

# How Retail Banks Can Put AI Agents to Work

By [Stiene Riemer](#), [Marc Pauly](#), [Bharat Poddar](#), and [Dan Sack](#)

**ARTICLE** MARCH 09, 2026 15 MIN READ

*This report is a joint initiative of BCG and OpenAI.*

Over the past several decades, retail banks have spent billions of dollars digitizing many aspects of their operations, from customer-facing front-office activities to the back-office systems—including credit assessment, onboarding, servicing, and dispute resolution—that support the customer journey.

Yet translating the output of these front- and back-office operations into action still depends heavily on the manual processes involved in reconciling information across systems, summarizing findings, and routing cases for downstream decision making. These labor-intensive activities represent a persistent source of cost, delay, and operational strain without delivering commensurate improvements in risk or compliance outcomes.

By shifting these tasks from human operators to supervised and auditable agentic systems, banks can significantly reduce customer onboarding costs and refocus human capacity on exception handling, judgment, and customer engagement. The same pattern extends to the back office, where automating routine checks and data ingestion or document creation allows banks to improve speed, consistency, and auditability, freeing up human talent to focus on higher-value advisory and oversight roles.

According to a recent BCG study, AI agents have the potential to increase banks' profitability by 30% and reduce costs by 30% to 40% by 2030. Early adopters of agentic AI will not just capture near-term efficiency gains; by moving first, they can increase productivity, improve economics, reduce cycle times, and deliver superior customer experiences—advantages that will become increasingly hard to replicate and compete against.

This report examines how banks can incorporate agentic AI throughout their operations, improving the customer onboarding experience at the front end while transforming the back-office systems and processes to reduce onboarding costs, shorten approval times, and reallocate staff capacity toward customer engagement and judgment-intensive tasks.

# The Promise of AI Agents

A recent BCG study shows that many retail banks have already begun incorporating AI into their operations. They are using generative AI bots to help customers manage their online banking activities and automating some aspects of key workflows such as fraud risk identification, regulatory compliance and risk management, and digital marketing.

Implementation of the new technology, however, has been slowed by several factors: Top management may lack commitment to the AI transformation. Efforts are often fragmented—there are many pilots but few efforts to scale them. Concerns about the reliability and accountability of AI's outputs remain. And it is in every bank's DNA to proceed cautiously when it comes to critical regulatory and compliance issues that could be impacted by AI.

Incorporating reliable and auditable AI agents into key activities can enable banks to overcome their concerns about AI's scalability and compliance issues. Agentic AI has the capacity to supervise and execute entire end-to-end processes across service, compliance, risk, and exceptions. Human managers can exert full control over the actions taken by properly designed AI agents, ensuring both reliability and accountability. At the front end, agentic AI could lead to a shift from digital self-service to intelligent assisted service. In the back office, AI agents could assist with manual review of applications, shortening approval timelines, and improving throughput without altering existing risk, compliance, and credit policy frameworks.

## Front-Office Transformation and the Customer Experience

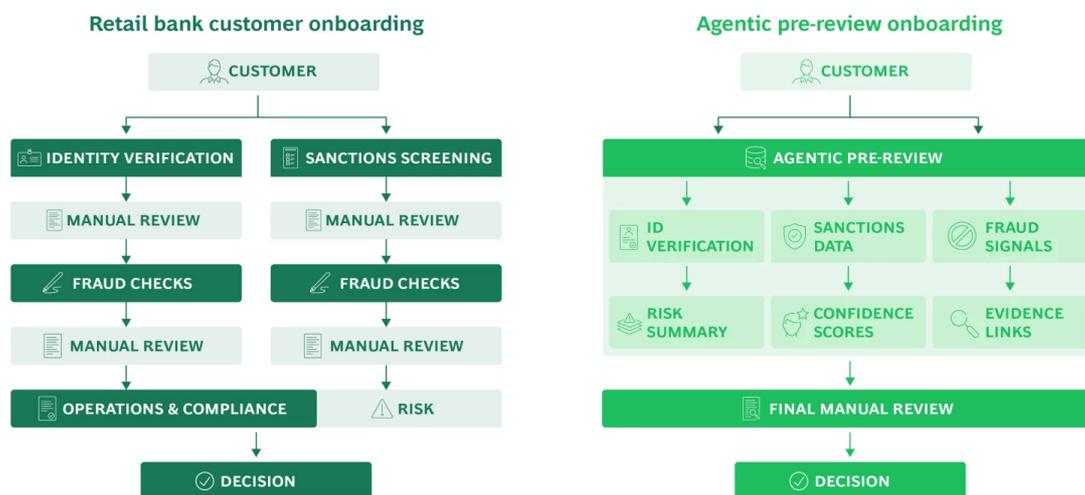
Over the past few decades, banks have successfully digitized customer experiences, automated isolated processes, and built a range of online, mobile, and branch-based digital channels. These changes have improved access and convenience for customers but have not materially transformed the institutions themselves. In the mobile era, for example, customers gained transparency and self-service but underlying operating models remained largely unchanged. Higher-value workflows stayed human driven, customer journeys remained fragmented, and automation was limited to rules-based tasks.

