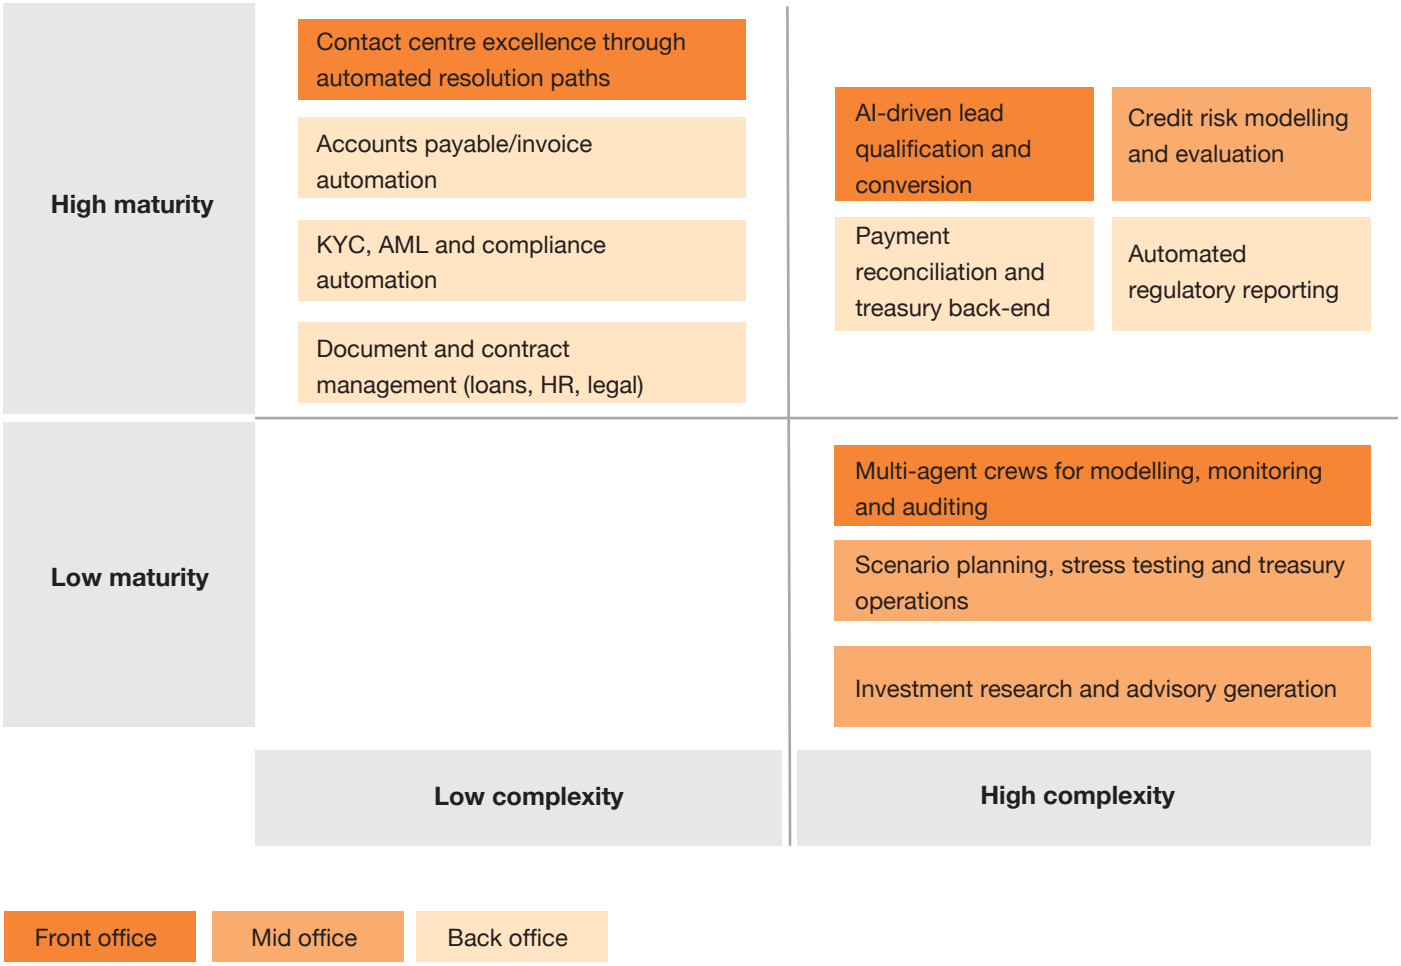

Source: PwC analysis

High maturity and low complexity (plug-and-play solutions):

Established, off-the-shelf solutions with numerous vendors and standard APIs for easy integration.

