Whitepaper

Redefining Intelligent Banking with Agentic AI

A Practical Guide to Autonomous, Trusted, and Scalable Banking Transformation

Authors

Gopal Parasnis

Saurabh Agrawal

Executive Summary

- Today, agentic AI is emerging as a transformative force in financial services. It is shifting the industry beyond generative AI's content-creation capabilities to autonomous, goal-driven intelligence that can understand objectives, orchestrate workflows, and execute complex, end-to-end banking processes.
- Yet, the adoption of Agentic AI remains highly inconsistent. While global spending on agentic AI is expected to hit $52.62 billion in 2030 and the financial services segment alone is forecast to grow at 41% CAGR, only 11% of companies have moved agents into production. Especially in banking, very few institutions have deployed it at scale, while many others remain stuck in pilot mode.
- This whitepaper bridges that gap. It examines practical ways banks can responsibly deploy agentic AI to redesign front-, middle-, and back-office operations. It explores proven governance models, practical architectures, human-in-the-loop approaches, and high-value use cases that help banks minimize costs, increase efficiency, strengthen compliance, and deliver exceptional customer experiences at scale.

Table of Contents

- Executive Summary 02
- Agentic AI in Banking: Shift from Assistance to Execution 04
- Key Challenges in Agentic AI adoption 05
- The Urgency of Agentic AI Implementation in Banking 06
- The Agentic Architecture: Front, Mid, and Back Office Integration 07
- Human-in-the-Loop Approach for Agentic AI Implementation 09
- Prerequisites for the Successful Implementation of Agentic AI Solutions 10
- Tech Mahindra's Proprietary Platform for Developing Use Cases 11
- Responsible AI Adoption in Banking 15
- Prioritization Framework 14
- Strategic Roadmap for Agentic AI Adoption in Banking 16
- Conclusion 18

Agentic AI in Banking: *Shift from Assistance to Execution*

Most banking leaders today understand the broad concept of agentic AI. However, few fully grasp how fundamentally different it is from the generative AI tools they are currently using.

Generative AI excels at assistance, drafting emails, summarizing documents, or generating reports. Agentic AI goes far beyond that. It is autonomous, goal-oriented intelligence that doesn't just respond but understands objectives, plans workflows, makes decisions, interacts with multiple systems, and executes complex, end-to-end processes with minimal human intervention.

Traditional AI systems were task-specific and rigid. They performed narrow functions but lacked the ability to adapt, orchestrate, or act independently. In contrast, agentic AI agents act more like intelligent digital colleagues, capable of handling entire workflows such as loan processing, fraud investigation, or portfolio rebalancing from start to finish.

While the underlying technology is advancing at a remarkable speed with new models and frameworks emerging almost monthly, actual adoption and deployment in banks remain surprisingly slow. Many institutions are still experimenting rather than integrating agentic capabilities into their core operations.

Key Challenges in Agentic AI adoption

Despite the clear potential of agentic AI, banks face several significant hurdles in adoption. The most common barriers include uncertain ROI and high upfront investments, a shortage of responsible AI implementation expertise, ambiguous data availability and quality, inadequate security, and a lack of talent and skill sets required for AI adoption, all of which hinder rapid technological advancement in banks.

To better understand these realities, Tech Mahindra conducted a comprehensive survey of 150 senior banking executives across APAC, the Americas, Europe, and the Nordics.4 The survey reveals how banks are currently approaching agentic AI adoption, where they are directing their investments and, most importantly, what truly separates the leaders from the laggards.

One in 4 banks are investing just enough while 1 in 3 are investing heaving in AI