Generative AI, particularly agentic systems that can reason and take action, now allows banks to provide tailored experiences at scale while guiding customers through complex tasks. Rather than navigating menus and forms, customers can interact with personalized financial assistants that understand conversational nuances. A customer might ask for help in opening a credit line, buying a home, managing cash flow, or resolving a dispute. Agentic AI systems can retrieve information

across accounts, check eligibility, analyze documents, initiate workflows, and coordinate with human specialists when needed, managing the entire life cycle of the request. (See Exhibit 1.)

**EXHIBIT 1**

**Agentic AI Systems Can Substantially Reduce the Manual Effort Involved in Customer Credit Onboarding**



Sources: OpenAI; BCG analysis.

Consider one of the most common and operationally expensive front-office functions at large retail banks: consumer credit onboarding. Despite years of investment in digital channels, the process remains both costly and slow. Its inherent inefficiencies degrade the customer experience, limit consumer access to credit, increase operating costs, and place undue strain on compliance and operations teams without a corresponding improvement in fraud prevention or credit risk outcomes.

The root cause is not a lack of automation. It is the ongoing dependence on humans to manually reconcile and summarize outputs from systems that are already trusted and validated. So even after identity, fraud, and credit checks are complete, banks often rely on human teams to stitch together the data before making a decision.

AI agents can overcome these challenges and speed up decision making by performing the initial analysis of the customer’s onboarding profile. The agent evaluates identity verification outputs, sanctions screening results, fraud signals, and credit bureau data to produce a structured risk summary and confidence assessment. Rather than replacing existing controls, the AI agent operates within the bank’s established risk and compliance framework and produces outputs that are explainable, traceable, reviewable, and auditable, creating a pre-review summary for downstream underwriting and human review. Traditional underwriting models continue to assess credit risk, affordability, and policy compliance.

After human review of these evaluations, the bank issues an approval, declines to provide credit, or restructures the offer. If the customer accepts the original or amended conditions, the bank proceeds with account creation, disclosure acceptance, and card issuance. Post-activation monitoring and early life risk controls continue unchanged. The result: a more transparent and accelerated credit review process that significantly improves the customer experience and frees up staff to focus on discrepancies, exceptions, and higher-risk cases.

## The Agentic Back Office

The agentic onboarding system described above provides the connective tissue between customer intent, risk evaluation, compliance controls, and fulfillment. To support the agentic process in generating the credit onboarding summary at scale, front- and back-office systems must operate as a single connected workflow.

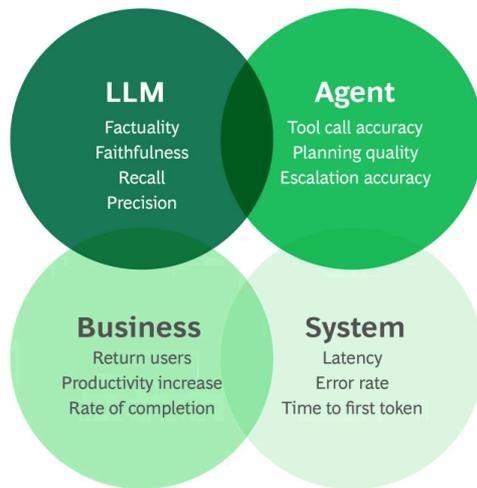
While AI can streamline the customer onboarding experience, the back office is where it will deliver the most immediate and substantial impact. The onboarding process is underpinned by a diverse set of structured, repetitive actions that require interpretation, judgment, and manual effort. Historically, most of these processes could not be automated because they required understanding documents, interpreting context, or handling exceptions. AI systems can now perform many of these tasks under supervision. They can interpret complex documents, extract information, analyze cases, escalate when appropriate, and maintain full audit trails.