High maturity and high complexity (enterprise-scale integrations):

Well-understood technologies that provide significant value but require deep integration with core systems and major project management effort.

Low maturity and high complexity (cutting-edge and experimental):

High-risk, high-reward research and development (R&D) initiatives. The technology is nascent and implementation is complex, often requiring custom development.

A. Back office use cases (Operations, compliance and audit)

Adoption of agentic AI in back office operations can help transform high-volume, rules-intensive processes by integrating document understanding, policy execution and tool orchestration. These AI agents can bring about high straight-through processing rates, manage exceptions autonomously and enable continuous optimisation.

1. Accounts payable/invoice automation: Agentic AI solutions can process invoices across formats (PDF, image, electronic data interchange (EDI)), validate them against purchase orders, check tax rules, schedule payments and handle exceptions — all without human input unless required.

Example:

- A bank's internal procurement team receiving thousands of invoices could leverage agentic AI systems to streamline extraction, matching, payment batch generation and reconciliation. This could potentially reduce the average turnaround time from multiple days to under a few hours.
- Potential benefits of this could include a reduction in invoice processing time by up to 80%,[1] compliance with vendor payment schedules and real-time visibility for finance teams.

2. Document and contract management (loans, HR, legal): Agentic AI solutions can scan contract, extract key fields (e.g. dates, clauses and parties), verify completeness and raise alerts for missing signatures or abnormal clauses.

Example:

- A financial institution managing loan origination process may utilise agentic AI solutions to read large agreements, auto-validate interest rate clauses and cross-reference with term sheets and CRM records.
- Potential benefits of this could include faster onboarding with approval from different departments (legal/ risk), lowering the operational risks, and assuring document accuracy.

3. Know your customer (KYC), anti-money laundering (AML) and regulatory compliance automation: Agentic AI agents can autonomously collect KYC documents, validate identity across multiple databases (e.g. PAN, Aadhaar, sanction lists), and flag suspicious activity using pattern recognition. Unlike the current automation layer, which follows static, rules-based scripts, an agentic layer can adapt dynamically to new regulations, learn from evolving fraud patterns and operate with greater decision-making autonomy, making it more effective in compliance-heavy environments.

Example:

- An agentic AI system can be utilised to scan transactions for anomalous behaviour (e.g. structuring, mule accounts), then auto-generate suspicious activity reports (SARs) with justification and route them to a compliance officer.
- Potential benefits of this could include faster reviews with up to 60%[2] drop in manual reviews, proactive risk flagging and timely filing to regulators.

1 https://economictimes.indiatimes.com/tech/technology/bpm-firms-see-ai-automating-traditional-work-but-boosting-high-tech-jobs/articleshow/118221515.cms

2 https://timesofindia.indiatimes.com/technology/tech-news/how-agentic-ai-can-fix-the-legal-industry/articleshow/123110705.cms

B. Mid-office use cases (Strategy, analytics, risk management)

In mid-office functions, agentic AI solutions can accelerate decision cycles by integrating data ingestion, scenario reasoning and policy governance. It delivers explainable recommendations for functions such as credit, treasury and risk management while ensuring regulatory compliance.

1. Credit risk modelling and evaluation: Agentic AI systems can co-ordinate multiple specialised agents for ingesting credit bureau data, behavioural analytics and for creating or validating ML models to automate and improve risk assessments.

Example:

- In retail lending, AI agents may pull data from internal CRM, social credit scores, utility payment behaviour and external credit bureaus. They can dynamically simulate scenarios such as interest rate changes and customer job loss to recommend safer lending thresholds.
- Potential benefits of this could include faster underwriting for SMEs and retail loans with adaptive real-time risk scores with improved credit decisions.

2. Scenario planning, stress testing and treasury operations: Agentic AI systems can integrate market data such as oil prices, consumer price index (CPI) and currency fluctuations, historical performance, and internal liquidity ratios to conduct stress tests and simulate treasury scenarios like funding gaps or foreign exchange (FX) exposure.

