

Beyond AI: Towards agentic banking




Technologies, regulations and payment architecture
shaping a new era of autonomous finance.

Table of contents

Authors' note	3
Executive Summary	4
How we studied the market – sources, method and limits of the analysis	5
What agentic banking is and where the market stands	7
Legal and regulatory boundaries	13
Technology and costs – what drives and what blocks adoption	21
New payment protocols – how Visa, Mastercard and Stripe are building infrastructure for agents	25
New threats – how criminals are turning AI against banks	32
The most profitable use cases	35
Action plan – what to prioritise and what to delay	38
Conclusion	40

Authors' note

"Beyond AI: Towards Agentic Banking" examines the most rapidly evolving transformation in the financial sector.



We have produced a practical guide that brings together three perspectives.



The regulatory perspective sets out what is permitted.



The technology perspective shows what is feasible.



The business perspective assesses whether any of it delivers a genuine return.

We deliberately challenge the widespread – and at times uncritical – enthusiasm surrounding artificial intelligence. Our focus has been on workable solutions that comply with European law and produce measurable results. Decision-makers in banks, technology firms and financial institutions should find this publication useful. It is designed to support the safe adoption of agentic banking, while preserving full control over risk, data and customer trust.



Michał Grela

Head of Growth

Executive Summary

Rather than fully autonomous agentic banking, 2026 will instead bring a wave of intelligent systems that can understand a goal, then plan and execute the necessary steps independently. These systems will, however, operate within boundaries set by legislators - acting only on the basis of customer consent and in accordance with banking procedures. As documented by the [EBA](#) and the [Bank of England](#), together with FCA findings, AI is being widely adopted across European banks, yet genuine autonomy remains minimal. Currently, only 2% of AI applications in the financial sector operate with full independence.

[Visa](#) and [Mastercard](#) are meanwhile using tokenisation and agent identity verification to build secure infrastructure that enables delegation of authority. Clearly defined limits and real-time monitoring of every transaction are central to this work.

Agentic banking already goes well beyond what advanced chatbots can do. AI agents do not merely answer questions - they can increasingly plan and coordinate complex financial processes without human intervention, though always within defined permissions and under supervision. Because these systems carry such significant decision-making responsibility, technical capability alone is insufficient to persuade customers to entrust them with their money. As [Deloitte](#) analysts emphasise, banks that succeed in this market will be those that, alongside deploying the innovation, can guarantee full control and safety - by refining their consent systems, making algorithms transparent and enforcing strict risk controls.

92%

European Union banks deploying artificial intelligence or conducting advanced pilot programs

2%

United Kingdom artificial intelligence systems operating with full autonomy while remaining systems function as assistants

+20 percentage points

Automated customer service process growth at NatWest

£100 million

Expected agentic artificial intelligence profit in 2026 for Lloyds

How we studied the market

– sources, method and limits of the analysis

This document is a pragmatic assessment of how the European financial sector is shifting towards agentic banking in 2026 and the years ahead. For the clearest possible picture, we set three independent perspectives side by side.

1.1

Three pillars of the analysis: regulation, deployments and costs

We reviewed documents from key supervisory bodies - the EBA, ESMA, ECB, Bank of England and FCA. These documents reveal the hard regulatory boundaries that limit what technology can do. The publication discusses in detail the requirements of the AI Act, PSD3, FIDA and the PLD.

We also looked at actual deployments in traditional banks such as Lloyds, NatWest, CaixaBank and ING. We examined fintechs including Revolut and N26, and the infrastructure offered by Visa, Mastercard and Stripe.

We then cross-checked these observations against analyses from leading consultancies (Deloitte, McKinsey, BIS) and real-world operating costs – including cloud spending and LLM maintenance budgets. These combined sources let us test market hype against reality and focus on solutions that are genuinely feasible, cost-effective and legal within Europe's tightly regulated environment.

1.2

Framework of the analysis: limits and assumptions

We clearly outline our assessment criteria below:

- 1. No common terminology:** The term "agentic banking" does not yet have a strict legal definition (the EBA uses the cautious phrase "so called Agentic AI"). Inconsistent market classifications of advanced chatbots versus genuine agent-based systems can distort statistics.
- 2. Evolving legislation:** This report is based on the legal situation as of 2026 and looks ahead. The main pillars (such as the AI Act) are already enacted, but implementing acts and technical standards for risk classification are changing frequently.


3. Geographic barriers: Our conclusions about the permissible agency of AI apply to the specifics of the strict European market (EU and UK). Initiatives from the US or Singapore make for interesting reference points, but they cannot be directly transferred to the EU because of different data-protection standards (GDPR) and authorisation requirements (PSD3).

1.3

Key indicators (market picture 2026+)

The following indicators illustrate where AI in banking stands today:

- 92% of EU banks are already deploying AI solutions, while the remaining 8% are preparing to start or running advanced internal tests - as confirmed by [EBA](#) data.
- 2% of AI systems in UK finance operate with full independence. Almost every machine-made decision or action still needs human verification or oversight, as both the [EBA](#) and the [Bank of England](#) confirm.
- +20 percentage points - the increase in fully automated case resolution in customer service through agentic banking, based on 21 optimised processes at [NatWest](#).



What agentic banking is

and where the market stands?

There is no separate definition of agentic banking in European regulations today. The EBA uses the careful label "so called Agentic AI," watching closely as AI takes root in banking and payments. [Deloitte](#) describes agentic banking as systems able to reason independently and carry out complex tasks within set business goals, while a [BIS](#) working paper draws a clear line between AI assistants - which support human work - and systems that automate specific tasks and partly replace people.

Agentic banking goes far beyond answering questions and making recommendations - it can pursue goals autonomously. The AI understands what a customer or employee needs, breaks a task into stages, retrieves the relevant data, uses banking tools and systems, proposes next steps, and in selected cases - within defined permissions - takes decisions without further human input. As [Deloitte](#) research highlights, it is worth distinguishing the three main models. A chatbot is reactive: it answers questions. A classic AI assistant - a copilot - already understands broader context. It can search for information, produce summaries or recommendations, even prepare next steps. But the human remains in control of the process. Agentic banking is fundamentally different from both.

Such a system plans an entire sequence of actions independently, draws on a wide range of tools and data sources, responds to unusual situations and identifies errors or rule violations. Its overriding aim is to resolve a customer's issue end-to-end - including carrying out the necessary actions and closing the process.

