Could agentic AI reshape post-trade operations?

Robin Hasson, head of reconciliation solutions at Smartstream, speaks to Zarah Choudhary about the role of ‘always on’ AI agents, the importance of human oversight, and why governance must remain central as automation evolves

The pressure created by shortened settlement cycles is pushing firms to rethink how operational processes are managed across the post-trade environment.

For Robin Hasson, head of reconciliation solutions at Smartstream, the move towards T+1 settlement is less about an increase in workload and more about the compression of time.

“Data integrity is all about completeness and accuracy,” says Hasson. “That’s not a new challenge. The challenge connected to T+1 is time. You’ve got to do the same amount of work, get the same level of quality, within the same timeframe.”

The value of agentic AI is not simply that it allows firms to process tasks more quickly, but that it enables problems to be identified and resolved before failures occur.

“Often in processing today, you find a problem in the data because a failure occurs,” he explains. “What we can start looking at with agentic solutions is much more analysis upfront, much more investigative and correction work upfront, which means you avoid the failure.”

He states that this shifts exception management away from reactive firefighting and towards preventative operations, allowing firms to address issues earlier than would traditionally have been economically viable through additional headcount or systems investment.

“The agentic improvement there is really to guarantee that quality, reduce the breaks, as well as achieve that end-of-day quality and integrity,” he notes.

As European markets continue preparations for a future T+1 transition following implementation in the US, firms are increasingly focused on compressed reconciliation and exception management windows. Agentic AI changes the operational equation because the technology is not constrained by traditional working patterns.

“If something goes wrong, you’re always waiting for a person to step in, investigate and solve,” he explains. “The key difference is that you’ve got always-on agents. Agents are always waiting and available to process and start to do that investigation.”

Rather than allowing investigation queues to build up while staff become available, Hasson says agentic workflows can instantly begin analysing breaks, gathering information, and recommending actions.

“The always-on agents streamline the evaluation, the analysis and the recommendation — potentially all the way through to correction, if you trust it,” he says.

However, he stresses that firms are still maintaining a “human in the loop” approach, particularly in regulated environments.

“You’re just using the human in the loop to validate, approve, and confirm that the action it’s taking is correct,” he explains. “At some point, I would expect the amount of human in the loop to reduce, but certainly for the moment it will remain high at the right points.”

Hasson also notes that the transition to T+1 itself has not necessarily created more exceptions.

“The move to T+1 didn’t change the number of breaks you have,” he observes. “The amount of work isn’t greater, it’s just the window’s compressed.”

At the same time, firms remain reluctant to significantly increase operational headcount or invest in large-scale new infrastructure projects.

“People have no appetite now to increase headcount,” he points out. “The agentic workflows are really seen as the future — not just for T+1, but changing how the operating model works fundamentally.”

He compares the model to a workforce that does not operate within standard business hours.

“Agents don’t need to sleep versus humans that obviously work business hours,” he says.

Beyond speed, the technology could also alter the role of operations analysts themselves. He describes agentic AI as a system capable of performing much of the investigative legwork traditionally carried out manually — including gathering information, analysing patterns, identifying potential causes, and presenting recommendations in a concise format.

“The greatest value is when you harness the AI capability with an expert — somebody who knows the situation and knows how to make decisions,” he says.

Rather than manually investigating each break individually, analysts increasingly become supervisors overseeing AI-generated recommendations and grouped exception patterns.

“You would no longer need to be doing the slow research or waiting for things to come back in,” he explains. “You’re left with the clean data, fully audited, with a recommendation.”

He adds that the user still remains responsible for decision-making.

“The user is still making that decision. They’re still deciding what to do. They’re still in control,” he reiterates.

Firms may eventually manage agentic infrastructure similarly to workforce planning today.

“If you think of your agentic tooling as almost a workforce in its own right, you can scale that capacity,” he suggests.

According to Hasson, firms could temporarily expand computing resources and token spend during periods of heightened operational activity, such as month-end spikes.

“It’s almost like an elastic workforce that you can grow and contract as you need the resource,” he explains.

The interview also explored the growing role of continuous learning and root cause analysis within reconciliation and exception management systems.

The industry has historically evolved through stages: identifying a problem, classifying the problem, understanding the cause, and then determining the corrective action. He believes agentic AI is particularly effective in environments where firms are handling increasingly complex and varied data sets that cannot easily be governed through static rule-based systems alone.

“The agentic systems are really good at spotting those patterns and then being able to work with you on a structure that delivers an action that you oversee and approve when needed”.

He points to use cases where AI-driven systems could automatically update static data, trigger corrections across systems, or notify counterparties of incorrect information with limited operational risk.

At the same time, he repeatedly emphasises the risks associated with fully autonomous execution.

“The risk of full automation is that you can’t always just trust AI. Hallucinations are absolutely a thing,” he notes.

During the discussion, he referenced reports of AI-assisted development tools accidentally deleting live environments and customer data, warning that regulated financial institutions cannot afford similar governance failures.

“In a regulated space, controls have to be thought through”.

He also warns that developers using AI-generated code without sufficient oversight could unintentionally introduce instability into production systems.

“If you don’t check everything every time really carefully, you can have systems breaking everywhere”.

For Hasson, governance, auditability, and data lineage ultimately remain central to the future adoption of agentic AI in financial services operations.

“It cannot be an afterthought,” he says.

Firms must maintain full traceability over decision-making processes and ensure actions taken by AI systems remain visible, auditable and accountable.

“Agents need to be accountable for their actions,” Hasson concludes. “If I’m the leader of that workforce, I need to know that it’s going to do the right thing.” ■

“The risk of full automation is that you can’t always just trust AI. Hallucinations are absolutely a thing”

Robin Hasson
Head of reconciliation solutions
Smartstream