Example:

- For its asset-liability management, a bank can utilise agentic AI agents to run a large number of parallel scenarios on regulatory capital adequacy in case of a sharp interest rate hike and then recommend rebalancing of asset portfolios.
- Potential benefits of this could include better hedging and capital plannings, faster and better stress testing, enabling autonomous decision-support during crisis.

3. Investment research and advisory generation: Agentic AI systems can pull structured and unstructured data (e.g. news, earnings reports, analyst calls, etc.) to generate real-time insights and produce research or pitchbooks personalised for the clients.

Example:

- A wealth advisor platform could use agentic AI solutions to generate client-ready market summaries and exchange traded funds (ETF) recommendations based on the latest macroeconomic data and investor goals, with compliance language auto-embedded.
- Potential benefits of this could include increase in productivity for analysts, real-time updates on portfolios and scalable personalisation for high net-worth individuals (HNIs).

C. Front office use cases

Agentic AI enhances customer engagement through intent recognition, contextual data retrieval and autonomous, compliant action execution. It enables faster resolution, proactive outreach and consistent omni-channel experiences that drive both service excellence and revenue growth.

1. Customer experience – contact centre excellence through automated resolution paths: Agentic AI systems can enhance customer service for customer issue resolution across multiple channels such as voice, chat and email. These systems can leverage contextual and natural language processing to interpret customer intent, resolve issues without human oversight and escalate matters only when needed.

Example:

- An organisation can enable its support channels by introducing an AI-based assistant that understands the context of a customer's question and pulls up the right data instantly. Agentic AI system can manage queries over both chat and voice, and adapt to customer's history and sentiment, making the overall experience more responsive.
- Potential benefits of this could include improved first-contact resolution with 70–90% containment for common queries, lower handling time and cost per contact, and up to 17% higher customer satisfaction.[3]

2. AI-driven lead qualification and conversion: Agentic AI systems continuously scan CRM systems, social media, and transaction data to identify high-value prospects, qualify them, and initiate first-level engagement through personalised outreach.

Example:

- A bank's wealth management team may use such systems to flag clients with sudden liquidity events (e.g. asset sales) and automatically schedule a call with a senior advisor, sharing a pre-built proposal.

- Potential benefits of this could include reduced lead conversion time, higher prospect engagement, and increased relationship manager productivity.

3. Multi-agent crews for modelling, monitoring and auditing: Agentic AI can deploy specialised agent crews where multiple agents can collaborate in real-time to interpret credit model insights, monitor risk and market impact signals, manage documentation, audit readiness and regulatory compliance to support complex decision making.

Example:

- Financial institutions can deploy an agent crew during credit approval with one agent explaining the model score and rationale and another checking for anomalies against market and customer signals, while a compliance agent prepares an audit-ready report, ensuring transparency before approval is granted.
- Potential benefits could include faster, more reliable client-facing decisions, improved compliance with audit-ready outputs and reduced model risk through along with collaborative monitoring.

3 https://cio.economictimes.indiatimes.com/news/artificial-intelligence/generative-and-agentic-ai-set-to-transform-customer-service-into-a-strategic-value-driver-for-businesses/119114383

D. Vibe coding across enterprise functions

Vibe coding empowers teams to design and deploy AI agents rapidly using composable skills, human-in-the-loop review and secure integration via MCP. It accelerates innovation while embedding governance and reusability across functions.

Workflow automation: Vibe coding allows cross-functional teams to design intelligent workflows by composing reusable agentic skills such as document parsing, validation and decision logic, into end-to-end automation pipelines. These workflows can span across business functions and integrate with enterprise systems via secure APIs.

Example:

- A business operations team can use vibe coding to automate the internal approval process for new product launches. Agents can extract key inputs from business requirement documents, validate them against policy checklists, notify relevant stakeholders and trigger downstream actions such as system configuration or compliance review, ensuring timely, policy-aligned execution.
- Potential benefits could include improved cross-functional alignment and compliance, with reduction in TAT for development and customisation, and reduction of dependence on technology team.

Figure 6 provides an example of loan processing in a financial institution.