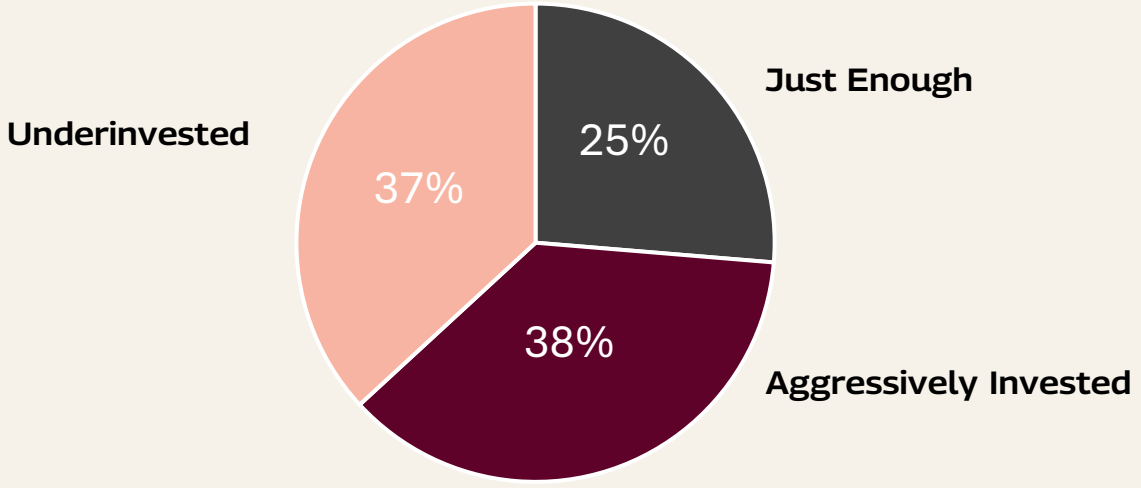

Figure 1: Key Findings from Building the AI-Driven Bank of Tomorrow Survey

The Urgency of Agentic AI Implementation in Banking

Agentic AI, together with generative AI, has matured into a powerful set of tools that are fundamentally reshaping key areas of banking. Right from customer experience (CX) and operational efficiency to risk management, compliance, and product innovation, Agentic AI is delivering new possibilities.

Today's banking reality also demands these capabilities. The current digital-native customers are looking for personalized, proactive financial experiences delivered in real time. Banks that adopt agentic AI early can dramatically elevate the customer journey, shortening onboarding timelines by up to 80%, [1.1]eliminating friction, and achieving near straight-through processing (STP). These institutions are not only delighting customers but also gaining a clear competitive edge.

The same principle applies to efficiency improvements in mid-office and back-end processes. With agentic AI, banks see 50-60% productivity improvements in key processes, making them more cost-effective. This is further echoed in global banks, which have deployed agentic AI for fraud detection, reducing false positives by 30-40% while improving compliance and fraud-detection accuracy.

The message is clear: The window of opportunity is narrowing. Banks that act decisively will unlock significant gains in efficiency, resilience, and customer loyalty. Those that hesitate risk losing market share to faster-moving competitors in an ecosystem increasingly defined by intelligence and speed.

The Agentic Architecture: *Front, Mid, and Back Office Integration*

As the financial industry undergoes a significant shift, customers and companies are also moving toward an agentic workplace in which AI agents can orchestrate end-to-end business processes. This is achievable only when AI agents are integrated across the value chain and are driven by a strong orchestration engine. While multiple technical approaches exist, the following framework, The Agentic Banking Ecosystem, is deliberately business-aligned. It is designed to deliver measurable ROI and minimize risk by enabling banks to begin with low-risk back-office automation and progressively scale toward more customer-facing applications. The approach consists:

- **Front-office Agents:** AI agents integrated with channel applications focused on customer interactions, personalization, and sales
- **Mid-office Agents:** AI agents managing complex tasks like underwriting and risk management
- **Back-office Agents:** AI agents that drive operational efficiency