The system's growing authority represents a significant shift, evolving from explaining problems and making recommendations all the way to preparing and executing transactions. At the highest levels of agency, the system must meet the requirements of the PSD2 directive - including strong customer authentication and obtaining explicit consent, as [EBA](#) guidelines make clear. In response, payment-system operators are introducing new frameworks for delegated authority that allow algorithms to act on behalf of users in line with security procedures.

Chatbot

The system solely answers questions. Human operators always execute actions.

Low autonomy

**Copilot**

The system analyzes data and proposes next steps. Human operators must explicitly approve actions.

Medium autonomy

**Agent AI**

The system autonomously plans and executes operations. It operates strictly within predefined limits.

High autonomy



Artificial intelligence evolution

2.1

Widespread deployment, but limited autonomy – the European market picture

Contrary to popular belief, Europe is not starting from scratch in agentic banking. The [EBA](#) reports that 92% of EU banks are already deploying AI, and the remaining 8% are at pilot stage or in internal discussions. In direct customer service, 55% of surveyed banks use generative or agent-based models - most commonly in detecting and reporting fraud or suspicious activity, supporting call-centre agents, customer self-service, financial-literacy tools, and digital assistants or voice bots.

Many experts view the UK market as the leader in agentic banking innovation. According to a [Bank of England](#) report on AI in UK financial services, 55% of AI systems have some level of automated decision-making, 24% are semi-autonomous, and just 2% were classified as fully autonomous. The same report reveals a limitation, too: many institutions only partly understand the AI technology they use, largely because they depend on external providers of models, data and cloud infrastructure.

The picture is, at first glance, contradictory. AI adoption is advancing rapidly, but the pace at which genuine system autonomy is growing remains moderate. [Deloitte](#) and [McKinsey](#) agree that simpler solutions still dominate, and real value appears only when a bank redesigns how it works from the ground up, rather than multiplying pilot programmes. Where that does happen, the potential is considerable.

Deployment phases (EU) **92%**

deploying

8%

testing

Artificial intelligence deployment levels

Autonomy levels (UK) **55%**

standard automation

24%

semi autonomy

2%

full autonomy

2.2

Europe: what banks and fintechs are doing today

1. The United Kingdom as a leader in customer-facing agentic banking

As noted earlier, the UK is widely considered the frontrunner in agentic banking for retail. Lloyds Banking Group announced the launch of the UK's first AI-powered financial assistant, designed for 21 million users of its mobile app. The bank explicitly sets this product apart from traditional virtual assistants - it delivers personalised support and financial coaching, using agent-based systems to break customer queries into smaller steps and plan actions. Lloyds stresses the importance of maintaining high standards for AI, including effective safety mechanisms and solutions that guarantee full explainability.

The group said its AI solutions generated around £50 million in 2025, with expectations of over £100 million in 2026. More than 50 AI systems were deployed in 2025, and 2026 is set to be the year of scaling agentic banking across the whole organisation.

NatWest is on a similar path. Working with OpenAI, it announced plans to use AI for more complex tasks - from identifying fraud and guiding customers through the entire claims process, to helping them understand their financial situation and plan goals. In February 2026, the bank said 25,000 customers would get access to a new financial assistant within its Cora system.

In 2025, Cora handled 12.9 million retail banking queries, half of which required no human involvement. As reported in NatWest's technology overview, deploying agentic banking across 21 customer-service processes increased the share of automatically resolved cases by 20 percentage points. That represents a marked improvement over what "classic AI" had achieved previously. The improvement stems from the AI's ability to understand the context of a question well enough to walk the customer through an entire resolution process.

2. The European Union and Poland

In the EU, CaixaBank launched an agent within its mobile app that helps customers compare products and cards, recommends the best options, and in the next phase guides them step by step through an application or purchase. The bank reports that 200,000 customers have already used the feature.

ING deployed an AI-driven chatbot across six markets, while separately using artificial intelligence to auto-assess consumer-finance applications. A dedicated AI module also detects unusual transaction patterns during KYC checks. Rather than chasing a single all-in-one mobile super-assistant, European markets are taking a different approach. Agentic banking in the EU is built as a network of specialised processes and models.

In Poland, AI adoption in banking is accelerating. PKO Bank Polski is the most prominent example - its 17 AI voicebots have held more than 70 million conversations with customers, and the voice assistant in the IKO app serves 5.6 million users, understanding 360 conversational topics. As [BANK.pl](#) reports, internally 11,000 bank employees use an AI-powered search engine, and 1,400 developers work with a coding assistant.

[mBank](#) rolled out BOT Marek, a voicebot that since May 2025 has been calling customers to verify suspicious transactions. [ING Bank Śląski](#) launched a pilot GenAI-based assistant for business clients in December 2025 - the first such implementation within the ING group - with full deployment planned for early 2026. Polish banks are approaching agentic banking cautiously but steadily, building internal capabilities especially in customer service and fraud detection.

3. Fintechs: caution around advice

[N26](#) is a good example of a careful approach. The company deployed a smart-insights module that generates personalised suggestions, but it is explicit that these are for informational purposes only. Recommendations are labelled as AI-generated, and the user can withdraw consent for the module at any time.

[Revolut](#) follows a similarly restrained line - its assistants help customers understand products, analyse spending and support budget planning. As with N26, the company states explicitly that all content is informational and does not constitute investment, financial, tax or legal advice.


Users can also switch to a human adviser at any point. More significantly, Revolut has moved directly into the payment layer. Revolut Pay became one of the first European financial instruments to achieve full compliance with the UK and European versions of Google's Agent Payments Protocol. This allows nimble fintechs to get ahead of traditional banks and capture the commerce space served by agentic banking faster. The approach reflects a carefully considered strategy: deploying modern services while remaining cautious about autonomous advice.

Fintechs may therefore enter the agentic-commerce space faster than traditional banks, even if they remain wary of "autonomous advice."

4. Nordea and Santander – what large banks are planning next

[Nordea](#) stresses that AI should strengthen customer relationships by giving advisers faster access to data and analysis. It also notes that human oversight will keep playing a central role.

Santander, for its part, talks about scaling AI solutions and transforming both its front office and back office, and in the longer run points to the need for a more conversational banking experience.