To successfully run complex agentic systems in the back office, banks need to implement two critical components for ensuring the reliability and scalability of the system. First, they need a repeatable way to observe and evaluate the AI agent's performance on the same kinds of document-heavy, exception-driven tasks their human teams do every day and be able to track the quality of these tasks over time as workflows and models change. Second, they need to build the in-house middleware layer that will establish a single, standardized control plane for all AI applications across the organization.

**Is it accurate?** To ensure the quality, correctness, and auditability of agentic AI outputs, they must be evaluated against the real-world back-office tasks they are designed to replace. This requires evaluation-driven development (EDD), an approach that measures the performance of an AI application across a range of defined dimensions and helps steer development priorities. EDD also provides concrete evidence to internal and external control partners that the system is being built, tested, and operated in a systematic and auditable manner. (See Exhibit 2.)

## EXHIBIT 2

### Evaluation-Driven Development Is Essential in Assessing the Reliability and Accuracy of a Range of Workflow Factors



Sources: OpenAI; BCG analysis.  
Note: LLM = large language model.

In the early stages of development, an AI agent must be evaluated to make sure it can reliably access the information required to perform its task. This initial “test” should encode clear examples of acceptable and unacceptable outputs and serve as the basis for teams to iterate on retrieval strategy, chunking, ranking, and prompt design.

For example, in a “know-your-customer” workflow, the application should retrieve the correct supporting materials at inference time; this would be measured through retrieval precision and recall. In parallel, evaluations should verify that model outputs are grounded in the retrieved evidence and do not introduce unsupported claims, measured through factuality and faithfulness.

As applications are expanded into tool-using and multistep workflows, how the AI agent performs must also be evaluated. In particular, its ability to construct a complete and policy-aligned plan becomes critical to overall system correctness. Without sound planning, even perfectly executed downstream results such as accurate “tool calls”—interactions with external systems like databases and APIs—or rapid response times can result in outcomes that are materially misaligned with the intended objective.

In an underwriting workflow, for instance, an agent may execute all tool calls correctly yet prematurely assume that a prospect’s credit risk is acceptable, focusing on income verification while failing to explicitly assess credit history. This represents a planning failure and introduces real business and regulatory risks—exactly the types of issues that EDD is designed to surface early, before deployment at scale.

Evaluation does not stop at individual model or agent behavior. The system must be monitored holistically as a production application, with continuous measurement of output quality, routing accuracy, exception rates, latency, and drift. This ongoing system-level evaluation ensures that the

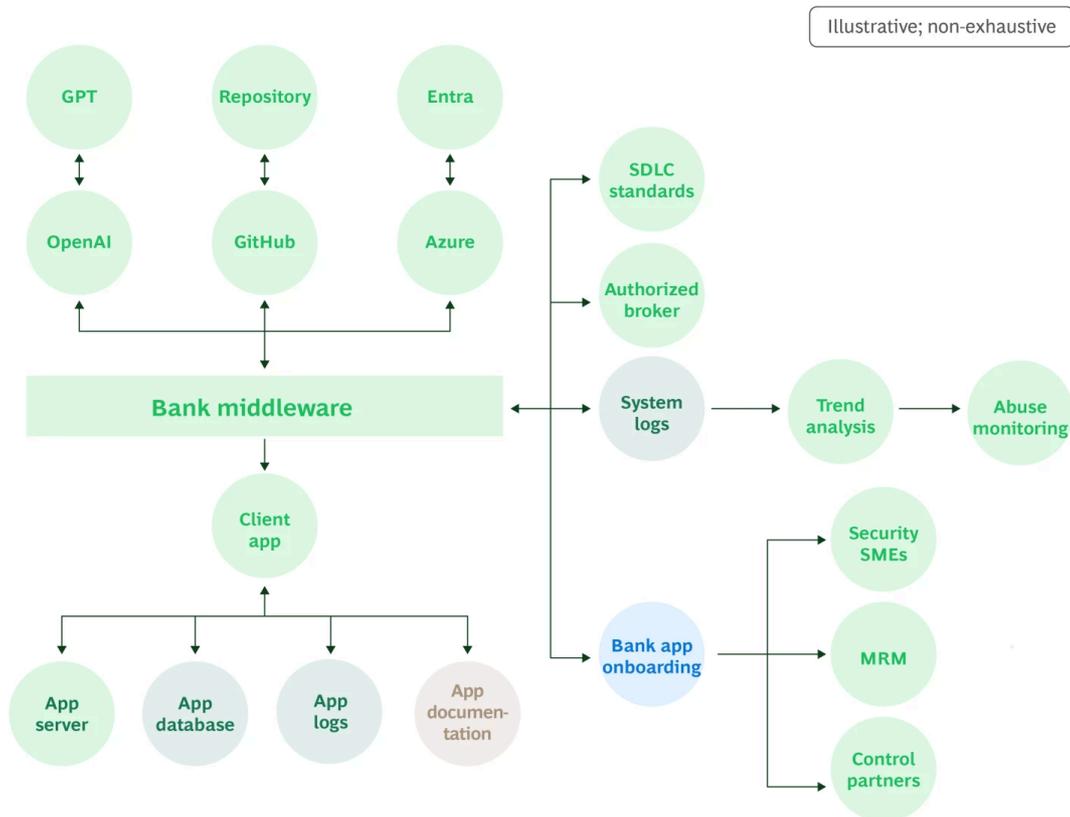
application continues to meet defined business, risk, and control objectives as data, usage patterns, and underlying models evolve.