Agentic Bank Ecosystem – Representative Agent View

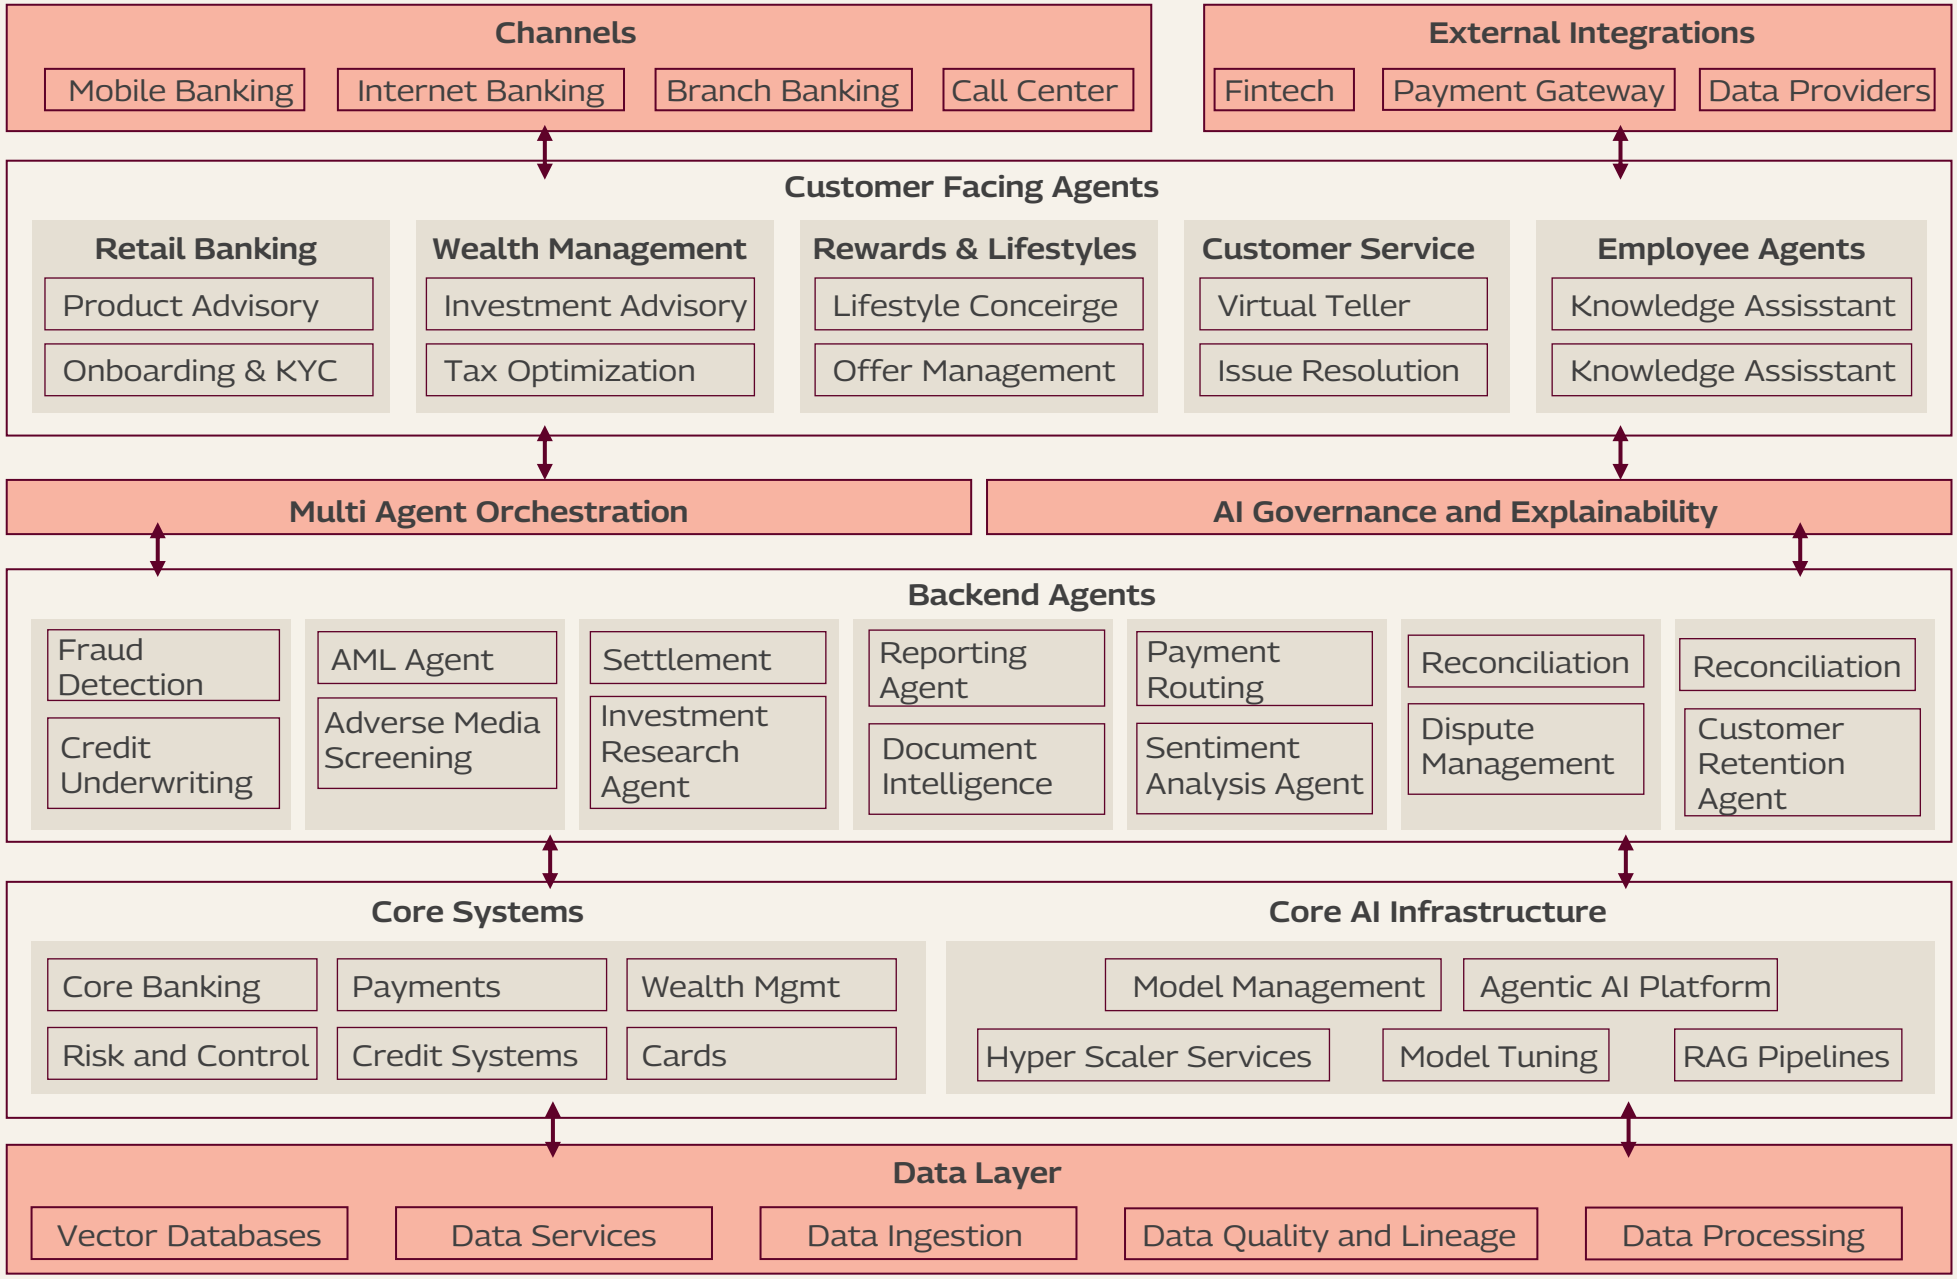

Figure 2: A Visual Representation of The Agentic Banking Ecosystem

Human-in-the-Loop Approach for Agentic AI Implementation

While agentic AI systems can autonomously monitor operations, analyze data, and execute decisions at high speed, they must never operate in complete isolation, especially in regulated financial services. The human-in-the-loop (HITL) approach provides the essential balance. It keeps AI agents within well-defined, trusted boundaries while allowing them to benefit from human expertise and judgment. When humans review, override, or approve critical decisions, the agents continuously learn and improve their accuracy over time

