When exceptions become the rule

As firms face growing data complexity, compressed settlement cycles, and increasingly interconnected markets, operational breaks are becoming harder to prevent and faster to resolve. Zarah Choudhary explores why reconciliation remains at the heart of operational resilience

For years, the industry has viewed automation as the answer to post-trade inefficiencies. Yet despite significant investment in straight-through processing (STP), reconciliation, and exception management remain among the most operationally intensive functions within capital markets.

The challenge is partly one of scale. As firms expand into new asset classes, onboard additional counterparties, and process greater transaction volumes, operational complexity grows alongside them.

Murray Campbell, principal product manager at AutoRek, says the impact is often felt most acutely within operations teams.

"For firms scaling or expanding their business the impact is felt most notably across operations, including reconciliations. Increasing transactional volumes, a growing number of counterparties, custodians, and third-party banks, plus expansion across new asset classes creates increasing challenges for reconciliation processes to handle."

Val Wotton, managing director and global head of equities solutions at DTCC, says the challenge is being amplified by wider market developments.

As firms contend with shortened settlement cycles, extended trading hours, and increasingly fragmented market activity, reconciliation must now occur across more counterparties, varying levels of automation, different market practices, and multiple time zones, leaving less time to resolve exceptions.

Steve Walsh, managing director of reconciliation at DUCO, points to the rise of private markets and digital assets as further contributors to operational complexity. Private markets often involve highly fragmented and unstructured data, while tokenised assets introduce new infrastructure and processing models that differ significantly from traditional markets.

Neil Vernon, chief product officer at Gresham, argues that reconciliation challenges are becoming more complex across virtually every dimension of the post-trade lifecycle.

Robin Hasson, head of reconciliation solutions at SmartStream, says the industry is facing several overlapping pressures simultaneously.

"Transaction volumes are growing at more than 25 per cent year-on-year, settlement timelines are shrinking, and regulatory

gorodenkoff/stock.adobe.com

expectations keep intensifying. Incoming regulatory initiatives are pushing firms to run more reconciliations across increasingly complex data, often on infrastructure that was never designed for this level of scale or speed."

He points to the European move to T+1 settlement, upcoming SWIFT standards changes, and the growing adoption of tokenised assets as examples of operational change occurring concurrently rather than sequentially.

"Transaction volumes are up, asset class complexity is up, regulatory requirements are up, and the counterparty ecosystem continues to expand in ways that defy easy automation."

The data challenge sits at the centre of the problem. Firms must reconcile positions, transactions, cash, collateral, corporate actions, and securities lending activity across multiple internal systems and external parties simultaneously.

As Vernon notes, the complexity of a single OTC derivative lifecycle can involve hundreds of data points changing throughout its lifespan.

Why breaks still happen

At its most basic level, a reconciliation break occurs when two records fail to match.

According to Campbell, timing differences between records, processing errors, and trade settlement issues remain among the most common causes.

"Reconciliation breaks are the result of a data mismatch between the two records in question. This is commonly the result of timing differences between when both records are updated, or as a result of inaccuracies on either record."

Settlement failures themselves are hardly a new challenge. The European Central Bank (ECB) has long identified operational risk, unmatched instructions, processing errors, and securities shortages among the primary causes of settlement failures. It notes that even relatively low fail rates can create cascading chains of settlement issues that affect wider market activity.

Karan Kapoor, head of risk and regulatory consulting at Delta Capita, argues that most operational breaks ultimately stem from wider data governance issues.

Data ownership often remains fragmented across multiple teams, systems, and workflows, creating inconsistencies throughout the post-trade lifecycle. Mismatched client static data, inconsistent payment information, and discrepancies between internal and external records frequently surface as reconciliation breaks, even though the underlying problem originates elsewhere.

Others point to standing settlement instructions (SSIs) as a persistent source of friction.

Vernon describes stale or incorrect SSIs as a "perennial source of breaks", while noting that the industry has been aware of the issue for decades.

Wotton similarly highlights inaccurate settlement instructions, reference data inconsistencies, and timing misalignments between counterparties as common causes of operational mismatches.

Hasson argues that data quality remains the single biggest source of operational breaks.

"Up to 67 per cent of asset servicing errors are directly attributable to data quality issues. That's not a problem with matching logic; it's happening upstream."

He identifies inconsistent counterparty formats, stale standing settlement instructions, and instrument identifier mismatches as some of the most common causes of exceptions, adding that many firms continue to suffer from a lack of front-to-back consistency across their operating models.

The issue has become increasingly significant as firms prepare for T+1 settlement in Europe. The EU T+1 Industry Committee has identified SSI standardisation as one of the key areas requiring additional market focus and guidance ahead of implementation.