2.3

JP Morgan, Wells Fargo and DBS – how banks outside Europe are ahead of the old continent

JP Morgan Chase is one of the most advanced banks in the world when it comes to AI deployment. The bank launched an internal assistant called LLM Suite, which within eight months of release was available to 200,000 employees. As CNBC reported, by September 2025 it had reached 250,000 users, half of whom say they use it daily. JP Morgan's technology budget rose to nearly \$20 billion in 2026, with roughly 25% of that spending tied directly to AI.

The bank is also developing agent-based payment solutions - in partnership with Mirakl, it launched an Agentic Commerce platform where AI agents make purchases on behalf of consumers.

Wells Fargo deployed an AI-based assistant called Fargo, which in 2024 handled over 245 million customer interactions - more than double the original forecast. Fargo's architecture runs on a zero-PII model, meaning no personal data reaches the language model. In August 2025, Wells Fargo expanded its partnership with Google Cloud and deployed the Google Agentspace platform for 215,000 employees, bringing AI agents into branches, investment banking and service centres.

Bank of America reports that its assistant, Erica, has logged over 2.5 billion user interactions, with 20 million people actively using it. Erica manages subscriptions, analyses spending and looks up routing numbers and transactions. It also handles transfers and payments. The bank still relies

on human oversight, though - complex cases are ultimately taken over by a staff member. This makes Erica an effective but not fully autonomous banking tool.

DBS in Singapore took a corporate-banking approach. Its AI-powered chatbot, DBS Joy, gives business clients round-the-clock virtual support and connects them with an expert for harder issues - an expert who also gets help from an intelligent assistant. From the start of testing through November 2025, Joy handled over 120,000 chats. About 4,000 corporate clients use it monthly. The bank has deployed more than 1,500 AI models covering over 370 processes, which in 2025 generated more than SGD 1 billion in profit.

Agentic banking is growing just as quickly in internal processes. Australia's CommBank, through its Coral project, links monitoring tools with code analysis. The system independently prepares change requests and proposes fixes, although engineers must ultimately approve every modification.

Key takeaways



Agentic banking does not just answer questions (like a chatbot) - it plans and executes processes from start to finish.



The UK leads large-scale deployment in Europe (Lloyds and NatWest in particular).



Continental banks and fintechs (Revolut, N26) focus on data analysis and deliberately steer clear of fully autonomous advice.



Agentic banking is no longer a technology curiosity - it is becoming the foundation of how entire organisations change.

Legal and regulatory boundaries

3.1

Regulation in Europe – the AI Act

Whether agentic banking can grow in the EU depends on meeting strict regulations and technical standards. What gets achieved in 2026 will be shaped by how well systems are adapted to EU rules on consumer protection and financial stability.

The European Artificial Intelligence Act (AI Act)

The law shaping this whole ecosystem is the European Artificial Intelligence Act (Regulation EU 2024/1689), which came into force in August 2024. The regulation's provisions are being introduced in stages, with the rules most significant for financial services taking full effect from 2 August 2026 - except for selected provisions on AI systems embedded as safety components, whose introduction has been postponed to 2027, as detailed by the [European Commission's digital strategy](#).

Banks care most about the risk-based classification of systems, and in particular the categorisation set out in Annex III. Under these rules, AI systems used to assess the creditworthiness of individuals or to determine credit scores are classified as high-risk - a classification with far-reaching consequences, as the [EBA's analysis](#) of the AI Act's implications for the banking sector makes clear.

This classification places several technical obligations on banks and providers. Mandatory human oversight severely challenges full automation. By 2 February 2026, the European Commission was required to publish detailed guidelines setting out how the high-risk classification rules should be applied in practice.

The [European Central Bank](#) lobbied for changes to the definition and scope of the high-risk classification. The ECB pushed to exclude traditional statistical models from this category. Its proposal covered generalised linear models - such as logistic or linear regression - that banks commonly

use in credit risk assessment. The ECB's argument is that, unlike deep neural networks, these models are highly explainable and transparent.



How the EBA and ESMA are aligning the AI Act with existing financial rules

In 2025, the [EBA](#) carried out a thorough analysis of the links between the EU AI Act and financial regulation, with the aim of avoiding new obligations that simply duplicate existing ones. The analysis covered in detail the relationships between the CRD, CRR, PSD2, PSD3, DORA requirements and other sector-specific rules. As [A&O Shearman's FinReg](#) commentary on the EBA mapping exercise notes, the European regulator concluded that all of these acts are fully complementary.

Deploying agentic banking means embedding intelligent algorithms into the structures that govern operational and technological risk. That kind of integration removes the need for separate control departments. Over 2026 and 2027, the EBA will standardise how national supervisors carry out oversight.

Separately, [ESMA](#) has tightened controls in investment services and capital markets. Tools used in advisory work, fraud detection and customer service fall under the EU's MiFID II directive. In 2026, as

Hogan Lovells summarises, guidelines were updated to bring algorithms within market rules.

Financial institutions must stay rigorous even as automation advances. They need to prevent discrimination and maintain data quality. Staff should avoid rubber-stamping machine-generated recommendations.

3.2

Who is liable for AI errors – the PLD, AILD and the digital fiduciary concept

Introducing autonomous solutions that sign contracts and trigger irreversible payments changes the approach to contract law. These systems independently manage investment portfolios, which forces an update to tort rules. Traditional regulations must now account for the automatic generation of declarations of intent.

European courts, legal departments and regulators are currently examining what is called the "liability gap." It stems from the nature of AI systems, which operate probabilistically and evolve over time. Neither the user nor the designer can predict their full range of effects. The EU's answer to this rests on two specific legal instruments.

The revised Product Liability Directive (PLD)

The thoroughly revised Product Liability Directive (PLD 2024/2853), having entered into force in December 2024, requires member states to transpose its provisions into national law by 9 December 2026 at the latest. From the standpoint of agentic banking technology, this act is exceptional - for the first time it explicitly extends the definition of "product" to cover software, including advanced AI systems. As [Squire Patton Boggs](#) explains in its analysis of the legal risks of agentic AI, the revised PLD is built on a strict-liability model.

An injured party - say, a bank customer for whom an AI system signed a deeply unfavourable contract through an unintended action - does not need to prove negligence or lack of care by the bank. It is enough to show that the AI system was "defective" and that harm was suffered. If an algorithm negotiates deposit terms on behalf of a customer but, due to model errors, signs off on a product with significantly worse conditions, the customer does not have to prove the bank failed to test the model - demonstrating the system was defective and that there was a loss is sufficient.

