

EBOOK

Agentic AI for Banking

Making AI real: From pilots to production
on Microsoft Azure

Introduction

Efficiency savings. Productivity gains. Hyper-personalization that directly translates to stronger customer loyalty and reduced churn. And [2.3x ROI \(on average\) in 13 months](#). That's the promise of AI-powered autonomous agents.

Many financial services organizations have already implemented early AI technologies. To date, though, these have primarily been AI assistants and automation tools used to support specific tasks rather than act independently (e.g., customer service chatbots and generative AI for communications support in the front office). Agentic AI represents the next step in the evolution of autonomous AI technology and a significant leap in its capabilities.

The Potential of AI for Financial Services

Agentic AI has the potential to deliver:

- Efficiency savings, innovation and [\\$3 trillion in corporate productivity gains](#)
- A [5.4% EBITDA improvement](#) for the average company
- An average [2.3x return on agentic AI investments](#) within 13 months
- Hyper-personalization that [directly translates to stronger customer loyalty](#) and reduced churn



The real challenge isn't building Agentic AI; it's proving its value."

Dmitry Tikhomirov

VP, Head of EPAM Global Microsoft Business Group



The Challenge With Many AI Projects

Around [half of all financial services organizations are running agentic AI pilots](#), yet just a fraction of those have existing deployments and many projects will fail.

These projects don't fail because of an inherent issue with the technology. In our experience, they fail because of the way the technology is applied.

The real challenge isn't building Agentic AI; it's proving its value. Organizations are rapidly discovering that simply inserting agents into legacy workflows yields only marginal gains. Instead, real ROI requires AI-first process redesign built around clearly defined business outcomes.

As banking trade publication The Financial Brand reported, ["Only 32% of banks realize significant returns from customer-facing AI investments despite 99% prioritizing these initiatives."](#)

“ Only 32% of banks realize significant returns from customer-facing AI investments despite 99% prioritizing these initiatives.”

The Financial Brand

Agentic AI Made Real, by EPAM

EPAM has been helping financial services businesses get beyond the pilot stage of Agentic AI with Microsoft Azure.

In this eBook, we'll explore what agentic AI is and why it matters for your organization. We'll examine how, through Microsoft tools such as Azure AI Foundry, Copilot Studio, and Agent Framework, you can enable Agentic AI at scale, driving ROI across retail, capital markets, commercial banking and more. And you'll see Agentic in action, as it helps financial organizations meet the requirements of Know Your Customer (KYC) regulations.



What Is Agentic AI?

Agentic AI represents an important step forward in AI technology, enabling advanced reasoning and semi-autonomous execution, to support, not replace, expert human judgement. In mortgage operations, for example, an AI agent might coordinate document collection, affordability checks and application tracking, enabling human underwriters to focus on exceptions or complex cases requiring their judgement.

What are 'Agents'?

In Agentic AI, agents are the equivalent of the people in your teams. Just like those people, you won't expect each individual to be an expert at everything – each will have certain specialisms, but together they will provide a full span of expertise across disciplines ranging from regulatory compliance and trading to KYC and risk.

Like any growing team, numbers will start small but as the agents prove their efficiency, they have the potential to scale, growing to hundreds or even thousands. And each agent will enhance each other's capabilities in a carefully managed, collaborative and self-evolving network.

Who owns the Agents?

As the organization commissioning Agentic AI, you do.

Enhancing Human Capabilities

While far more autonomous than previous AI technologies, AI agents work in combination with humans. For example, an AI agent could monitor transactions for signs of fraud, prioritize alerts and recommend actions, while compliance and fraud specialists would investigate flagged cases and make final decisions.

In fact, AI systems that require humans to steer and orchestrate agents [outperform those that don't](#).



Real ROI requires AI-first process redesign.”

Dmitry Tikhomirov

VP, Head of EPAM Global Microsoft
Business Group



How To Maximize ROI From AI

Four steps to maximizing ROI from Agentic AI:

1. Take an AI-first approach to process redesign

Rather than attempting to insert agents into legacy workflows, which will yield marginal gains and limit ROI, start with an AI-first redesign of processes.

2. Embed cost governance from day one

Agents need to demonstrate their value. Ensure business KPIs are evaluated for and aligned with every individual agent, embedding cost governance from day one.

3. Bake in security, accountability and transparency

Embed safety and resilience, traceability, bias and drift monitoring, auditable logs, and secure-by-design controls to reduce technical debt and cyber risk.

4. Build for the future

Build with tomorrow's technology in mind and design processes and systems to accommodate emerging AI capabilities, rather than retrofitting or making do with old and outdated systems. Modular, composable architecture will play a significant role in this, enabling businesses to add, swap, and combine agents, models and tools without significant disruption or expenditure.



The path to becoming an AI-frontier organization isn't only about adding technology. It is about redesigning processes, empowering people and building a modular, governed ecosystem that can evolve safely."

Dmitry Tikhomirov

VP, Head of EPAM Global Microsoft Business Group



Agentic AI In Action: Across The Banking Industry

Meet the Agents

Which agents you employ and how you employ them will vary depending on the business division in which they operate. Here's a comparison of agents in retail, capital and commercial applications.