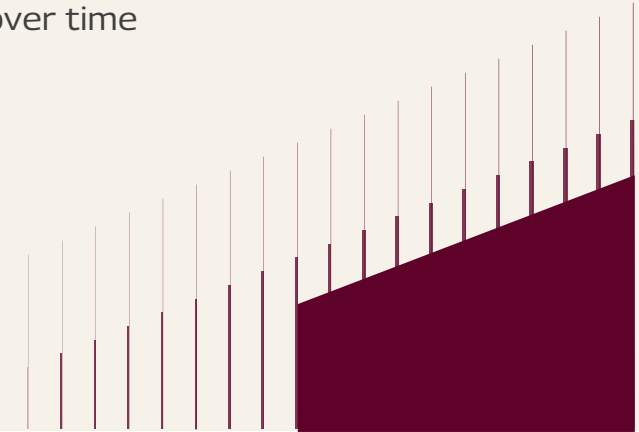

This collaborative model delivers the best of both worlds: the speed and scale of autonomous AI combined with the accountability, ethical reasoning, and contextual understanding that only humans can provide. It acts as a critical safeguard against unintended consequences such as bias, errors, or regulatory non-compliance.

For every implementation use case, banks should establish a clear distinction:

- What decisions can AI agents make independently
- When and how humans must be involved
- The non-negotiable boundaries that agents must never cross

This design not only builds trust with regulators but also accelerates adoption, as leaders see agentic AI as an augmentation of human expertise rather than a replacement.