Anticipating that injured consumers know far less about how black-box models are built than financial corporations, lawmakers introduced an important procedural aid: a presumption of causation in cases involving highly complex technology. A bank deploying an agentic banking solution


becomes directly exposed to claims even when the faulty behaviour originates in the AI model. Because courts consistently reject the idea of giving artificial intelligence any form of legal personality, the full burden of material liability rests with the commercial entity controlling the deployment environment.


The Artificial Intelligence Liability Directive (AILD) and the digital fiduciary

The PLD covers damage arising from a "defect" in the product itself, but the question of pursuing compensation on fault-based grounds for harm caused by AI-assisted actions remains open. The AI Liability Directive (AILD) was designed for exactly this. It introduces a presumption of fault and obliges operators to disclose algorithmic documentation, making it easier to pursue claims. However, at the start of 2026, work on the AILD stalled. The European Commission is considering withdrawing the proposal due to a lack of political consensus. This deadlock creates a legal vacuum - if the directive fails, injured parties will have to fall back on the outdated national laws of individual EU member states, as [legal analysts](#) examining AI's regulatory implications warn.

As a result, the financial industry has put forward the concept of a Digital Fiduciary. Under this concept, banks offering autonomous advisory systems carry two core duties. The first is a duty of care, requiring the system to perform at least as well as a human adviser. The second is a duty of loyalty, requiring the algorithm to act solely in the customer's interest without promoting the bank's own products. As [ZwillGen](#) notes in its analysis of the fiduciary concept in AI, this framework is gaining traction precisely because formal regulation has stalled.

Concretely, this means writing penalties into contracts with cloud providers. Such clauses shift financial liability for AI errors and data breaches onto the technology supplier. They also cover situations where a system signs a contract on behalf of a customer without their consent - a risk dimension examined in [Baker Donelson's](#) 2026 AI legal forecast.





3.3

Open finance – the driving force behind technology, through PSD3 and FIDA

Advanced agentic-banking ecosystems in the financial sector could not develop without a deep, parallel rebuild of the infrastructure for sharing and processing data. AI systems need real-time data access by definition - without it, they remain nothing more than sophisticated conversational algorithms.

The EU's new legislative package fundamentally changes how financial information is exchanged, shifting the focus from the narrower concept of "Open Banking" to a full "Open Finance" model.

PSD3 and PSR: the evolution of the payment framework

In November 2025, after months of negotiation, the [European Parliament](#) and the [Council of the EU](#) reached a political agreement on the final shape of the Third Payment Services Directive (PSD3) and the new, directly applicable Payment Services Regulation (PSR). Both acts will likely enter into force in 2026, but given the planned 18- to 21-month transition period for technical implementation, full enforcement of the new rules on the market will come no earlier than the second half of 2027, as [Taylor Wessing's](#) fintech outlook explains.

These rules overhaul the environment for payment-initiation services and account-information services. The PSR focuses on significantly improving

strong customer authentication (SCA) protocols, introducing transparency requirements, giving non-bank payment providers - such as e-money institutions - unconditional access to payment systems, and above all tightening the liability of payment firms for fraud.

As a result, the infrastructure for intelligent client-payment solutions will integrate much more deeply with APIs, reducing regulatory inconsistencies and letting providers offer cross-border services without technical barriers, as the [Taylor Wessing](#) outlook for 2026 details.

FIDA (Financial Data Access): the information foundation for agentic banking

While PSD3 improves payment infrastructure, FIDA gives AI systems the operational data they need - a full picture of the consumer. As [Capco's](#) FIDA primer explains, the regulation takes a major step: it extends the account-focused data-sharing principles known from PSD2 to the entire range of financial-sector products. On the explicit request and with the consent of the customer, institutions will have to open access - via standardised APIs - to data on insurance, pensions, mortgages and, in time, crypto-asset portfolios.

The regulation formalises a new category of market participant: Financial Information Service Providers, who will be the main operators of advanced AI solutions. A system with a full customer profile would gain the ability to autonomously model and optimise finances - instead of just reporting payment history, it could use FIDA to spot unfavourable insurance policies and automatically calculate how much a customer would gain by breaking a fixed deposit in favour of a more attractive fund. The next step could be negotiating terms and executing the capital transfer in a fraction of a second, freeing the consumer from doing the analysis themselves.

That kind of transparency and instant data mobility worries traditional institutions, which fear losing the customer relationship to analytical algorithms and technology interfaces supplied by BigTech players.

GDPR and SCHUFA: the issue goes beyond model performance

Another significant boundary comes from GDPR rules, particularly around automated decision-making and profiling. The [EDPB](#) upholds its existing guidelines in this area. The Court of Justice of the EU's ruling in the SCHUFA case showed that the automatic generation of a credit score by a credit bureau amounts to automated individual decision-making. This holds when the bank treats that score as the deciding factor in granting credit.

More than just model quality is at stake here. What matters is whether a human actually analyses the case, or merely rubber-stamps a decision the system has already generated. If a credit analyst receives a complete recommendation from an AI and simply approves it, the machine is the real decision-maker. The human is only confirming a ready-made result.

Accordingly, in 2026, AI systems that prepare credit analyses are feasible to deploy. It is far harder, though, to implement a process in which the algorithm independently grants a retail loan - a constraint underscored by the EBA's AI Act analysis.

MiFID II and wealth management: AI on the condition of profile fit and best interest

ESMA reminds firms using AI in investment services for retail clients that they must still meet MiFID II requirements. These rules demand acting in the client's best interest and maintaining proper conduct standards. This is precisely why ESMA lists specific situations where AI may be used. Algorithms can support customer service, detect fraud, manage risk and assist in investment advice. That list, however, does not give a green light for opaque, AI-generated recommendations.

This regulation explains the caution among fintechs described earlier. The disclaimers from N26 and Revolut about the informational nature of their AI products stem directly from MiFID II requirements. Market trends indicate that wealth management and advisory work can accommodate agentic banking solutions. But such deployments require full transparency and AI governance to keep recommendations continuously aligned with the client's profile.

Key takeaways



The AI Act treats credit-scoring algorithms as high-risk systems.



Under the PLD, banks and software developers are liable for AI errors on a strict-liability basis.



The growth of agentic banking depends directly on PSD3 and FIDA, which give algorithms access to data.



GDPR still sharply limits situations where a machine can make a decision based on profiling alone.



