

Can SmartStream move bank operations from AI pilots to governed execution?



Written by Foo Boon Ping

Sep 30, 2026 | 6 min read

Banks are experimenting with artificial intelligence at very different speeds, but SmartStream's Nicholas Smith said their common requirement is governance. As trading hours lengthen and operational data grows more complex, the firm is using orchestration, agentic AI and tighter data controls to move automation into production without requiring banks to replace legacy systems first.

Miami - Artificial intelligence (AI) is moving into financial operations at the same time as the operating window itself is becoming less forgiving. Nicholas Smith, Chief Business Officer at SmartStream pointed to plans for longer weekday trading hours on Nasdaq as one example of the pressure this will place on settlement cycles, data quality and exception management. Tokenised assets are adding another data and control problem, while banks remain under pressure to contain operating costs.

That combination changes the question around AI in the back office. The issue is no longer simply whether a model can identify a break or suggest an action, but whether a financial institution can allow it to act inside a controlled production environment, explain what it did and connect it to systems that may have been built long before generative AI existed.

Smith said banks are approaching that transition at very different speeds. Some institutions remain cautious and want to see how the market develops, while others want to move aggressively. The common requirement, he said, is to remain within regulatory expectations and maintain transparency over automated decisions.

Governance becomes the common denominator

Smith said the divergence in banks' appetite for AI is already visible in client discussions. At one roundtable, he encountered institutions that were deliberately conservative and others that wanted to be industry leaders. Yet every participant was focused on governance, whether it was buying AI capabilities or developing them internally.

SmartStream has introduced what it calls GRIT, emphasising governance and transparency over what an AI tool is examining, how it reaches decisions and what outcomes it produces. SmartStream's product announcement ahead of Sibos defines the framework more fully as Governance, Return on investment (ROI), Integration and Trust. Governance covers centralised controls, auditability and human accountability; ROI focuses on autonomous exception management and matching efficiency; integration uses outbound application programming interfaces (APIs) and model context protocol (MCP) connectivity; and trust is built around running multiple reconciliation domains on a common real-time engine. The framework provides a broader product context for the operating issues Smith described, rather than changing the underlying requirement that banks remain accountable for AI used in production.

The governance question becomes more important as automation moves beyond identifying exceptions. SmartStream's September release of Smart Reconciliations Premium applies agentic AI to investigate, progress and resolve reconciliation exceptions, with actions logged and subject to human accountability. Smith described Smart Agents more broadly as a workflow orchestration layer that can work across SmartStream's reconciliation, liquidity, collateral and corporate action solutions while connecting with banks' existing systems.

The distinction matters as banks move from proofs of concept into production. Smith said he is seeing more pilots and proofs of concept than at any point in his career. SmartStream said separately in June that Smart Agents had been tested in tier 1 pilot deployments, with assistive and autonomous modes, human-in-the-loop oversight, full audit trails and step-by-step explainability. Smith's broader point was that banks want suppliers to demonstrate that the technology can do what has been promised, while operating teams need enough exposure to understand what is genuinely possible.

Poor data remains an operational constraint

The usefulness of that automation still depends on the information entering the process. Smith repeatedly returned to data quality as a source of operational failure. He cited research indicating that a large proportion of exceptions originate in poor data and argued that the same problem affects regulatory reporting, where institutions need increasingly complete and timely information.

AI can help repair missing or faulty data more quickly, Smith said, and in some cases identify and correct an issue before it becomes visible downstream. The business value is therefore not only in reducing the labour required to investigate a break. It is also in preventing weak data from interrupting a process later in the transaction lifecycle. SmartStream has previously applied explainable machine learning in its Affinity reconciliation technology, where suggested matches are accompanied by the data attributes behind the prediction, providing an existing example of the transparency Smith said banks now require from AI.

This becomes more important as market structures change. Smith pointed to longer trading hours as one source of pressure on settlement processes and exception management. He also expects tokenisation to create new demands around reference data as institutions need to manage information on a wider range of digital assets.

Smith cautioned that tokenisation can expand the range of assets represented digitally without removing the need to understand their risk. He argued that inadequate reference data for tokenised assets could itself create a new risk profile for banks. The implication is that more automation and more digital assets increase, rather than reduce, the importance of governed data.

Orchestration offers a route around legacy replacement

For many banks, however, the immediate obstacle is architectural. Smith said some institutions still operate highly complex legacy technology estates that cannot be replaced quickly because of the time, cost and implementation risk involved. He therefore expects orchestration layers to play an important transitional role.

SmartStream is supporting MCP connectivity so that its software can connect with banks' own large language models, AI tools and upstream systems. Smith said this allows institutions to use existing technology while adding an orchestration layer rather than waiting for a wholesale platform replacement. That approach is consistent with SmartStream's established Transaction Lifecycle Management architecture, which the company says is designed to integrate with existing processing capabilities rather than force wholesale replacement of legacy systems.

Smart Agents is part of that orchestration approach. Smith said it connects through MCP connectivity to banks' upstream systems as well as SmartStream solutions, allowing institutions to use AI across operational workflows without first replacing the underlying architecture.

Production AI will depend on controlled execution

Smith's account suggests that the near-term path to operational AI is likely to be incremental rather than a clean replacement of bank infrastructure. Institutions can add intelligence and orchestration above existing systems, but the technology still has to operate within controls that make its decisions visible and accountable.

That also changes what banks need to prove in their pilots. A model that performs well in isolation is not enough if it cannot connect to production data, operate within permissions, leave an audit trail and hand decisions back to people when judgement is required. Smith said the speed of development over the past year has made experimentation necessary, but banks are also using those experiments to educate operating teams and establish what they are prepared to automate.

As markets move towards longer operating hours and a broader mix of digital assets, the pressure to reduce manual exceptions will increase. SmartStream's proposition is that banks do not need to wait for their legacy estates to disappear before acting. The more immediate question is how much execution they are prepared to entrust to AI once governance, data quality and integration are strong enough to support it.