Prerequisites for the Successful Implementation of Agentic AI Solutions

Implementing agentic AI at scale requires a strong foundational setup. Banks must address five critical prerequisites before rolling out agentic solutions.

Data Foundation and Infrastructure

Agentic AI agents need reliable, real-time access to data across the organization. Key requirements include:

- Access to a unified data layer
- Availability of APIs for agents to interact with
- Modernization of legacy systems to work with current integration technologies

Governance and Guardrails Framework

Strong governance is non-negotiable in banking. Organizations must:

- Define clear roles, responsibilities, escalation paths, and oversight mechanisms for AI agents
- Ensure every decision made by an agent is traceable and logged for audit
- Establish a human in the loop approach for escalation and edge cases

Talents and Technical Skills

Success depends on people as much as technology. Banks should build:

- Cross-functional teams with AI expertise
- In-house team with agent development knowledge and capabilities

Impact Measurement Framework

To justify investment and track progress, banks need:

- Tenets of the impact framework for measuring value
- A framework to define KPIs and measure benefits

Technology Stack Selection

Choosing the right platform and tools is critical. This process includes:

- Identification of the right technology to suit the current needs and environment
- Platform setup and identification

Addressing these prerequisites early significantly reduces implementation risk and accelerates the journey from pilot to enterprise-wide deployment.

Tech Mahindra's Proprietary Platform for Developing Use Cases

To accelerate the journey from concept to production, Tech Mahindra has developed Orion, a comprehensive, modular, and centralized AI platform specifically designed for the financial services industry. Orion integrates built-in AI governance through TechM VerifAI and enables banks to train and fine-tune LLMs, build RAG-based knowledge repositories, and design, orchestrate, and deploy autonomous, agent-driven workflows.

Developed as a joint initiative between Tech Mahindra and NVIDIA, Orion offers AI in a strategic approach that enables seamless adoption across financial services. Pivoted on Tech Mahindra's AI Delivered Right theme, it minimizes the challenges and fast-tracks the realization cycle for enterprises.

Agent AI for Bank: TechM unified Platform- POV

Channels

Agentic AI maturity assessment | Existing tooling assessment | Developer Tools | Change Management

Agentic AI|LLM Evaluation & Bench Marking | Feedback & Improvement(RLHF) | Experimentation Workbench

Enterprise-Ready Agents- Banking

Credit|Loans|Deopsits
- Credit Underwriting Agent
- Credit Risk Monitoring Agent
- Product Recommendation Agent

Trade Finance
- Doc Processing Agent
- Finance Optimizer Agent
- Trade Compliance & Sanctions

Payments
- Payment Routing Agent
- Payment Fraud Prevention Agent
- Cross Border payment Compliance

Wealth & Asset Mgmt
- Client Financial plan Agent
- Investment Portfolio creation & Optimization

Customer Onboarding
- Identity Verification Agent
- Account Set up Agent
- Orchestration Agent

Customer Service
- Identity Verification Agent
- Complaint resolution Agent
- Financial Advisory Agent
- Relationship Manager Agent

Back Office Ops
- Reconciliation & Settlement Agent
- Document Processing and Management Agent
- Opertaions Risk Control Agent

Corp Functions(HR, IT, Procurement)
- Talent Acquisition and Develoment Agent
- Vendor Management Agent
- IT issue Resolution Agent

Hyper Sclae Cloud
Microsoft Copilot Studio
Amazon Bedrock- GCP

Applications/Solutions
SAP, Sales Force solutions
Enterprise Tool Intergration

Automation & AI-ML Models
RPA-Automation |AI-ML Models |Process Mining Maps

Data Layer
Database|Data Lake |Data Platforms

Channels

Figure 3: TechM Unified Platform for Banks

Sample Use Cases

Use Case	Purpose	Agent-Based Agentic Workflow	Role of Agents	Role of Human-in-the-loop
Portfolio Management	Rebalancing and optimization of a client portfolio for investment bankers	Research Agent, Client Data Agent, Portfolio Analysis Agent, and Recommendation Agent	Provides insights for wealth advisors and maintains real-time portfolio optimization	Builds trust, understands personal goals, and offers emotional intelligence
KYC/Onboarding	Automate client onboarding	KYC Agent Document Extraction Agent	Conducts data extraction, identity validation, and AML compliance	Builds rapport, validates information, and ensures compliance
Customer Engagement	Help resolve customer queries	Mail Processing Agent and Response Preparation Agent	Manages routine queries, offers 24/7 support through agents	Handles nuanced queries and conflict resolution
Credit Assessment	Perform credit assessment during onboarding	Document Analysis Agent, Behavior Analysis Agent, Credit Assessment Agent	Analyzes a large volume of financial, behavioral, and credit data	Evaluates intent and business viability
Risk and Fraud Detection	Identify fraud risks	AML Agent, Fraud Agent	Monitors accounts and transactions 24/7, detects anomalies, and initiates proactive action	Applies situational judgment