Technology and costs

– what drives and what blocks adoption

4.1

Why agentic banking demands a rebuild of bank IT

Bringing intelligent models into financial architecture forces banks to overhaul their technology foundations. This goes beyond purchasing new software. Agentic banking needs sufficient computing power and data security. Banks must control their AI operating costs while connecting modern tools to systems that are twenty or more years old.

4.2

Cloud partnerships and ICT oversight under DORA

Processing queries through language models and handling many processes at once require large, scalable computing power. Both CPUs and GPUs are needed. That demand explains why nearly all banks are entering strategic partnerships and outsourcing AI systems to the cloud. Institutions use services from the biggest providers on the market - AWS, Microsoft Azure and Google Cloud.

Cloud providers are strictly categorised as Critical Third-Party ICT Service Providers (CTPPs) under the DORA regulation. This classification lets supervisors monitor the concentration risk that the ECB and EBA have flagged. Having almost the entire sector rely on a handful of technology giants threatens the stability of the whole financial system.

For this reason, supervisors such as Germany's BaFin require banks to implement rigorous AI governance in cloud projects. System architecture must deliver the availability written into SLA agreements and allow regular resilience and penetration testing (TLPT). Banks also need proven contingency plans. These procedures allow banks to move systems back to in-house servers quickly if an external provider fails, ensuring continuity of service.

4.3

Data-protection architecture in multi-agent environments

Banks use RAG (retrieval-augmented generation) architecture to feed isolated AI models with internal procedures and client data in real time. This setup creates a high risk of leaking information protected by banking secrecy. Errors in this process lead directly to GDPR violations.

In agentic banking systems - where one algorithm might classify a customer query while another fetches an investment recommendation - context and permission leakage, known as Capability Bleed, is one of the biggest engineering challenges.

If the process responsible for investment recommendations gains unintended access to data from the KYC module - income information or credit history, for example - that kind of context leak could breach banking secrecy and violate GDPR's data-minimisation principle.

Data security demands that AI models have no direct access to entire databases. Systems must be built on strict role-based access control (RBAC), configured directly in the mechanisms governing external function calls.

On top of that, data sent to the model must pass through filters that strip out personally identifiable information. This protects against memory-poisoning attacks. It means confidential customer data will not be permanently stored by the AI model, preventing leaks in the event of a system breach.

4.4


How much an AI agent costs – strategies for cutting expenses

Traditional SaaS customer-service platforms have a stable subscription model, with charges per user (per seat). The unit economics of agentic banking are defined quite differently - almost entirely by the cost of processing tokens through native language-model APIs.

These costs climb rapidly because an autonomous banking system does not perform one operation - it runs a multi-step analysis. The algorithm checks for errors, verifies conditions and repeats reasoning cycles, increasing the volume of data transmitted at each step. This quickly raises the cost of handling a single transaction or conversation. Sudden pricing changes from LLM providers compound the problem, destabilising the budgets of projects built solely on external APIs.

European banks have responded by adopting intelligent query routing (Agent Mesh Topology). Simple operations go to smaller, open models, while expensive, powerful tools are engaged only for difficult analyses and final verification. Accurately tracking these hidden costs is what directly determines whether the whole project turns a profit.

Cost architecture (Agent Mesh)



4.5

Legacy core banking systems as the main barrier to AI integration

Outdated IT architecture is a critical obstacle in Europe. The core banking systems of many banks were built in the 1990s or at the turn of the millennium. These solutions rest on rigid, monolithic structures that are very hard to modify. Because any change carries high risk for a massive customer base, banks are reluctant to undertake deep modernisation.

LLM models work in a completely different way - they call functions through modern API standards. Traditional, sluggish systems rarely allow that kind of smooth communication. This gap forces institutions to rebuild their infrastructure and move to flexible cloud-based solutions.

For most institutions, that means separating core banking operations from agile communication layers connected by modern data buses. Without this change, agentic banking is useless - it cannot check an account limit or retrieve customer data in a standardised format, as [FinTech Futures](#) notes in its assessment of production-scale AI agents in banking.

Key takeaways



Agentic banking runs on the cloud, which creates a heavy dependence on just a few large technology providers.



Having multiple AI models work together creates a risk that sensitive data leaks from one process to another.



Because API costs are rising, banks cut expenses by routing simple tasks to cheaper models (Agent Mesh).



Core banking systems are outdated and block the full potential of new technologies.



New payment protocols

– how Visa, Mastercard and Stripe are building infrastructure for agents

5.1

The race for a standard – payment authorisation protocols

Mass adoption of agentic banking needed universal, highly secure machine protocols capable of safely carrying sensitive payment data and purchase intent across a threat-filled online environment. The year 2025 and early 2026 brought a wave of announcements from the world's biggest players, all competing to set global authorisation standards for artificial intelligence.

Payment protocol table



Visa TAP

Network level agent identity verification at the content delivery network or edge layer



Mastercard Agent Pay

Tokens featuring hardcoded business rules and limits



Stripe ACP

Merchant assigned permissions including transaction amount and time limits



Google AP2

Multistage microcontracts requiring separate consent for cart additions and final payments

Visa Trusted Agent Protocol (TAP)

Visa's solution, the Trusted Agent Protocol (TAP), focuses on providing secure infrastructure for verifying the intent and identity of AI tools directly at the HTTP protocol layer - minimising the need for deep, costly code changes at online stores. TAP is built on a cryptographic architecture that uses digital signatures backed by public-key infrastructure.

The trust mechanism in TAP relies on several connected verification elements found in API request headers and bodies. The foundation is a recognition signature. This unique identifier confirms that the system contacting the merchant's server is an authorised, registered commercial process. It stops the algorithm from being mistaken for a data-scraping bot, which protects it from being blocked by WAF firewalls.

The second element is cryptographic customer recognition. This module passes the merchant a secure token (often a JSON Web Token) containing anonymised data about the device, location or loyalty programme. It is completed by verification of the public key stored as a kid variable.

The final layer is secured payment data. If the system is executing a payment, it transmits a cryptographic hash and the necessary information - such as a shipping address and billing data. This lets the transaction go through without exposing the full card number (PAN) in the language-model window, preventing data leaks.

A partnership with firms like Akamai and Cloudflare allows queries to be verified at the network edge. Systems can check signature validity directly in CDN nodes. This lets them reject attack attempts instantly, before any harmful data gets near the bank's infrastructure.