(Note that evaluation can also be carried out through evaluation frameworks such as OpenAI’s GDPval, a publicly available platform that measures the performance of any AI application on standardized, well-specified knowledge work tasks across occupations, including financial services, and scores its output using expert human judges. The tasks are documented publicly on Hugging Face, an independent platform that allows teams to run the same prompts, inspect intermediate outputs, and extend the evaluation framework to their own workflows and data sources.)

**Does it scale?** A strong and stable in-house middleware layer is critical for large, highly regulated organizations such as banks. It enables them to uniformly enforce software development standards, authentication and authorization policies, and onboarding requirements for internal control partners such as model risk management (MRM) and compliance. Rather than allowing each team to implement bespoke integrations and control mechanisms, the middleware acts as a consistent “front door” through which all AI workloads must pass. (See Exhibit 3.) This enables banks to avoid control drift and can demonstrate that every AI application is built, deployed, and operated under the same governance framework.

**EXHIBIT 3**

The Bank’s Middleware Gathers All the Necessary Inputs for the Onboarding Workflow in One Place



Sources: OpenAI; BCG analysis.  
 Note: MRM = model risk management; SDLC = software development life cycle; SME = subject matter expert.

Just as important, the middleware layer centralizes and simplifies the need to observe and audit AI agents in action, more strictly enforcing their conformity to policy. It provides a single place to capture standardized logs, usage metadata, and decision traces across all AI clients, enabling effective monitoring, investigation, and regulatory reporting. Because identity, permissions, data access rules, and safety guardrails are enforced consistently before and during each model interaction, banks can manage sensitive data and detect in real time any anomalous or abusive behavior on the part of the AI agent. This level of visibility allows banks to respond quickly to problematic incidents, apply targeted controls or kill switches, and limit the blast radius of potential failures.

The middleware layer makes AI adoption scalable and sustainable within a highly regulated environment. By brokering connectivity to outside providers of AI models and abstracting individual applications away from vendor-specific implementations, banks retain strategic flexibility while maintaining uniform controls. New AI applications can be onboarded more efficiently using reusable patterns and shared infrastructure rather than rebuilding the process from scratch each time. The result is a platform approach that allows banks to expand AI usage across business lines while maintaining consistent risk management, operational discipline, and regulatory confidence.

**What's the result?** A solid, stable back office, with consistent evaluation processes and strong middleware, is the foundation of the AI-enabled bank. Properly implemented and run, it will allow banks to reduce operational cost and complexity, increase speed, strengthen compliance, decrease error rates, and reallocate human talent toward higher-value advisory and relationship roles. The back office becomes a digital engine that supports efficiency and superior customer outcomes and serves as the critical enabler of front-office transformation. Without a modernized back office, intelligent customer experiences cannot scale. Together, the front and back offices form a single AI-native system. This is the foundation of the bank of the future.

## Internal Adoption and Workforce Capability

The adoption of AI agents as a crucial and consistently valuable part of retail banks' operations will require a long-term transformation involving a major re-architecting of core customer journeys and operational processes. To achieve this, banks need to create cross-functional teams that bring together AI engineers, architects, platform teams, safety and compliance experts, and domain specialists who can design AI-native processes and ensure they operate safely at scale.