Agents in Capital Markets

Execution, client intelligence and risk and regulatory compliance are the core aspects of agentic AI for capital markets. From data streamlining and client behavior profiling to identifying trading opportunities and monitoring and managing real-time credit, market and liquidity risk, agents are making their impact felt across the capital markets domain.




Agents in Retail Banking

Retail banking is already highly digitized, but execution remains fragmented across products and channels. Agentic AI addresses this by introducing domain-aligned agents across customer engagement, payments, lending, risk and operations. These enable real-time decisioning and coordinated execution across millions of interactions, rather than isolated, rules-driven processes.



Agents in Commercial Banking

Commercial banking relies on coordinated decisions across lending, payments, trade and risk, often fragmented across products and teams. Agentic AI addresses this by embedding domain-aligned agents into the client lifecycle, enabling consistent, structured decision-making around capital, cash and risk, and improving how banks manage relationships rather than individual transactions.



Agentic AI In Action: KYC Workflow

So how does Agentic work in a real scenario?

[Know Your Customer \(KYC\)](#) processes are one of the most complex and resource-intensive areas in financial services, requiring manual data checks, regulatory oversight and significant operational effort.

EPAM's multi-agent Know Your Customer (KYC) workflow shows how agents can work together to automate decision-making, reduce operational cost and improve both speed and transparency.

How we give Agents Autonomy

In EPAM's KYC workflow, multiple agents work together to extract data from multiple sources and make decisions based on their analyses. The agents' work begins with a human prompt to process a KYC verification. Then, the agents take over ...

AI Agents in a KYC Workflow

Orchestrator agent

The team lead responsible for kicking off the process.

Document analysis agents

Responsible for investigating customer IDs and documents, and biometric verification.

Screening agent

This global investigator carries out internal verification and checks the watchlist and PEP databases. Operates in parallel with the document analysis agent.

Risk assessment agent

Conducts advanced analysis to calculate a risk score (low, medium or high) for a particular client. A "low" or "high" score advances the case to the final decision stage. Human-in-the-loop oversight is vital for cases deemed "medium risk" and ensures there's a transparent, traceable log of the decision-making process.

Case manager agent

The final decision maker responsible for creating a comprehensive overview of a customer and approving or rejecting the case. The agent automatically approves "low" risk scores, rejects "high" risk scores and relies on human intervention for "medium" risk. The organization sets the risk appetite.



The Impact Of Agentic

When we've applied Agentic AI to financial services projects, the results have been compelling:

- We developed an agentic co-pilot to simplify tax monitoring for tax researchers, achieving a 93% reduction in unnecessary alerts and data inconsistencies.
- We built a digital platform with an AI virtual assistant that automates tasks, enables rapid information retrieval and delivers personalized interactions to improve employee satisfaction and engagement.
- When we developed an AI-powered underwriting assistant for a leading insurer, data extraction time fell by 50%.
- When we updated 17,000 lines of legacy code for an FS client, a project that would have taken a full squad three months was completed by two people in two weeks.
- Our data suggests FS transformation projects that harness agentic capabilities are, on average, 20% faster and cheaper than those that don't.



How EPAM Applies Microsoft's Agentic Ecosystem

Microsoft's agentic AI ecosystem enables adoption at scale. It is particularly well-suited to financial services organizations, combining agentic capabilities with world-leading data integration, enterprise-grade security, regulatory compliance capabilities and powerful automation tools.

Microsoft's comprehensive suite of tools include:

- **Azure AI Foundry**
- **Copilot Studio**
- **Agent Framework**

As Microsoft's Partner of the Year in 2025, no one is better placed to bring agentic AI to your organization than EPAM.



Familiarity and ease of use

One of the main advantages of Microsoft's Agentic AI platform is its integration of existing Microsoft features and processes into an intuitive, easy-to-use platform. This ensures most developers will be familiar with much of its functionality.

For instance, Microsoft's open-source agentic AI development kit, Microsoft Agent Framework, builds on Semantic Kernel and enables developers to create AI-powered features using established languages, such as C#, Python and Java. It is also designed to fit into existing applications rather than replace them, enabling developers to integrate it with web APIs, microservices and cloud apps. The result is a development kit that feels familiar and enables developers to leverage existing architectures, languages and engineering practices.

Seamless integration with existing Microsoft technologies

Copilot Studio enables financial services organizations to automate complex, highly regulated workflows – such as claims handling, fraud detection and compliance – while maintaining strict data privacy and regulatory standards.

Finally, Microsoft Azure integrates seamlessly with existing Microsoft technologies, enabling financial services to easily transition to cloud services and agentic AI technologies without extensive rebuilds.

The platform is designed for use with .NET and Visual Studio development tools, Windows Server and Active Directory, Microsoft 365 and data platforms such as Microsoft SQL Server. This simplifies the development and deployment of agentic AI, lowers costs and eliminates several potential barriers to technology adoption such as complex system integration, application migration challenges, data compatibility issues and staff retraining requirements.

