Post-trade is the place for agentic AI workflows

By Elizabeth Lumley, editorial contributor, FinTech Futures

In the next few years, most post-trade operations analysts won't be bogged down with processing exceptions. Instead, they will supervise agents, configure workflows, manage counterparty relationships at a strategic level and handle cases that require human authority, according to Yogesh Shenai, director of product management at Smartstream.

Supporting the smooth running of the ongoing global expansion of T+1 settlement in the post-trade space is precisely the type of 'problem' agentic AI has been waiting to solve, he says.

T+1 settlement, in which securities transactions must be finalised within one business day of the trade date, is intended to reduce market risk and free up capital. North American markets transitioned in May 2024, while the UK, EU and Switzerland will implement it on 11 October 2027.

However, Shenai says the firms that have transitioned found their operations hadn't changed much: "they'd just become faster at doing the same manual things in a shorter window," he says. The problems exposed weren't data or infrastructure problems; they were workflow issues, he adds.

"The right information exists, the right processes are known, but executing them fast enough and consistently enough across hundreds of simultaneous exceptions is beyond what manual operations can deliver," he adds.

Shenai explains that, post-T+1 in North America, many buy-side firms missed trade-date affirmation deadlines because monitoring and chasing were still manual. Meanwhile, break investigation – identifying why a trade failed to match, enriching missing data fields before the trade escalated, and correcting it in time – proved "too slow" when done by humans across multiple systems simultaneously, he adds.

Gaps in validation

Exception queues weren't prioritised by settlement urgency, so the most time-critical breaks weren't always the ones getting attention first, Shenai goes on to say. However, when settlement failed, the downstream workflow of counterparty notification, partial settlement decisions and penalty tracking remained largely manual, he says. Meanwhile, standard settlement instruction (SSI) applications unveiled gaps in validation

and confirmation timing under the compressed window, he adds.

The UK, the EU and Switzerland have the advantage of watching North America's transition as they prepare for 2027. However, those transitions could be supported by adopting agentic AI tools, says Shenai.

When a human picks up an exception today, they're "starting cold", he says. That means they are logging into multiple systems, reading email history and reconstructing what happened. In an agentic model, the escalation itself is redesigned, he adds, as agents enact "pre-emptive triage".

An agentic workflow will monitor live trade data against known resolution patterns in real time – identifying signals that historically precede a break, says Shenai. For example, these can include a counterparty that consistently responds late, an asset class with chronic SSI application delays, or a trade type where confirmation mismatches cluster. When it detects a trade that fits that profile, it triggers an investigation that immediately queries systems, validates SSI application and cross-references records.

"That is exactly what an agent does," says Shenai. "It doesn't flag a problem and wait; it investigates, communicates, enriches, prioritises and resolves – autonomously, across systems, at a speed that the T+1 window actually allows."

This means that post-trade analysts receive a pre-built case, allowing them to make a decision rather than conducting a lengthy investigation. "That matters enormously when the clock is running," he adds.

Reducing manual touch rates

An agentic workflow can also change the way email communication is handled and automated, says Shenai. Email remains one of the primary communication channels between operations teams and counterparties. However, an AI agent can draft and send outreach to counterparties, monitor for responses, read the reply in natural language, assess its sentiment and intent, and bring that determination back into the workflow, making the email thread a structured data input, he adds.

According to Shenai, automating routine exception handling reduces manual touch rates by 60-80% across in-scope workflows. Meanwhile, an investigation that takes a skilled analyst 30 to 45 minutes can be completed by an agent in less than two minutes, he adds.

AI agents can act autonomously, with no human interaction, for high-volume, well-patterned exceptions where the workflow is proven and the action is bounded or augmented, where the agent holds and waits for explicit human sign-off before executing.

However, activities that require human input should be "deliberately designed as human-only from the outset", says Shenai. "Not a fallback when the agent gets stuck, but a conscious workflow design choice." Those activities can include a relationship call, a counterparty escalation carrying reputational weight, or a judgement call with legal implications.

However, there are cases where an agentic AI solution could face barriers.

> "Agents don't have shift changes. They don't have time zones. They can investigate, communicate and resolve exceptions at trade speed across any market, 24 hours a day."
>
> Yogesh Shenai, Smartstream

Those include instances of counterparty behaviour with no historical footprint; disputes requiring contractual or legal interpretation; complex multi-leg trades or failures caused by third-party infrastructure – such as a custodian system outage or a regulatory hold. In these cases, the agent recognises the barrier and escalates. "Knowing when to stop is as important as knowing how to act," Shenai says.

Much of that knowledge lies with experienced operations analysts, adds Shenai.

Today, "the most valuable thing" an experienced analyst carries is their judgement about which process to follow, which counterparties need handling carefully, which exceptions are routine, and which are signals of something larger, he says.

By building agentic workflows, firms can embed that knowledge into the workflow design itself, says Shenai. That experience can now be encoded, inspected and transferred, allowing new team members to inherit institutional knowledge from day one and letting that knowledge improve over time as the workflow is refined.

While the world is moving towards T+1 settlement, the development of agentic AI workflows could lay the foundation for a future T+0 environment, says Shenai.

"Agents don't have shift changes. They don't have time zones. They can investigate, communicate and resolve exceptions at trade speed across any market, 24 hours a day," he says. "The post-trade operations layer is actually one of the areas most ready for T+0."