Mastercard Agent Pay and the Web Bot Auth architecture

Mastercard is developing its Agent Pay Acceptance Framework. The solution guarantees clear identification of every AI system making a payment, so that each transaction is authorised and traceable. The platform uses the Web Bot Auth standard (compliant with IETF RFC 9421). This lets merchants tell a trusted algorithm from a malicious bot before a payment goes through.

A central element of Agent Pay is the dynamic payment tokens generated within the Mastercard Digital Enablement Service (MDES) ecosystem. Unlike traditional card data, these tokens have security rules built in - for instance, they may only allow payment to a specific merchant, within a set time window and up to a set spending limit. As a result, even if an AI model makes an error (say, through a hallucination or an attack), the token itself prevents an unauthorised transaction. Mastercard is working with the FIDO Alliance on a global payment-verification standard. The solution

requires biometric confirmation of an operation by the consumer. This gives every command executed by the algorithm undeniable proof of informed human consent.

The platform-agnostic approach – Stripe ACP and Google AP2

From the standpoint of companies supporting e-commerce, the leading standard has become the Agentic Commerce Protocol (ACP), published by [Stripe](#) in close collaboration with OpenAI.

An online store must expose dedicated interfaces (RESTful API or MCP) to work with an external AI algorithm under ACP rules. This lets the system send commands to create and update a shopping cart.

Stripe's security model is based on generating a special token that the AI assigns to the merchant on behalf of the customer, setting a spending limit and an expiry time at the same time. A private billing module processes this token so that the LLM itself never sees sensitive data, effectively preventing abuse and permission hijacking. Stripe supplemented the protocol with technical guidelines requiring stores to publish configuration files that act as machine-readable dictionaries defining permissible navigation rules.

Google's Agent Payments Protocol (AP2) follows a different philosophy: an advanced construction of cryptographically secured authorisations that act, in effect, as a series of micro-contracts. The first step is an authorisation allowing the AI process to search for options within a set budget. When the solution picks specific products, a separate record approves the cart. Only then does the appropriate signal trigger the authorisation network and close the whole sequence.

The agentic commerce ecosystem is also being built by companies outside the traditional card networks. [Klarna](#), the European leader in buy-now-pay-later, launched an AI assistant built on OpenAI technology. In its first month, the algorithm handled 2.3 million conversations, which was two-thirds of all chats in the support department. The tool did work equivalent to 700 full-time agents and cut problem-resolution time from 11 minutes to under 2. As [CX Dive](#) reports, by Q3 2025 Klarna's AI assistant was doing the work of 853 agents, generating annual savings of around \$60 million.

[FIS](#), one of the world's largest financial-infrastructure providers, in January 2026 launched the industry's first agentic-commerce platform allowing banks to securely process transactions initiated by AI agents. The platform is built on a Know Your Agent (KYA) framework for identifying AI systems and works with Visa Intelligent Commerce and Mastercard. According to [Payments Dive](#), AI adoption among banks has grown eightfold compared with 2023.

5.2

Visa and Mastercard – same goal, different methods of building trust

Visa and Mastercard are not banks, but they are working on the hardest problem in agentic banking. The challenge is letting an algorithm act on behalf of a customer without breaking authorisation procedures, liability rules or fraud controls. The solution must be transparent for card issuers and convenient for merchants.

Deloitte points to Visa, Mastercard and PayPal as players experimenting with next-generation commerce. The next wave of AI autonomy in finance will be built on settlement systems and credentials supported by communication protocols. The language models themselves become secondary.

VISA

Edge authorization via
Cloudflare or edge
computing and the TAP
protocol

Tokenization ensuring artificial
intelligence systems lack access to
real card numbers, hard
transaction limits, and full
immediate customer consent
revocability



Fast Identity Online
standard for verifiable
intent and trusted
agentic token
concepts

Visa: from card tokenisation to a new economy

Visa Intelligent Commerce's test environment launched in April 2025 in partnership with Anthropic, IBM, Microsoft, Mistral AI, OpenAI, Perplexity, Stripe and Samsung. Broad market rollout is planned for the current year, 2026. The list of partners shows the scale. Thanks to this solution, agentic banking tools from different ecosystems can directly execute transactions through Visa's infrastructure.

Instead of passing raw card data, users rely on digital credentials built for AI handling. The system acts on behalf of a customer only after their explicit activation. The consumer keeps full control, setting spending limits and conditions for each operation. By transmitting commerce signals directly to the Visa network, the institution can effectively block fraud and settle disputes. The whole mechanism is based on delegated authority - each permission has an assigned cost policy, which the account holder can revoke at any time, as described in Visa's Intelligent Commerce announcement.

Technical documentation from [Visa](#) indicates that AI processes require strict budgets and tokenisation, because the market is still working out the rules for liability when errors or misuse occur. Visa requires the ability to instantly revoke or modify access when transaction conditions change. Human oversight is essential, particularly for large amounts or legally restricted operations. Even the biggest infrastructure providers accept that full autonomy without a safety brake is not the end goal.

In October 2025, [Visa](#) and Cloudflare introduced the Trusted Agent Protocol. This lets stores tell verified AI systems apart from harmful bots. Testing already involves giants like Stripe, Shopify and Adyen. The change matters for card issuers because automated payments need to confirm two things at once: customer consent and the credible identity of the software placing the order.

In September 2025, [Visa](#) released its MCP Server and Acceptance Agent Toolkit. This set of tools lets AI models connect securely to the Visa Intelligent Commerce system. It also supports no-code modes, enabling invoicing or generating payment links through simple text commands. Visa is building a complete operational platform for agentic banking, moving beyond a simple payment product. In pilot deployments - including work with Aldar in the United Arab Emirates - AI models are already independently paying recurring bills, such as service charges.

Mastercard: Agent Pay and the attestation architecture

The [Mastercard](#) Agent Pay Acceptance Framework requires every AI instance to be registered and verified before executing a payment. Systems must be unambiguously identifiable, and the operations themselves use dynamic credentials that guarantee traceability. Merchants can recognise authorised processes at the data-transfer stage through the Web Bot Auth standard. Rather than forcing a wholesale rebuild of the internet, Mastercard embeds verification mechanisms into existing payment forms using special codes.

The company is building a clear model for verifying software because simply having card data is no longer enough. The network now checks which specific system is making the payment, on what basis, and with what spending limit. This direction shifts traditional "yes or no" authorisation into a complex assessment of the context in which the customer delegated authority to a machine.

