

AI's impact on post-trade

Thomas Steinborn, chief product and technology officer at Smartstream, speaks with Matthew Challis about how AI is changing post-trade operations, why the biggest barrier to modernisation sits outside technology, and why humans stay accountable as autonomy grows

Can you give us some background about what you do at Smartstream and what it entails?

I am chief product and technology officer at Smartstream. I joined about a year ago from a background in data and AI, and my focus has been the same throughout: how we bring AI and data maturity to our financial services products. At Smartstream, that means applying them to post-trade operations, where we draw on decades of domain knowledge.

What has been the main theme that has dominated Sibos this year in terms of the conversations that you have had?

Most of what we do is post-trade operations, and the biggest part of our portfolio is reconciliation. The biggest theme is, unsurprisingly, AI. Everybody is talking about it, and the conversation has moved on. People have played with it, many of them with us, and now want to take their pilots into production and realise the benefits. The other question is the business case. This is not about putting AI everywhere; it is about where it creates value. In post-trade, that is operational efficiency. About 98 per cent of transactions clear straight through, but the remaining 2 per cent, the breaks and mismatches, absorbs roughly 70 per cent of operational effort. We focus on shortening cycle times for resolving breaks and running investigations, whether with other departments of the bank or all the way out to counterparties.

In collateral management and other post-trade capabilities, we are also seeing more conversation about tokenisation of assets, and how it helps banks lower margin costs, improve cost ratios, and reduce risk.

Is there a discrepancy between what banks are saying to you about AI and what they are saying on the commercial side of things? Is there a difference between the institutional conversation and the public, commercial conversation surrounding AI?

Publicly, the message is consistent: AI is great, we have a plan and we want to bring it in. When we work directly with clients, the picture is more nuanced. Most have built their first agents and are trying to move them into production, with or without us. Roughly 15 per cent of our customers are at that point, which is in line with the industry overall. The US is a little ahead, APAC follows, and Europe is a little behind. We also see a split between the Tier-1s moving ahead versus everyone else.

Everybody is bullish on the outside and more thoughtful on the inside. Much of that is about risk and compliance, because agents are not yet fully manageable on their own. The recent news involving OpenAI and Hugging Face is exactly what a bank never wants to see in its own environment.

Other than prohibitive costs for clients, do you think sometimes the fear of hallucinations and AI misconduct is a concern?

It is far less of a concern than it was. I would have given you a different answer a year ago. These models are probabilistic: they give you the most likely answer. But we can now manage that risk much better, for two reasons. The models have advanced, and Smartstream brings the knowledge we have built up over decades of reconciliation as context. On top of that, we have put guardrails in place. We do not simply let agents reason and decide on their own; we combine them with our existing deterministic technology stack. That balance turns it into a manageable risk, as opposed to something you could not really control last year.

With all of the things that have been introduced over the past year, do you think the adoption of AI technology is moving at a greater pace than you expected, or is it staying about the same? Is it scaling relative to the guardrails and safeguards being put in place?

It is very much within our expectations. As a vendor, we always want to move faster and stay a little ahead of the curve, but we partner with our customers and meet them where they are. Some are ready and we are moving together; with others we are still defining the business case. Our focus is aligned, and things are moving as planned.

What has been the biggest operational challenge that banks have presented to Smartstream?

Outside AI, it is market pressure. We have moved to T+1 and some markets are looking at T+0. At many Tier-1 banks, including our customers, the operations team in London also serves Singapore and New York, so round-the-clock operation is the focus. Independent of T+1, clients also want near real-time processing

because it lowers their margin costs. There is real pressure to operate in near real time whether or not regulators require it.

Banks do not change systems on a whim, so we are taking a nuanced approach. We are making our products much easier to upgrade. Part of that is AI-related, though not because of adoption: AI makes it much quicker to find security vulnerabilities, so our products have to be easy to maintain, update, and upgrade. At the same time, customers still need to make changes through their Sarbanes-Oxley Act (SOX) compliance controls before anything goes into production, so we work through that in partnership.

Software that has been in production for 10 years eventually becomes unmanageable. Change is hard, but we are pushing in this direction. Our managed service helps: customers who embrace it benefit from upgrades we carry out for them.

You did mention HuggingFace, do you foresee the wider adoption of AI potentially having a threat to security, or the speed at which these things can be dealt with?

It is no longer theoretical. AI is here to stay, and our own work shows how much faster you can identify issues and, in turn, how much faster you can close them. With our customers, we are working on redefining security postures so that we can respond faster. None of us wants to be exposed to a zero-day vulnerability, whether for reputational damage or for the risk of losing customer data.

As I like to say, AI needs to fight AI. Cybersecurity has always been an arms race, and this is another tool in the box.

What else do you believe to be the biggest blockers for legacy systems when it comes to modernisation?

The biggest issue sits completely outside technology. Even if we can build the upgrade and roll it out, it is still change management. We have a very large customer base, and every new version, even a small change like a button moving from top left to bottom right, is a change for someone.

Operations teams can be 1,000 to 1,500 people. Reconciliation is a cost driver rather than a revenue driver, so every second saved matters, and if people have to hunt for a button, that matters. That is the simplest example; a full new user experience means onboarding and retraining people. The cost of training is one of the biggest blockers in this journey.

Would you say then it is a short-term cost for long-term gains into the future?

Not necessarily. It depends on how you modernise. If you only update every 10 years, which is not uncommon, the delta is massive, and so are the training and change-management costs.

I want to move our customers to continuous upgrade. The changes we deliver become much more incremental, so the change-management cost falls and is spread out. The two go hand in hand.

Looping back to AI, what do you believe is the biggest gap between what banks are looking for from the technology and what is actually feasible and deployable today?

What banks are looking for is, in effect, digital employees, and it will still be a team sport. Teams will not be made up of agents alone; there will be a human in the loop at a minimum for approval steps. It becomes a virtual team of digital and human colleagues, in which everyone becomes a builder or a manager, orchestrating their digital peers. That is the dream.

How much is reality? Some of our customers are deploying it, but the bottleneck is observability. How are the agents operating? What can they be accountable for, and what not? What decisions do they make, and do they tell us? A chat history is not full transparency. What happened under the hood? Which systems did they call, how many times did they try, and did they try to circumvent something that did not work, as in the Hugging Face example?

Agents are goal-oriented, and they should be. That also means they can create their own plans: if step one fails, they work out steps C, D, E, and F. In a banking environment we can and should allow some of that, otherwise we are back to the old days of writing down every single step. The challenge is how much autonomy to give the agent, in its planning as well as its execution.

We approach it as a ladder. An agent first assists, then recommends, then acts with approval, and only then acts autonomously within defined bounds. Autonomy is granted task by task and withdrawn the moment behaviour drifts. And the question every bank should be able to answer is: what is your kill switch? ■