



Thomas Steinborn (Background photo: Krizjohn Rosales / Pexels)

Kill the exception queue: post trade's real bottleneck is the 2%

© SMARTSTREAM 📅 2026-09-28

Post-trade operations have become a story about the two percent.

Contributed by Thomas Steinborn, chief product and technology officer, SmartStream

 **smartstream** Around 98% of transactions now clear straight through – matched, confirmed, and closed within minutes, untouched by human hands. The remaining 2% consume roughly 70% of operational effort. Those are the breaks: the mismatches, the failed instructions, the unexplained cash differences that take days rather than minutes to resolve. That is the uncomfortable arithmetic of modern operations. We have automated matching brilliantly. We have barely automated exception management at all.

For years, that imbalance was tolerable, because the queue was survivable. It is becoming less so. T+1 has compressed the window for investigation to a fraction of what it was. Volumes are climbing 15 to 20%, and volatility multiplies mismatches faster than analysts can clear them. ISO 20022 brings richer, harder data; DORA raises the bar on resilience. Meanwhile, hiring and training a reconciliations analyst takes months. A volatility event takes hours. Capacity cannot be summoned at the speed the market now breaks.

Look closely at how post trade actually fails, and the same fault line appears each time. Instructions fail in flight while markets keep moving, so funding gaps open before the break is detected. A break raised in Asia waits for EMEA to open, and the cut-off passes while the queue sleeps. Backlogs roll into the next cycle. Different pressures, one common failure: exceptions that need a decision faster than people can make one.

The execution gap

This should be the moment AI proves its worth. On the evidence so far, it hasn't – not because the models are weak, but because of where they have been pointed.

The 2026 picture is sobering. Some 88% of organisations use AI somewhere, yet only around 5.5% report material EBIT impact. MIT's NANDA research found 95% of enterprise generative AI pilots delivered no measurable P&L effect. Spend has tripled to US\$37 billion; roughly a quarter of initiatives met expected returns.

The gap is not intelligence. It is execution. Copilots advise, dashboards report, models recommend – nothing acts. Each of those tools still terminates in a human who must open the case, query three systems, interpret a standard operating procedure (SOP), decide, execute, and close. Add a copilot and you make each step marginally faster. You do not remove the step. The queue survives, and so does the cost.

That is why exception management is the natural first home for genuinely agentic AI. Exception handling is not creative work. It is procedural, documented, rules-bounded, and evidence-generating – precisely the profile of work an agent can be trusted to complete rather than merely comment on. The pattern is simple to state and hard to achieve: plan, decide, act, escalate. Read the exception, its context, and the applicable SOP. Choose a resolution path within policy. Execute the workflow and the counterparty communication. Hand to a human at defined risk thresholds. Detection to closure, with no queue in between.

Autonomy is only worth what its governance is worth

Here operations leaders are right to push back, and the data supports them. Only 23% of organisations are scaling an agentic system, even as most experiment with them. Gartner expects 40% of enterprises to demote or decommission autonomous agents by 2027 over governance gaps found after incidents. Since August 2026, under the EU AI Act, high-risk obligations – logging, traceability, human oversight – have been legal requirements rather than aspirations.

The lesson is not to slow down. It is that trust must be engineered, not assumed. In practice, that means three things. Guardrails: agents act only inside defined workflows, thresholds, and approval limits that operations own and can change, with least-privilege credentials and allow-listed actions. Evidence: an immutable audit trail on every step – inputs, reasoning, action, outcome – explainable, replayable, and regulator-ready by default. And graduated autonomy: a ladder from drafting, to recommending, to acting with approval, to acting within bounds, climbed only as accuracy proves out and revoked the moment behaviour drifts. Accountability stays with a named human owner; it is never delegated to the agent.

Firms taking this route report 50 to 70% reductions in exception-handling effort, resolution times an order of magnitude faster, and complete audit coverage – the last being the reason the first two survive contact with risk and compliance.

The exception queue was never a technology constraint. It was a human-throughput constraint we industrialised around. That constraint is now the binding one. The firms that remove it will not be the ones with the best-informed analysts. They will be the ones whose operations no longer depend on a queue at all.

To find out more, visit [Smartstream](#).