Mastercard is combining these solutions with biometrics and passkeys through its Payment Passkey Service. Partnerships with Microsoft, IBM, Braintree and Checkout.com allow these features to be deployed in both retail and B2B settings. Payments triggered by automated systems are, for Mastercard, the next step in moving away from manually entering data towards programmable credentials.

What Visa and Mastercard share

Visa and Mastercard are pursuing the same objective: conditional permissions. They differ in method - Visa relies on edge-network verification and the TAP protocol, while Mastercard favours tokens with embedded rules and the FIDO standard. Both share a firm principle, however: an AI system receives only as much authority as the customer or bank grants, and any consent can be revoked immediately. Unrestricted access to funds is ruled out.

Three conclusions follow from this strategy. Agentic banking will grow faster in online commerce and on cards than in ordinary bank transfers, because the infrastructure there is already ready. Banks can strengthen their position if they adopt new verification systems and rigorously enforce access policies. Success will be decided by effective control over what the customer has authorised, and for whom - the language models themselves will be secondary. As Mastercard's Agent Pay announcement makes clear, this control infrastructure is the real competitive differentiator.

The wider ecosystem: PayPal, Stripe and commerce protocols

Visa and Mastercard operate alongside many competing solutions. Stripe is developing scoped payment tokens and the ACP protocol, while PayPal is building commerce services on its own infrastructure, combining identity verification with buyer protection. PayPal works with Mastercard and AI-tool developers. From a European angle, Revolut Pay's move towards the AP2 standard also matters. Payments in the agentic banking model will probably draw on several protocols at once. A closed model tied to a single provider's technology is unlikely.

Key takeaways



The market is creating new payment standards for machines: Visa TAP, Mastercard Agent Pay, Stripe ACP and Google AP2.



Tokenisation is the bedrock of security, hiding sensitive card data from artificial intelligence.



New solutions let a machine execute a payment only within a set time window and pre-agreed limits.



Agentic banking is easier to plug into card-payment infrastructure than into traditional bank transfers.



New threats

– how criminals are turning AI against banks

6.1

Adversarial AI – the evolution of threats to digital infrastructure

Hacking methods are changing along with the wider economy. In 2026, rigid mass-automation scripts are giving way to autonomous, decentralised attack systems. Traditional rule-based defences are failing because criminal activity increasingly mimics the traffic of authorised users and internal banking tools.

New categories of AI-driven fraud include the following:

Polymorphic phishing involves hijacking a company's email domain. Rather than sending crude messages with suspicious links, the system spends months analysing correspondence history. This lets the algorithm learn the victim's writing style and internal jargon, so that at a critical moment - such as during payment approval - it can generate a forged message with a swapped account number. Defences that only react to infected pages let such a message through, treating it as a natural part of a genuine business thread.

Predictive payment impersonation uses spyware to analyse the victim's payment calendar, striking precisely when settlements are due. A malicious algorithm sends a message suggesting a change in invoice terms just before the real document arrives from a regular supplier. The resulting confusion in accounting departments means the forged data passes through checks without proper scrutiny.

Hardware-assisted KYC bypass methods and **deepfake-as-a-service** offerings effectively undermine the credibility of remote onboarding and biometric identification. Attackers use specially prepared virtual systems that trick financial apps by feeding a high-resolution fabricated image directly into the phone's system camera driver.

Synthetic identity farms built by autonomous AI models spend months constructing credible histories for fictitious accounts. The algorithms make small repayments on time and manage deposits to build up high credit scores. Once the system's trust is won, the synthetic identities carry out a coordinated strike - simultaneously taking out loans worth tens of millions of euros and moving the funds through dispersed international transfers.

Prompt injection hits the conversational layer of banking assistants directly. Bots that lack proper filters execute instructions hidden in user queries. Such attacks force systems to bypass financial blocks or disclose sensitive data, leaving the bank with full legal and reputational liability.

Over the 2025–2026 period, financial organisations began replacing one-off identity verification with continuous, real-time monitoring of operations. The new security standard requires authorisation of the payment tool itself that is integrated with the banking system - a practice described as Know Your Agent.

New adversarial artificial intelligence cyberattack vectors



Polymorphic phishing

Artificial intelligence generates emails perfectly replicating the writing style, industry jargon, and correspondence history of the victim ensuring flawless grammar and perfect context



Predictive impersonation

Precision attacks where artificial intelligence strikes during a specific settlement window posing as a trusted vendor exactly when accounting departments expect an invoice



Deepfake KYC Bypass

Injecting artificial intelligence generated video directly into virtual smartphone cameras to bypass liveness verification checks



Synthetic farms

Automated bots massively building perfect credit histories for fabricated identities to execute a singular massive fraud after one year

6.2

ESG reporting and sustainability (CSRD) in the face of AI deployment

Behind the apparent simplicity of a digital, cloud-based AI deployment lie significant environmental costs, making sustainability one of the more demanding challenges for bank auditors as the EU presses for lower carbon footprints. The new EU Corporate Sustainability Reporting Directive (CSRD) widens the circle of institutions covered and imposes strict, mandatory disclosure standards (ESRS) linking financing and business activities to climate indicators on a large scale from 2024 - and from 2026, this extends to non-financial entities in Europe connected to the sector.

Cloud-service providers generate enormous energy demand, expected to grow from 25 GW to as much as 120 GW by 2030. This level of operation places a heavy carbon and material footprint on the servers that banks rent to run AI platforms. These emissions fall squarely within Scope 3

calculations, which cover the complex web of links beyond the financial institution's physical control. Under CSRD rules, banks must report and carefully estimate all these indirect logistical costs.

Banks that are required to report Green Asset Ratio (GAR) indicators are, however, increasingly turning to algorithms for help. Systems using RAG technology rapidly check lists of business clients, analysing thousands of pages of financial reports and audits without involving staff. This automation, however, consumes substantial computing power, generating high energy demand. These hidden energy costs mean banks must calculate their real environmental impact more precisely if they are to meet EU standards and pass ECB inspections.

Key takeaways



Traditional defences are sometimes powerless against attacks run by AI.



Cybercriminals strike with precision - hitting on specific settlement dates and perfectly copying the victim's email style.



Fraudsters use deepfakes on a massive scale to bypass cameras during identity verification.



AI in the cloud consumes vast amounts of energy, which hurts banks' statistics and makes ESG reporting harder.