Responsible AI Adoption in Banking

Responsible AI (RAI) practices are essential in the banking sector to ensure that AI agents' decisions are fair, explainable, transparent, and accountable. For instance, when an AI agent approves or declines a loan, the decision must be explainable and free from bias.

As banks increasingly adopt AI technologies to enhance efficiency and customer experience, implementing Responsible AI frameworks mitigates legal, ethical, and reputational risks.

The regulatory landscape further structures adoption. Central banks worldwide are already formulating frameworks; for instance, the EU AI Act and the RBI's FREE-AI framework mandate AI systems to be auditable and nondiscriminatory. Moreover, responsible AI practices ensure that AI models are tested, validated, and monitored, aligning with model risk management practices.

They ensure that automated systems operate in accordance with ethical standards, avoiding situations where AI produces unfair outcomes. This approach also involves having human-in-the-loop systems to validate and supervise AI decisions, especially in critical, high-impact decisions.

TechM VerifAI solution addresses these issues with a comprehensive approach to validating and verifying AI project outcomes. The solution brings together a strategic framework, industry-recognized tools, and experts with niche skills. It includes:

- Pre-built and tested 360-degree validation framework with customizable validation metrics
- Pre-built solutions and implementations for faster value realization
- Microservices-based architecture and API integrations to third-party solutions
- Templatized integration into delivery frameworks to ensure easy governance

Prioritization Framework

Given the hype around the agentic AI implementations, the use cases are plenty across the whole lifecycle of the banking functions. To ensure that the AI initiatives are aligned with the strategic goals of the organization and not solely driven by technology enthusiasm, organizations need to have a framework that can help them prioritize the AI initiatives

TechM's proprietary use case prioritization framework assists financial enterprises in methodically prioritizing use cases based on the multiple dimensions outlined below.

- Business Value
- Implementation Feasibility
- Data Quality and Availability
- Cost of Execution
- Ethical AI Considerations

