

The EU accidentally wrote the AML industry's AI strategy

AMLR doesn't mandate the use of AI. But its requirements lead to exactly the architecture AI needs to be useful.

Peter Lange

Head of Sales, Pingwire

This summer, Johan Övergaard, Manager Transaction Monitoring at Swish, shared an article with me that had been published in FinTech Global. It asks whether our defences can evolve at the same pace as AI-driven financial crime (see *Is financial crime entering an AI arms race*).

In short, the article's answer is yes: the arms race is already under way. What made it refreshing reading, though, is that it doesn't settle for the easy conclusion that autonomous AI will replace both system infrastructure and people and solve AML and fraud prevention flawlessly. Instead, it focuses on the practical application of AI: how the technology actually fits into existing methods and what is happening right now in AML and fraud. That is the perspective I wanted to explore further, and this article is the result.

I also believe the EU has created a regulatory framework that drives AI adoption, most likely without intending to. AMLR doesn't mandate the use of AI. But it does require broader customer due diligence, that more data points are actually used in ongoing monitoring, synchronous real-time monitoring where feasible, and traceability at every step. Those requirements happen to lead to exactly the architecture AI needs to be useful: data that is consolidated, connected, accessible in real time and traceable.

As Head of Sales at Pingwire, I work every day with AML and fraud systems built for companies covered by anti-money laundering regulation. That means I meet many companies and institutions facing challenges as they try to renew their ways of working and their system infrastructure to fight financial crime more effectively.

Existing rules aren't useless, but they aren't enough

Static, non-sequential rules with fixed thresholds (e.g. rolling sums, amount limits, net value of funds moved in and out) are not sufficient against schemes designed specifically to evade them. They still serve a purpose by setting the boundaries. What is needed are smart rules that continuously learn each customer's behaviour and flag deviations from it: rules built on machine learning that also connect data to build a profile. This is something many of the companies and institutions I meet are asking for.

This is also a central point in the FinTech Global article: crime is designed to blend in. A transaction that looks completely normal in isolation only gives itself away in other ways: against the customer's own history, against the behaviour of comparable customers, or when data is connected at network level. A money mule account can imitate a normal cus-

tomers in a snapshot. It is much harder to simultaneously fake a consistent history over time and keep every connection plausible, including transactions, devices, IP addresses, counterparties, ownership structures and more.

AI is already delivering real value in AML processes

The market is also asking for AI support in case management. I think many have realised that case management is a repetitive process where similar typologies are investigated over and over, which makes it ideal for LLMs to streamline. This applies both to the pre-investigation presented to the analyst and to tailoring it to each company's specific investigative needs. The process has several steps that increase efficiency and can also make investigations more legally robust by reducing manual errors: 1. Data compression, 2. Data analysis, 3. Risk identification, 4. Recommended investigative steps and 5. Outcome. It helps the investigator quickly understand the context of what has happened, what the data looks like and how it connects, and what needs closer investigation. It also produces clear, explainable investigation notes to store in the case.

But it doesn't stop at pre-investigation. We are already seeing AI agents used more broadly in operational work: to prioritise the case queue so that the alerts carrying the highest risk rise to the top, to calibrate rules and thresholds, and to dig into data and pull together the material an analyst would otherwise gather manually from several different views.

At Pingwire, this takes two concrete forms. The first is our AI case analyst, which reads through the material in a case and delivers a short description and summary directly in the case, tailored to the institution's guidelines. It is enabled per account and governed by permissions, so only the roles that should have the feature can use it.

The second is Copilot, an AI assistant that accompanies you throughout the platform. It can answer questions about how the system works and how rules and risk classes are set up, and retrieve things like customer data, cases, transaction statistics and hit history for PEP, sanctions and adverse media. Copilot works with the logged-in user's own permissions, which means it can never see more than the analyst is allowed to see. And every answer passes an additional, isolated AI check before it is shown to the user.

In my view, this is how AI support needs to be built in this domain right now: with broad access to data, but narrow and traceable authority, controls, and a decision that stays with a human. Things are moving fast, however, and even though supervisors are currently pushing back and holding firmly to the principle that decisions should rest with humans, I believe we will soon see decisions actually handed over to AI. Not just AI in an advisory role, but a real increase in the degree of automation in AML work. What makes this possible is that the technology for governing AI already exists. AI governance in practice is about being able to write policies, set technical guardrails, manage access rights and ensure traceability. It doesn't limit AI, it steers it. Above all, it lets you define where the line goes: which decisions

AI may make on its own and which must still pass through a human. This was also an insight from Nordic Fintech Week (NFTW), where banks, lawyers and industry professionals reasoned along exactly those lines.



It doesn't limit AI, it steers it."

Another topic that is increasingly being discussed is cost. Every AI session consumes tokens, whether it's a case analysis, a question about customer data or a question about how something works in the platform. In Pingwire, that consumption is measured per session, and that's deliberate. Without measurement, you can't answer the only question that really matters: what does an AI-handled case, or a question to AI about how the platform works, cost? And what did it cost before, when the case was analysed manually and the user had to search through a manual?