Figure 6: Loan processing – traditional process vs AI-driven process

Loan origination	Loan processing	Loan servicing	Loan collection
Traditional process (How it happens): 1. Customer approaches bank/ agent to apply for loan. 2. Eligibility checks are performed manually. 3. Paper-based forms are filled in which are prone to errors. 4. Physical meetings, phone calls and panel discussions for clarifications 5. Manual KYC verification and credit bureau checks 6. Underwriting decisions take longer due to manual turnaround.	**Traditional process (How it happens):** 1. Customer queries are handled by call centres. 2. Agents resolve issues manually leading to long wait times 3. Limited self-service options for customers 4. High operational cost due to manpower requirements	**Traditional process (How it happens):** 1. Loan agreements prepared manually 2. Multiple rounds of compliance checks can cause delays 3. Repayments tracked through manual accounting 4. Reconciliation requires manual adjustments 5. Inconsistent follow-up on delays	**Traditional process (How it happens):** 1. Payment follow-ups are done manually (calls, physical visits) 2. Reminders sent late or inconsistently 3. High collection costs due to manpower involvement 4. Escalation handled only after overdue builds up
AI enhancements (On top of the process): • AI-driven screening and instant credit scoring for faster eligibility check. • OCRs digitise and validate forms to reduce errors and turnaround time. • Chatbots/conversational AI engage customers instantly and resolve FAQs. • Automated KYC (optical character recognition, face match, liveness detection) reduces manual dependency. • Predictive risk analysis to assess creditworthiness. • Personalised engagement for better customer experience.	**AI enhancements (On top of the process):** • AI-powered self-service (chatbots, IVR, virtual assistants) • Real-time resolution of simple queries, reducing dependency on call centres • Sentiment analysis helps in identifying unhappy customers quickly • AI agents suggest insights for cross-sell and up-sell opportunities • Reduced cost and improved customer satisfaction	**AI enhancements (On top of the process):** • AI decision engines recommend approvals based on pre-set policy rules • ML predictive models reduce manual interventions • AI predicts defaults before they occur, allowing proactive action • Automated multi-channel repayment reminders (text messages, calls) • Faster reconciliation with AI-driven matching	**AI enhancements (On top of the process):** • Predictive AI identifies early risk indicators for overdue customers • Risk-based and rule-based collection prioritisation • Automated smart reminders trigger fund release or repayment nudges • Escalation to human officers only for high-risk or complex cases • Lower cost, faster recoveries and reduced defaults

Source: PwC analysis

03

Agentic AI in BFSI – a case study

A leading BFSI group adopted agentic, multilingual voice and document processing AI to enhance service, onboarding and collections — achieving scale, accuracy and cost efficiency. Below are some of the impact areas from the implementation of agentic AI solutions.

Customer onboarding feedback automation:

- Automated 50,000+ daily customer calls with 65%+ connect rates,[4] retaining context over multiple attempts and triggering follow up actions directly in core systems.
- This helped resolve:
  - High volume manual-calling for feedback from newly onboarded customers, resulting in data leakages.
  - Inconsistent feedback collection, limiting actionable insights for process improvement.
  - Lack of automated routing to relevant teams for timely resolution and corrective action.

Merchant onboarding and KYC:

- Verified 100,000+ merchants annually across 300+ document formats and multiple languages at 95%+ accuracy[5] — slashing lead times and manual QA through seamless CRM integration.
- This helped resolve the following challenges:
  - Large amount of time and efforts involved in manual verification of over thousands of monthly applications by agents.
  - Quality assurance and automation gaps causing delays and scalability issues
  - Diverse document formats and languages complicating resource allocation and processing.

Collections outreach:

- Contacted 300,000+ overdue customers monthly via intelligent voice bots, securing repayment intent from 13%+[6] of connected calls while cutting cost to collect.
- This helped resolve challenges like:
  - Customers forgetting re-payment due to forgetfulness, lack of reminders or missing information.
  - The unscalable nature of manual collection calls and inconsistencies associated with them.

The outcome:

Consistent, policy compliant engagement at national scale, faster turnaround on critical processes and measurable uplift in operational efficiency demonstrates how agentic AI helps organisations in transforming their processes from process automation to intelligent, end to end orchestration.

4 PwC analysis
5 Ibid
6 Ibid

04
Considerations and challenges in implementation

A. Application and drivers of adoption

How banks can integrate agentic AI

The financial services industry is changing fast and, unlike traditional AI, agentic AI can learn, adjust to new situations and work across systems with minimal input which results in greater efficiency, smarter decision-making and more personalised customer experiences.