Liz Collegio, head of product management at Meritsoft, highlights counterparty behaviour and fragmented communication channels as additional contributors.

"SSI submissions don't arrive uniformly, and teams without visibility into counterparty patterns waste time chasing trades that will self-resolve while genuinely at-risk trades go unattended."

She adds that reliance on email for settlement instructions and exception management continues to introduce delays and operational errors.

Automation has helped — but it has not solved the problem

The industry has undeniably benefited from automation. Match rates have improved, manual processing has declined, and firms can handle greater volumes than ever before.

Yet automation has not eliminated complexity.

Hasson believes automation and complexity have evolved in parallel rather than one replacing the other.

"Automation has meaningfully improved straight-through processing rates and without it the volume growth we've seen over the past decade would have required unsustainable headcount increases."

However, he argues that many firms have automated individual processes without achieving the end-to-end consistency needed to eliminate breaks entirely, leaving operational teams responsible for managing increasingly complex exceptions.

Campbell observes that firms often assume that automated systems can absorb increasing complexity indefinitely, creating new operational pressures when systems reach their limits.

In some cases, firms have "automated themselves into a problem", investing in solutions that addressed a specific challenge at a particular point in time but lack the flexibility to support new asset classes, growing volumes or evolving business requirements.

Kapoor argues that automation delivers the greatest benefit when supported by strong data quality, governance and interoperability across the wider ecosystem. However, the industry remains fragmented, with firms automating at different speeds and operating different post-trade models, creating new dependencies between counterparties and infrastructure providers.

Wotton makes a similar observation, noting that many firms now operate hybrid environments where newer automated processes coexist with legacy systems, creating additional operational dependencies even as manual intervention declines.

Vernon makes a similar point.

"Automation has largely succeeded in shifting complexity rather than eliminating it. The residual exception tail — the breaks that require investigation, human judgement, and counterparty communication — remains stubbornly resistant."

This phenomenon has become a defining feature of modern post-trade operations. Automation handles the straightforward majority of transactions, leaving operations teams focused on the small percentage of exceptions that tend to carry the greatest operational and regulatory risk.

Automation still requires oversight

Despite growing automation levels, firms remain cautious about removing human oversight entirely.

Vikash Rughani, head of solution design for trade processing and lifecycle at OSTTRA, notes that most market activity is already highly automated, with controls focused on monitoring exposures and managing exceptions through defined tolerances and business logic.

The challenge for firms is determining where human involvement continues to add value and where greater autonomy can safely be introduced.

Several interviewees emphasised that auditability remains critical. Wotton points to enhanced monitoring, escalation frameworks, and governance controls as essential safeguards, while Walsh argues that well-governed automation can improve operational oversight by creating stronger audit trails and reducing reliance on manual data handling.

As automation expands, firms are increasingly adopting exception-led workflows, where operational teams focus on investigating prioritised breaks rather than reviewing every transaction individually.

The continued burden of legacy infrastructure

Legacy systems remain a major obstacle to operational efficiency.

Campbell argues that older platforms are often difficult to integrate with modern workflows and frequently require spreadsheet-based workarounds to complete processes.

Kapoor notes that many financial institutions continue to operate fragmented technology estates built over decades, often involving multiple platforms, duplicated data sources, and manual processes. These environments make it difficult to achieve the real-time visibility increasingly required under compressed settlement cycles.

Vernon is equally critical, describing many legacy reconciliation platforms as solutions designed for "yesterday's volumes and message formats".

These systems often struggle to support intraday reconciliation, modern asset classes, or multiple data sources simultaneously. They also risk undermining the benefits of newer industry standards.

ISO 20022, for example, was designed to improve data consistency through richer, structured, and machine-readable messaging. Campbell notes that firms using the standard effectively are seeing faster matching and fewer manual breaks, although inconsistent implementations continue to limit the benefits.

Vernon agrees that ISO 20022 has materially improved reconciliation where implemented properly, but warns that many firms still truncate or remap the richer data into legacy formats, reducing its effectiveness.

Several interviewees echoed this view.

Wotton says the full benefits of ISO 20022 will only be realised once adoption becomes more consistent across market participants, while Kapoor argues that standardised messaging alone cannot solve fragmented internal data architectures and inconsistent operational processes.

AI moves from theory to practice

As firms continue searching for efficiencies, AI is increasingly being viewed as the next frontier for reconciliation and exception management.

Rather than replacing reconciliation engines, many firms are deploying AI to enhance existing workflows.