Governance, Trust & Control

In a highly-regulated industry like financial services, governance is not optional. Rather, it's the foundation on which any successful AI deployment is built. EPAM embeds governance from the outset of every project, ensuring agentic AI is implemented from the beginning in an ethical, accurate and transparent manner.



Bias and fairness

With agentic AI, explainability and governance must be built into any AI project from the very start to protect fairness and eliminate bias. Throughout the project timeline, engineers must also provide oversight and constantly monitor for bias, supported by technology.

Microsoft's advanced tooling allows for continuous bias monitoring. Azure's AI Content Safety, Responsible AI Framework and supporting tools facilitate consistent, comprehensive monitoring from prototype to production. Additional tools, such as the PyRIT Risk Identification solution and the open-source Fairlearn responsible AI dashboard, can complement the bias monitoring process.



Transparency and traceability

Agentic AI performs complex reasoning using structured and unstructured data from multiple sources. It's constantly iterating and refining its process and outputs. It's vital, therefore, that the decisions it makes are transparent, traceable and explainable, so underwriters and other insurance professionals can quickly and effectively establish data lineage.

Azure Agent Service makes this easy. It constantly monitors all agentic AI processes, listing the data sources and tools an agent uses and the reasoning it performs, to provide comprehensive oversight.



Accountability

When an AI agent raises an issue, who's responsible for taking it further? When is human intervention or authorization required? Where does accountability sit?

The answer lies in the Decision Authority Framework, a comprehensive tool suite that clearly defines roles and responsibilities in organizational decision-making, removes ambiguity and ensures accountability is enshrined in a considered and codified manner.



Security

AI agents interact with a wide range of data sources and are integrated into complex systems. As with any digital system, effective security is fundamental.

- Azure Entra is an identity and access management (IAM) solution for user identities, apps and devices.
- Microsoft Entra Agent ID extends Entra's comprehensive security capabilities to AI agents, enabling teams to safely build, govern and protect agent identities.
- Purview enables financial organizations to discover, classify and govern sensitive data, as well as automate data lineage and manage compliance risk.
- Built-in AI tooling helps secure the AI system, such as Prompt Shields, Groundedness Detection, safety system messages, safety evaluations and risk and safety monitoring.

Microsoft's solutions are already certified to standards and regulations, including ISO 27001, SOC 1/2/3 and EU GDPR, helping financial organizations meet the industry's strict data privacy, security and regulatory requirements.

Where To Start For Banking?

Start specific. Build incrementally. Scale with confidence.

01

Build leadership buy-in

In the financial services industry, executives set risk appetite, so the first step in any agentic AI project is to listen to executives' concerns and build understanding of benefits and challenges.

02

Get your governance ready

Strong governance is a prerequisite for any agentic AI deployment. Embedding AI safety and oversight before agents are deployed, including defining what agents will do, where they operate and how they are monitored, establishes clear boundaries and stronger protections.

Start with enterprise risk and governance processes, supported by Azure tools such as Microsoft Fabric for data lineage, AI Content Safety for filtering inputs and outputs, Agent Service for orchestration and monitoring and PyRIT for testing security posture.

03

Start specific

Many organizations make the mistake of rolling out a generic AI tool and expecting it to perform a specific function. AI agents can't do it all. If your business is highly diverse, you will need a diverse range of specialized agents operating as part of an orchestrated system.

Initially, apply AI agents to narrow, well-defined use cases. This simplifies the journey, improving the likelihood of successfully bridging the gap between prototype and production.

04

Scale incrementally

The journey from prototype to production is a step-by-step process. Progress is incremental. The most successful projects start by developing a solution that works for one individual. Then they gradually expand.

This step-by-step approach breaks the problems associated with scaling into smaller, more manageable challenges. It also ensures you scale on strong foundations. Once the initial agent is in production and has achieved its objectives, move on to the next use case and the next team, expanding your multi-agent system one agent at a time.

For more information on getting started on your AI journey, take a look at our 36-month AI differentiation roadmap:

[EPAM's 2026 Banking Playbook | EPAM](#)

05

Know your business

Every company has its own pain points, technologies and data landscape.

Select use cases that reflect these realities, focusing on a specific business challenge with clear outcomes and high-value data. Start with a low-risk use case and take it all the way to production.

This not only demonstrates the value of agentic AI but also ensures compliance processes are in place for safe scaling.

EPAM, Agentic AI & You

EPAM helps financial organizations design, embed, and orchestrate agentic AI solutions built on Microsoft Azure with strong governance, transparency and scalability.

As a Microsoft Global Systems Integrator and Innovate with Azure AI Partner of the Year, we enable organizations to deploy agentic AI at enterprise scale, focus on the use cases that deliver the greatest business value and drive sustained AI adoption across the organization.



Discover how EPAM and Microsoft are helping financial companies to move from AI experimentation to production.

Find more information [here](#) or [contact us](#)



Global

41 University Drive, Suite 202
Newtown, PA 18940, USA

P: +1-267-759-9000

F: +1-267-759-8989