Often referred to as an AI center of excellence (CoE), this team is responsible for establishing and disseminating consistent best practices, tools, frameworks, and technologies across the organization. While business teams continue to own the end-to-end delivery of their individual use

cases, the CoE acts as a group of trusted experts collaborating closely with the business and shepherding the successful execution of high-priority initiatives.

For organizations early in their AI adoption journey, the CoE can help map out which use cases should be prioritized through a standardized intake process. Business units submit proposals that include the objective, baseline metrics, required systems and tools, data sensitivity, and an initial assessment of control requirements. The CoE then scores candidates using a transparent rubric across technical feasibility, business value, and strategic fit; delivery readiness with respect to data, process ownership, and platform capacity; and control requirements, including implications for information security and privacy, compliance, and MRM. The CoE and business unit leaders then select a small initial portfolio to deploy, with clear owners, timelines, and control-partner checkpoints.

The CoE should maintain an approved AI tool chain and a repeatable process for evaluating models, frameworks, and software development kits as they evolve. The team should be fluent in modern evaluation and guardrail methodologies and actively identify opportunities to reuse capabilities built elsewhere in the organization. Just as importantly, the CoE should anticipate known risks and propose mitigating controls that satisfy both internal and external control partners. (See Exhibit 4.)

The CoE should also create a consistent engagement model for external partners involving a single intake path for technical deep dives, control alignment, and escalation of complex cases. This speeds up learning and improves deployment quality across priority use cases.

The CoE provides the foundation for transformation across both the front and back office, allowing banks to scale AI safely and effectively across the institution.

## Leading the Way

Banks that intend to lead in the AI-native era will need to rethink their organizations across four dimensions:

- **Redesign core customer journeys.** Customer journeys must be redesigned so AI agent-enabled systems can speed up end-to-end workflows while staying within policy constraints and clear accountability boundaries.
- **Implement proper governance mechanisms.** New governance and risk practices must be put in place to support AI-assisted workflows with human oversight, auditability, and alignment to established control functions, including compliance, information security, privacy, and MRM.

- **Build flexible operating models.** New operating models must be designed for continuously evaluating, testing, and iterating AI agents as usage and data evolve.
- **Create cross-functional teams.** Set up an AI CoE that brings together experts on technology, product, operations, risk, and compliance to design and deliver AI agents safely at scale.

# Conclusion

As retail banking enters an age of AI-driven transformation, the competitive reset is underway. Banks are already moving from pilots to production, and customer expectations for faster, more personalized service are rising. Speed matters: financial institutions that act early will be better positioned to shape standards, build reusable platforms, and compound learning across use cases—and gain a considerable competitive advantage over their slower peers.

For banks that move fast, the customer journey will soon look materially different. Instead of navigating menus and forms, customers will engage goal-driven financial assistants that understand context, retrieve relevant information across systems, initiate workflows with appropriate permissions, and coordinate with human specialists when judgment or exceptions are required. The outcome is not just better self-service, but faster, more consistent completion of real financial tasks.

The transition is already happening. The question for every bank is how quickly it chooses to build the capabilities, and the operating discipline, required to lead.

# Authors

## Stiene Riemer

Managing Director & Partner  
Munich



## Marc Pauly

Partner & Director  
Frankfurt



## Bharat Poddar

Managing Director & Senior Partner  
New York



## Dan Sack

Managing Director & Partner  
Stockholm



## ABOUT BOSTON CONSULTING GROUP

Boston Consulting Group partners with leaders in business and society to tackle their most important challenges and capture their greatest opportunities. BCG was the pioneer in business strategy when it was founded in 1963. Today, we work closely with clients to embrace a transformational approach aimed at benefiting all stakeholders—empowering organizations to grow, build sustainable competitive advantage, and drive positive societal impact.

Our diverse, global teams bring deep industry and functional expertise and a range of perspectives that question the status quo and spark change. BCG delivers solutions through leading-edge management consulting, technology and design, and corporate and digital ventures. We work in a uniquely collaborative model across the firm and throughout all levels of the client organization, fueled by the goal of helping our clients thrive and enabling them to make the world a better place.

© Boston Consulting Group 2026. All rights reserved.

For information or permission to reprint, please contact BCG at [permissions@bcg.com](mailto:permissions@bcg.com). To find the latest BCG content and register to receive e-alerts on this topic or others, please visit [bcg.com](https://bcg.com). Follow Boston Consulting Group on [Facebook](#) and [X \(formerly Twitter\)](#).