Campbell identifies several practical applications, including ingesting and normalising unstructured data, identifying potential record matches, and assisting users with exception investigations.

Meritsoft is seeing adoption in exception management, where AI can extract settlement requests and fail notices from emails, enrich the information, and route it directly into post-trade systems.

Meanwhile, Gresham highlights intelligent matching, automated rule generation, exception classification, and predictive break detection as some of the most practical production use cases emerging today.

SmartStream is seeing similar momentum in exception-heavy workflows.

According to Hasson, AI is proving particularly effective in reconciliations, cash breaks, and settlement exception management, where autonomous agents can investigate discrepancies, gather supporting data, and route cases for resolution.

"We're seeing up to a 29x reduction in time per break, 30 to 60 per cent faster resolution times and up to 70 per cent reduction in manual investigation workload, without adding headcount."

He also argues that AI is increasingly being applied to reconciliation configuration and optimisation, helping firms automate data mapping, matching rule creation, and ongoing tuning of reconciliation processes.

Kapoor says many firms are focusing on predictive analytics and intelligent exception handling, using AI to identify potential breaks before they occur, prioritise exceptions based on historical patterns and recommend remediation actions.

Rughani believes AI's greatest value lies in helping users understand the root causes of disputes more quickly, enabling operational teams to focus on genuinely complex investigations rather than routine queries.

Marketnode's Vihang Patel highlights document intelligence, exception triage, golden-record adjudication, and conversational support for operations teams among the most practical use cases emerging today. However, he stresses that auditability remains non-negotiable in regulated environments, with AI outputs requiring the same standards of traceability and control as any other operational process.

The wider industry is increasingly viewing AI as an augmentation layer rather than a replacement for existing reconciliation systems.

Research suggests that AI can improve exception classification, root-cause analysis, and data normalisation while leaving deterministic matching and controls intact.

T+1 raises the stakes

If there is one initiative forcing firms to address operational weaknesses, it is T+1.

The move to shorter settlement cycles significantly reduces the time available to identify and resolve discrepancies.

"The reduction in operational time available within which to investigate and resolve discrepancies" is the biggest impact of shorter settlement cycles, according to Campbell. Manual processes, he argues, are simply unable to cope.

Kapoor argues that T+1 has not necessarily introduced new operational challenges but has effectively eliminated the industry's tolerance for delay.

"What was once manageable within a longer operational window now requires near real-time visibility and remediation," he says.

Wotton similarly notes that T+1 requires earlier trade affirmation and much tighter coordination between front, middle and back office functions.

Hasson argues that shorter settlement cycles expose the industry's historical reliance on end-of-day processing.

"The shift exposes something firms have been able to work around under T+2: a deep dependency on end-of-day batch reconciliation."

He believes firms will need to move towards more parallel and near real-time processing models, particularly as Europe prepares for T+1 implementation across multiple markets, currencies and central securities depositories.

Collegio describes T+1 as the "forcing function" that makes straight-through processing non-negotiable.

Walsh goes further, describing compressed settlement cycles as the single biggest operational change facing the industry. Tasks that were previously completed over a 24-hour period increasingly need to be addressed on trade date, leaving little room for poor data quality, inaccurate settlement instructions, or false reconciliation breaks.

The EU T+1 Industry Committee has similarly emphasised timely matching, automation, and real-time processing as core requirements for successful implementation.

For many firms, the challenge extends beyond T+1. Looking further ahead, T+0 settlement and 24/7 trading environments could require continuous reconciliation models capable of operating in real time.

As Vernon notes, batch processing architectures are fundamentally incompatible with a T+0 world.

The next phase of reconciliation

For all the industry's progress in automation, reconciliation remains fundamentally a data problem. Growing transaction volumes, expanding market ecosystems, and shorter settlement cycles are exposing weaknesses in operational models that were designed for a different era.

Several interviewees argue that the next phase of transformation will focus less on processing transactions faster and more on preventing breaks from occurring in the first place.

Delta Capita's Kapoor believes reconciliation is becoming increasingly proactive and predictive, supported by intelligent workflow orchestration and stronger data foundations.

DUCO's Walsh expects firms to focus on identifying and eliminating exceptions further upstream in the trade lifecycle, while Gresham's Vernon argues that long-term success will ultimately depend on solving underlying issues around data normalisation, completeness, and consistency.

The next phase of transformation is unlikely to come from eliminating human oversight altogether. Instead, firms are increasingly focusing on combining automation, AI, and better data quality to reduce the volume of exceptions reaching operations teams in the first place.

As settlement cycles continue to compress and operational expectations increase, the firms best positioned for the future may not be those that automate the most, but those that manage exceptions the best. ■