The best examples from the market

In 2026, the best results come from deployments that combine large scale with concrete returns, fit within clear frameworks and are not subject to the strictest legal requirements. The standout areas are customer service, fraud prevention, internal processes and AML controls. Full autonomy in credit or wealth management still runs up against too many procedures, as [EBA](#) research confirms.

Here are the areas with the greatest business potential:

Fraud detection and claims.

This is currently the most effective application of agentic banking. Banks have clear performance indicators here, and algorithms can independently gather data, block suspicious operations and prepare a case for a staff member. [NatWest](#) is building along this line, among others.

Day-to-day service in the app.

Explaining charges, analysing history, managing subscriptions or changing card limits - these are tasks where AI copes without staff help. Systems like Erica ([Bank of America](#)) and Cora ([NatWest](#)) demonstrate that most routine queries can be closed without human involvement.

Budget management and financial coaching.

This is a priority for retail banking because of the high value it delivers to customers and the low deployment cost. As long as the system does not pose as a certified adviser, it stays on the safe side legally. Solutions are being developed by [Lloyds](#), [N26](#), [Revolut](#) and Bank of America, among others.

Staff and back-office support.

Automating office work delivers the quickest return on investment. Algorithms summarise calls, search for procedures and draft responses, which cuts service costs significantly. As [McKinsey](#) research shows, properly redesigning these processes can produce step-change gains in productivity.

KYC verification and AML alerts.

AI models are ideally suited to multi-stage document analysis and detecting unusual transfers. [ING](#) uses these technologies to track suspicious patterns, and [Deloitte's](#) analysis points to anti-money-laundering as one of the most advanced areas of AI use.

Credit-process coordination.

AI does not yet make independent lending decisions, but it takes over the entire bureaucracy. Systems collect documents, fill in forms and prepare data packages for credit assessment. [McKinsey](#) estimates this kind of support raises work efficiency by 40 to 80%.

Virtual assistants for corporates and SMEs.

The experience of [DBS Joy](#) confirms that in corporate banking, machines can handle most of the ongoing relationship with a client. Under this model, the system resolves routine problems, and for complicated cases it hands the client over to a specialist who benefits from the system's suggestions. In Europe, this approach brings down the high cost of serving small businesses while still giving them fast answers.

Low-risk automated payments.

Modern purchase-authorisation methods are not yet widespread in the EU, but they represent a key strategic direction. In 2026, the most realistic scenarios involve payments to trusted merchants, defined purchase categories with spending limits, and paying household bills. This is the direction [Visa](#), [Mastercard](#), [Stripe](#) and [Revolut](#) are all heading.

Profit \ Readiness	Readiness		
	Low	Medium	High
Low			
Medium			Quick implementation Summary of expenses, delegation of fixed household bills, coordination of loan applications (data collection).
High	Regulatory blocks – Let go Fully autonomous decisions on loan approval, autonomous management of assets and investments.		Do it Fraud detection, call centre assistant, daily support in the application, preliminary KYC/AML analysis.

Key takeaways



Agentic banking works best today in fraud detection and handling high volumes of customer queries.



Rather than giving advice (which is legally risky), AI should focus on spending analysis and guiding the customer.



AI as a call-centre assistant cuts costs and avoids the risk of a machine taking a poor decision autonomously.



AI can gather documents and prepare a credit process, but the final decision must always be made by a person.

Action plan


– what to prioritise and what to delay

Rather than building one assistant for everything, banks should split permissions across four stages: data analysis, generating recommendations, preparing an operation, and executing it within limits. This division makes it easier to identify who is responsible for errors and to carry out audits - which is critical under the AI Act and under EBA standards for payment services.

When designing products, it is best to focus on high-volume, low-legal-risk tasks. Fraud detection, everyday service, spending analysis and customer verification (KYC/AML) are the areas where it is easiest to measure returns and set hard limits for algorithms, as EBA data on AI deployment in the sector confirms.

In payments, banks need to create their own consent-management system built on tokenisation, modern access keys and limits. It is important that these solutions go into the bank's main systems from the start. Isolating them as side "innovation projects" means they quickly fall out of step with real market demands - a point underscored by the ECB's instant payments regulation, which sets clear infrastructure expectations. These consent frameworks must also align with the principles that Visa has established for its Intelligent Commerce infrastructure.

The ability to instantly interrupt an automated process and have a staff member take over must be available



from day one. Adding these safeguards only after a failure sharply raises the cost and risk of deployment. This principle is consistently emphasised by Lloyds, DBS and Bank of America in their respective deployments, and reflected in BIS research on agentic AI, as well as in Deloitte's recommendations for banks adopting agentic systems.

Clear separation is also needed between financial education and investment advice, and between data collection and final credit decisions. Disclaimers in terms and conditions are not enough - the product must be designed so that its evolution does not threaten business stability. ESMA has made this a central element of its AI guidance for investment firms.

At the strategy level, cloud outages or model errors should be a topic for the board and risk committees, not just for the IT department. Since innovation depends on external suppliers, banks need contingency plans ready and must enforce quality rigorously at every stage of the partnership - a concern the EBA raises directly in its survey of AI deployment across EU banking.

Key takeaways



Banks should deploy AI gradually - from simple analysis through to strictly limited payments.



Product teams should focus on simple, high-volume operations where a machine mistake costs little.



Banks must design systems from the start so that a human adviser can stop the machine at any moment and reverse its decision.



Dependence on cloud and AI suppliers is now a board-level concern, not just an IT matter.

Conclusion

The success of agentic banking technology depends on trust, which means that deploying it alone does not guarantee mass adoption. Without customer acceptance, even the most autonomous solutions will not gain traction. In 2026, the strongest prospects lie with deployments that take over specific tasks within clearly drawn boundaries. Fully replacing the banker remains a distant vision - and a legally impossible one - so banks are concentrating on fraud detection, customer service, KYC/AML verification and case-stage management. The directions set by [Visa](#) and [Mastercard](#) confirm that the future of automation in finance rests on precise permissions that the customer or bank can withdraw at any moment - as [EBA](#) research on AI adoption in EU banking underscores.

The institutions that come out ahead will be those that best refine how they delegate tasks to algorithms. This approach will determine how quickly and safely a bank can expand AI responsibilities as acceptance grows among regulators and customers. Over the next two to three years, the most important test of maturity will be the ability to manage these permissions smoothly while keeping every decision fully transparent - a standard the [EBA](#) has already begun formalising in its supervisory guidance.