Drivers of adoption: Cost efficiency, customer experience and compliance

Agentic AI is gaining traction across financial services, driven by its ability to improve operational efficiency, enhance customer engagement and ensure stronger regulatory compliance.

B. Implementation considerations

Advanced AI tools alone are not enough and success depends on strong planning, oversight and integration. Since autonomous AI agents act independently, organisations must enforce governance, compliance and risk controls to ensure adherence to regulations, operational safety and fitment with existing processes. This safeguards daily operations while keeping AI reliable and trustworthy. Some of the aspects an organisation must consider before adopting agentic AI technology for their operations are:

1. **Governance frameworks:** Set up oversight teams to monitor agent behaviour, review key decisions and run risk assessments pre- and post-deployment. Build systems with secure access, legal safeguards and ongoing monitoring, supported by training the staff for quick response to anomalies.
2. **Compliance and regulatory alignment:** Align with international and local regulatory norms such as the Digital Personal Data Protection (DPDP) Act, General Data Protection Regulation (GDPR) and Health Insurance Portability and Accountability Act (HIPAA) through legal reviews, system logs, updated policies and vendor contracts that clarify roles, responsibilities and escalation protocols.
3. **Ethical AI and risk management:** Establish an ethics panel with legal, technical and policy experts to ensure fairness, transparency and societal responsibility which extends beyond compliance.
4. **Scalability:** Assess infrastructure and data silos that hinder integration. Upgrade legacy IT and computing resources to support autonomous AI at scale.
5. **Uncertainty:** Continuously monitor evolving AI behaviour, adapting governance and risk controls to manage unintended outcomes, bias and model risks.
6. **People:** Address skill gaps in BFSI through targeted hiring and upskilling. Counter resistance with transparent communication and inclusive change programmes.
7. **Initial model training:** In practical scenarios, most data is scattered and unstructured, requiring months of human effort to clean, organise and train agentic AI systems so they can move beyond demos and perform reliably in real-world environments

05
Way forward

Agentic AI is transforming BFSI through autonomous, multi-agent systems powered by MCP. With widespread adoption across front, mid, and back office, agentic AI is assisting BFSI organisations in aspects such as scalable orchestration, contextual intelligence, and compliance-ready automation, signaling a shift from pilots to enterprise-grade integration across the financial ecosystem. Going ahead, the influence of agentic AI would be a dominant theme across the BFSI space, and the industry expects Agentic AI to significantly impact the payments landscape.

PwC's agentic AI practice

At PwC, we aim to facilitate and support organisations in the BFSI sector in their transformation journey – from training their workforce on agentic AI to helping them identify relevant use-cases, streamline their workflows and develop customised AI agents for their business needs.

Our AI tools make work easier by creating business requirement documents (BRDs) and business process model and notation (BPMN) diagrams, predicting demand and modelling future operations. Built for strategy, operations and compliance teams, they save time, reduce manual effort and help drive faster, smoother transformation across the enterprise.

About PwC

We help you build trust so you can boldly reinvent

At PwC, we help clients build trust and reinvent so they can turn complexity into competitive advantage. We're a tech-forward, people-empowered network with more than 370,000 people in 149 countries. Across assurance, tax and legal, deals and consulting we help build, accelerate and sustain momentum. Find out more at www.pwc.com.

PwC refers to the PwC network and/or one or more of its member firms, each of which is a separate legal entity. Please see www.pwc.com/structure for further details.

For more information, please contact:

Vivel Belgavi
Partner and Leader – Financial Services Advisory
PwC India
vivek.belgavi@pwc.com

Mahesh Parab
Partner – Agentic Automation
PwC India
mahesh.parab@pwc.com

Mihir Gandhi
Partner and Leader – Payments Transformation and FinTech
PwC India
mihir.gandhi@pwc.com

Geetika Raheja
Partner – Payments Transformation and FinTech
PwC India
geetika.raheja@pwc.com

Zubin Tafti
Partner – Payments Transformation and FinTech
PwC India
zubin.tafti@pwc.com

Contributors

Vignesh Venkataraman, Pranav Gupta, Mayank Desai and **Rahul Chemburkar**

Editor: Rubina Malhotra
Design: Harshpal Singh

pwc.in

HS/October 2025 - M&C 48925