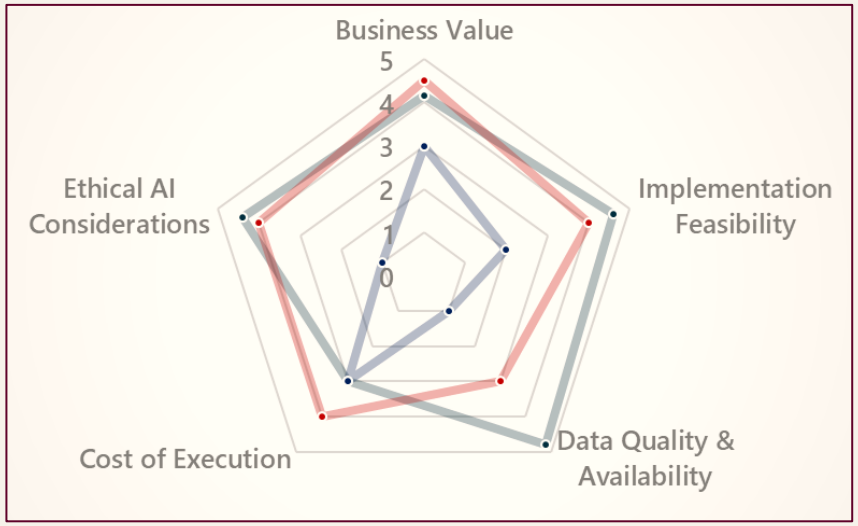

Figure 4: Sample Use Case Prioritization

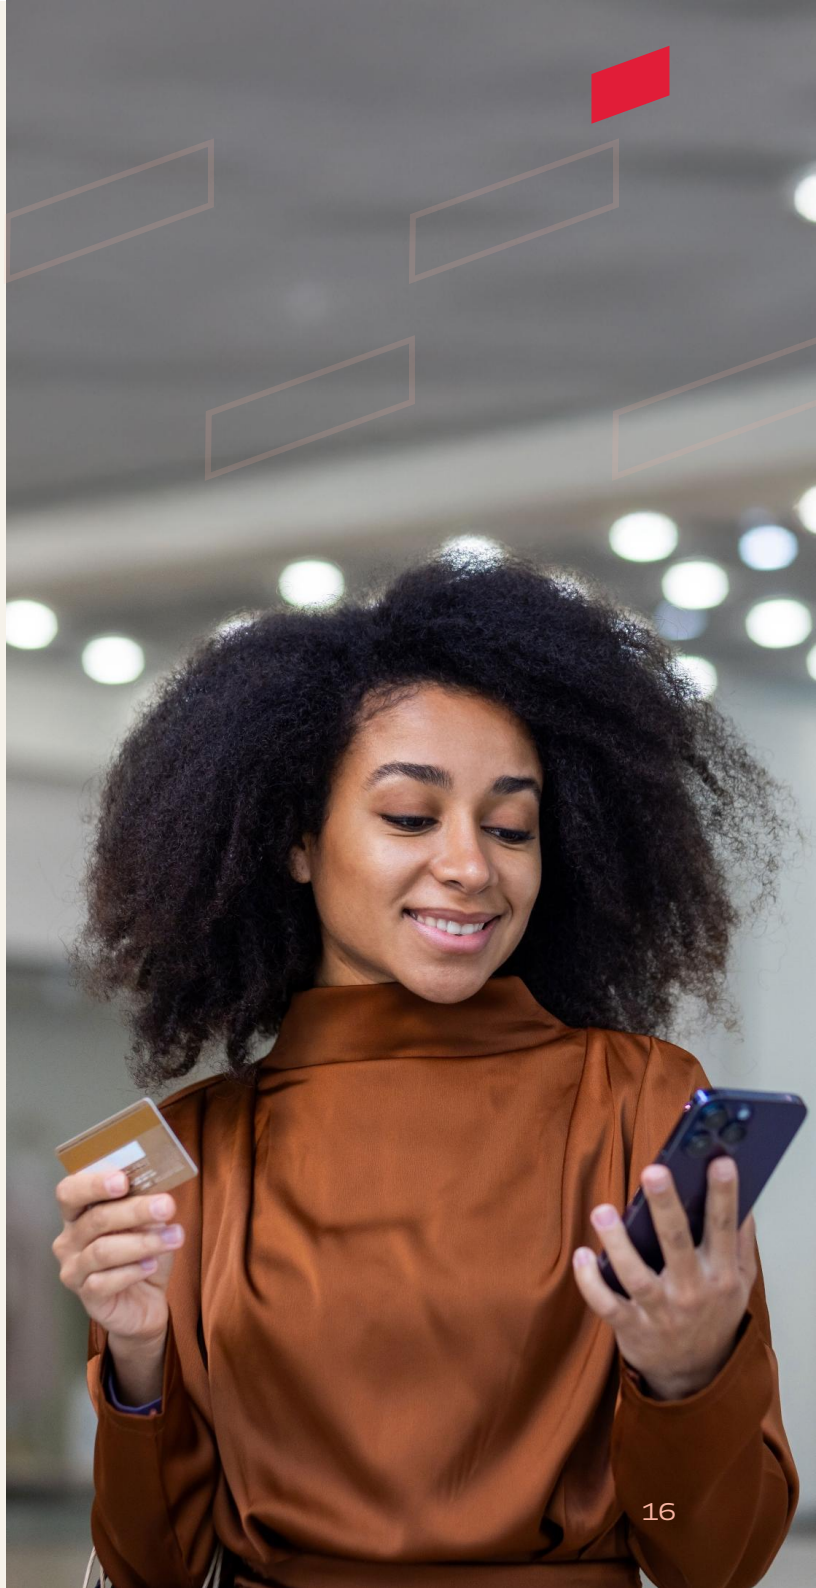

Strategic Roadmap for Agentic AI Adoption in Banking

To successfully implement agentic AI, financial institutions must follow a phased approach:

Phase 1: Assessment and Planning

- Identify high-impact use cases aligned with business goals.
- Conduct a technology audit to assess data readiness and infrastructure.

Phase 2: Pilot Implementation

- Deploy agentic AI in a controlled environment, such as fraud detection or customer onboarding.
- Partner with Tech Mahindra to leverage pre-built AI frameworks and accelerators.

Phase 3: Scaling and Optimization

- Expand AI deployment across functions and integrate with existing systems.
- Use real-time analytics to monitor performance and refine algorithms.