That is where the value becomes concrete. A few kronor in tokens instead of twenty minutes of analyst work is a calculation anyone can run and track. And the cost per case should fall over time as the data model, the prompts and the way of working mature.

The regulatory direction is clear

The statements I have picked up point to a clear regulatory direction. EU-level authorities such as AMLA, as well as national supervisors, encourage the use of AI but emphasise AI governance and that decision-making should rest with humans. Although Recital 58 of the AI Act excludes AI systems used to detect fraud from being classified as "high-risk AI", the combined effect of legislation such as the AI Act, GDPR and DORA is that AI in AML and fraud work in practice needs:

- **Risk assessment:** a documented assessment of the AI system's risks before it goes into production.
- **Documentation and traceability:** a thorough description of what the model does and what data it relies on, plus a log showing what formed the basis for each outcome.
- **High data quality:** assured quality and representativeness in the data the model is trained on and operates on.
- **Transparency and explainability:** the ability to justify afterwards why a specific outcome or flag occurred.

- **Human oversight and involvement:** a human makes the actual decision; AI supports it.

How much the AI Omnibus and Digital Omnibus will ease this remains to be seen. European policymakers have, after all, realised that the EU has fallen behind in the global AI race.

For the foreseeable future, I at least don't see an autonomous AI agent with limited human oversight running the entire cycle: carrying out the business-wide risk assessment, implementing models and scenarios, automatically investigating and deciding cases with consequences for payment execution and individual customers' access to financial products, reporting SARs/STRs to FIUs, back-testing, and then going back to adjust the business-wide risk assessment, and so on. What we do need, right now, is AI support to speed up this cycle so we can catch up with criminal money flows. Criminals can use AI systems without restraint to find and test new ways to launder money and commit fraud.

Criminals are already operating at AI speed

Digitalisation and new technology have also changed customer behaviour: online banking means KYC is completed digitally in minutes and payments are executed internationally within seconds. LLMs can create synthetic identities, forged KYC documentation and tailored cover stories. They can also probe AML systems to map thresholds and find gaps to exploit. Techniques that used to require time, expertise and coordination can now be carried out quickly, at scale and with a high degree of realism.

Traditional systems with static rules, a fragmented architecture that leaves data in silos, and daily batch processing of transactions simply can't keep up with AI-generated criminal flows. AML teams are drowning in false positives and can't see the full risk picture. The underlying architecture of traditional AML systems, built on fragmented set-ups and siloed data, is no longer fit for purpose.

The tempting shortcut and why it isn't enough

It is tempting not to replace the infrastructure at all, and instead add agentic AI as an intelligent layer on top of existing systems to streamline and improve case management by extracting, compressing and surfacing the data available across those scattered systems. For institutions with an entrenched architecture that rely on traditional systems, this could be seen as a leapfrog. Analysts no longer have to switch between numerous systems and browser windows, gather information manually and piece together connections from fragmented data sets, a time-consuming process with a heavy cognitive load. But it doesn't reach the root of the problem: the scale, speed and sophistication of AI-generated financial crime.

To get there, you have to start at the foundation: the data model. It needs to be flexible, in the sense of easy to change, and able to take in large volumes of data points of different kinds without rebuilding the system every time a new data source is added. On top of it sits an API that mirrors the data model, making it easy for the institution to send in the data it has and actually wants to use. And within the model, data must be connectable at entity level, so that a rules engine can calculate on the relationships fast enough for synchronous real-time monitoring.

Only then can you truly benefit from agentic AI in AML and fraud systems, apply behavioural analytics and stop payments in real time. Both behavioural analytics and stopping payments in real time, where possible, will become requirements under AMLR and its accompanying guidelines (see, among others, Article 26 AMLR and AMLA's guidelines on ongoing monitoring of a business relationship).

Traditional systems were designed when a customer was a row in a database table and monitoring ran in overnight batches. A new system doesn't have to account for decades of changes to that structure, or for all the integrations built on top of it that in turn required layers of data cleansing and transformation to fit together.

The uncomfortable question

The FinTech Global article argues that the next generation of crime prevention will be determined less by any single breakthrough technology and more by the ability to connect and operationalise a broad set of technical capabilities in an AI-native way. I share that view, but it forces an uncomfortable question for those of us who build AML and fraud systems and deliver them as SaaS. In a world where agentic AI drives efficiency and companies can prompt their own features with tools such as Cursor, GitHub Copilot, Claude Code or OpenAI Codex, what is our reason to exist?

Our answer, and in my view a do-or-die for all B2B SaaS, is that expertise must go deeper within your own domain, not broader across more. We need to be experts in AML, both regulatory and operational. Experts in AI, leading the way in testing and applying the latest LLMs to AML and fraud work, and experts in embedding those models safely and in a regulatorily sustainable way into a system built for AI from the ground up. That way, customers know we will always make development better and more efficient.