Phase 4: Continuous Improvement

- Establish governance frameworks to ensure ethical use of AI.
- Invest in employee training to foster AI-human collaboration.

Throughout this journey, Tech Mahindra's AI center of excellence provides end-to-end support, from strategy formulation to deployment, ensuring seamless adoption and measurable outcomes.

Road to Implementation

Phase 1:
Discovery and Assessment (0-3 Months)

- Analysis of existing systems, identify high-impact use cases (e.g., onboarding, reporting automation).
- Define KPIs: Cost reduction, customer satisfaction, compliance adherence.
- Deliverable: AI readiness report and prioritized use case roadmap.

Phase 2:
Pilot Deployment (3-6 Months)

- Deploy low-risk Agentic AI applications (e.g., compliance reporting, automated notifications).
- Integrate with existing systems using Tech Mahindra's ORION Framework.
- Implement AI guardrails
- Deliverable: Pilot results with 20-30% efficiency gains.

Phase 3:
Scale and Optimize (6-12 Months)

- Expand to complex workflows (e.g., onboarding, wealth management).
- Implement multi-agent systems using Collaboration Framework.
- Deliverable: 30-50% reduction in manual effort, 20-60% productivity gains.

Phase 4:
Continuous Improvement

- Enable continuous learning with real-time feedback loops.
- Optimize using VerifAI for governance and performance monitoring.
- Deliverable: Fully autonomous, compliant AI ecosystem.

Figure 5: A roadmap to AI Adoption Conclusion

Conclusion

Agentic architecture provides a practical framework for developing autonomous, scalable, and compliant banking systems. By integrating governance, traceability, and continuous learning at every level, banks can achieve new heights of operational agility while maintaining trust

Adoption isn't about replacing humans - it's about enabling institutions to act with speed, precision, and resilience at scale. With proper safeguards, agentic systems become a strategic asset, not a risk.

End Notes

- MarketsandMarkets. (2025, April). AI agents market by agent role, offering, agent system - global forecast to 2030 (Report Code TC 9168). https://www.marketsandmarkets.com/Market-Reports/ai-agents-market-15761548.html
- Mordor Intelligence. (2026). Agentic AI in financial services market size, share & 2031 growth trends report. https://www.mordorintelligence.com/industry-reports/agentic-artificial-intelligence-in-financial-services-market
- KPMG. (2025). AI quarterly pulse survey. https://kpmg.com/us/en/articles/2025/ai-quarterly-pulse-survey.html
- Tech Mahindra. (2025, October 27). Tech Mahindra releases survey report on "Building the AI-Driven Bank of Tomorrow" [Press release]. https://www.techmahindra.com/insights/press-releases/tech-mahindra-releases-survey-report-on-building-the-ai-driven-bank-of-tomorrow/

About the Author

Gopal Parasnis heads digital transformation for the BFSI vertical at Tech Mahindra. Gopal has around 25 years of experience in the IT industry and has been part of several large-scale digital transformation initiatives across the banking and financial services clients. He has led and advised on several mainframe modernization journeys. He has a deep understanding of the latest industry and technology trends in the space sector.

Gopal Parasnis

Head, AI and Digital Transformation, Financial Services

Tech Mahindra

Saurabh is a part of the Digital Transformation Office for BFSI at Tech Mahindra. He has around 20 years of experience in the IT industry and has executed large-scale API transformation projects for global banking clients. He has a deep understanding of the API landscape and technology trends.

Saurabh Agrawal

Practice Head-Banking & Financial Services

Tech Mahindra

About **Tech Mahindra**

Tech Mahindra (NSE: TECHM) offers technology consulting and digital solutions to global enterprises across industries, enabling transformative scale at unparalleled speed. With 147,000+ professionals across 90+ countries helping 1100+ clients, Tech Mahindra provides a full spectrum of services including consulting, information technology, enterprise applications, business process services, engineering services, network services, customer experience & design, AI & analytics, and cloud & infrastructure services. It is the first Indian company in the world to have been awarded the Sustainable Markets Initiative's Terra Carta Seal, which recognizes global companies that are actively leading the charge to create a climate and nature-positive future. Tech Mahindra is part of the Mahindra Group, founded in 1945, one of the largest and most admired multinational federation of companies. For more information on how TechM can partner with you to meet your Scale at Speed™ imperatives, please visit https://www.techmahindra.com/.

*Figures as per Q4, FY 26.

www.techmahindra.com
www.linkedin.com/company/tech-mahindra
www.x.com/Tech_Mahindra