This brings another B2B SaaS trend with it: custom solutions, tailored to the customer, far from the one-size-fits-all model that has traditionally defined B2B SaaS. When the AI tools mentioned above make it possible to develop, design and launch a prototype in a short time, the same adaptability will be expected of us as B2B SaaS companies. For us at Pingwire, that means a product vision built on a robust, non-negotiable core: a flexible data model with an API that mirrors it, and a rules engine that performs powerful calculations and responds to calling systems in milliseconds. Above the core sits a fully customisable layer.

Customers can order different versions of this layer, or connect their own AI agent via Pingwire's MCP and change whatever needs to follow their own process, such as case statuses, fields and workflows.

The same MCP also enables the reverse: working with Pingwire's structure, processes and data directly in the interface of whichever agent you prefer. As recently as NFTW, several banks said in panel discussions that **the best interface is no interface**.

Build vs Buy

But if building with AI is so accessible, why not do it yourself? Especially in a world where AI gives instant access to specific knowledge such as syntax, API documentation and regulation at negligible cost. Just as you can buy software, couldn't you acquire the necessary regulatory and programming knowledge yourself to build the AML process you want, with the traceability and explainability needed to hold up in supervision and audit?

I believe there are three arguments against this: AI falls short on consistent behaviour, control and sufficient quality. Consistent behaviour means the AI system maintains the same logic, tone and behaviour across many different tasks. That's a challenge because AI output is based on probabilistic text generation, which makes exact, repeatable results harder to guarantee. Control refers to the fact that when we use AI tools, we fundamentally rely on prompts, which makes the result hard to control. Many have probably experienced how adjusting a small detail in an AI output quickly turns into a time-consuming, difficult task. Consistent behaviour and control are technical problems, and progress is so fast that I think they will be solved soon. Quality is a different kind of issue: it's not about the model's capability, but about who makes the trade-offs.

When it comes to quality, AI can imitate and reuse existing work, but it can't apply genuinely human taste or judgement. For lack of a better term, and in line with the phrase that has become established, you could say AI lacks personal taste.

My conviction is that as long as a human is needed to make trade-offs, frame the right problem and see the whole picture and the vision, specialists in software/AI and in the subject domain will be needed. If you accept that, nothing in the Build vs Buy question has really changed since before AI arrived. You could always build yourself if you hired the right specialists, or choose to buy software from a vendor that already has specialists in the field.

And this is where the advantage of buying specialist expertise becomes clear. It's not just about writing code and applying AI, but about someone having already built, tested and defended the risk assessment, documentation, traceability and human oversight listed above, before multiple supervisors, across multiple institutions, over time. That experience can't be prompted into existence in an afternoon.



That experience can't be prompted into existence in an afternoon."

Architecture decides how well AI can be used, and AMLR is the catalyst

The arms race is real, and we need to use AI to stop criminals from laundering money. Who succeeds will not be decided by who has the most autonomous AI agent, but by who builds the architecture, and the combination of human and AI, that actually holds when money moves very fast.

Companies and institutions subject to anti-money laundering regulation must implement AMLR by 10 July 2027. Many have already started, but for anyone with experience of implementing regulation, the deadline is tight. For a regulatory project of this size, the application date is in practice already just around the corner.

AMLR contains much of what this article has discussed. Customer due diligence requirements become broader and stricter, and ongoing monitoring requires more data points to be used. On top of this comes the requirement for synchronous real-time monitoring as soon as an organisation is able to do so, along with expanded reporting obligations, where AMLA itself is investing heavily in the infrastructure to receive large volumes of data.

What I hear in meetings with companies and institutions is that their AMLR gap analyses reveal a need to broaden their KYC standard. More data points must be collected and used effectively in ongoing monitoring. Many then realise that their existing processes are not fit for purpose and that the system support is missing.

So companies and institutions are looking for systems with a flexible data model and an API that mirrors it, so data can be received and processed fast enough for real-time monitoring. And this is where something interesting happens. As a side effect of compliance, they put in place exactly the architecture needed to get the most out of AI.



As a side effect of compliance, they put in place exactly the architecture needed to get the most out of AI."

The database is more than a storage place. An AI agent can't work with data it can't reach, can't interpret relationships that aren't modelled, and can't leave a trail in a system that doesn't record what happened. The data model and its API therefore become the layer agents must pass through: their operating system, the memory they draw context from, and the history that makes their work auditable after the fact.

It is in this sense that AMLR acts as a catalyst for better use of AI in the future, whether or not AI was at the top of the institution's agenda.

I'm curious how others in the industry see this. How are you thinking about applying AI in AML and fraud work, and how much change to processes and system support do you think AMLR will demand of the industry as a whole? Share your perspective in the comments or reach out to me directly.

Peter Lange

Head of Sales, Pingwire

>pingwire_

Source: FinTech Global, "Is financial crime entering an AI arms race".