



From fragmentation to integration: Rethinking the securities finance lifecycle

Tony Camarota, head of Spire at EquiLend, discusses eliminating the divide between the front and back office, the significance of a unified model, and the role Spire 3.0 plays in tackling fragmentation

Securities finance has long operated with a structural divide between front office trading and post-trade operations. Trading desks make pricing, allocation, and execution decisions in real time, while post-trade teams manage books and records, lifecycle events, and operational control.

Each side has historically run on separate platforms built for its specific role, and that structural divide is intentional — back office controls ensure the front office is not taking on undue risk, while front office decisioning protects risk-adjusted returns

While this separation has worked historically, newer market dynamics like decreased settlement time, are increasingly creating friction with this model.

The biggest disconnect today is information continuity.

Those pressures are not easing. T+1 is the baseline now with T+0 in view, amalgamated activity across securities borrowing and lending (SBL), repo, and swaps each demands distinct workflows, and trades increasingly touch a mix of internal, external, and EquiLend systems.

Every additional touchpoint is another place for breaks if the data is not continuous

Front office desks operate in real time, while post-trade teams often work off delayed or transformed versions of that same data. In practice, that shows up in familiar ways: trades are re-interpreted downstream instead of flowing cleanly from execution to lifecycle. Manual reconciliations emerge between what the desk believes it traded and what books and records reflect. Inventory, availability, and restrictions are checked too late in the process, rather than shaping decisions upfront. The lack of a single source of truth for positions, static data, and lifecycle events leads to inefficiency and introduces risk via unreconciled breaks or incorrect or stale data.

For market participants, this creates a constant cycle of explaining and correcting breaks that originated much earlier in the workflow. It leads to lost opportunities, increased operational effort, and more generally limits the ability to scale.

Just as importantly, it creates a disconnect between decision-making and accountability. The trading desk makes the decision, but operations inherit the consequences, often without full visibility into how or why the trade was structured in the first place. That misalignment drives inefficiency across the entire lifecycle. It also exposes the front office directly: trading against bad inventory means foregone revenue, and trading into fails strains counterparty and client relationships.

When trading and lifecycle environments are run separately, those issues compound. Operational workflows depend on handoffs, file-based transfers, and manual touchpoints, particularly around allocations, returns, and lifecycle events. The same trade can exist in multiple states across systems, creating reconciliation work, audit complexity, and reporting risk. At the same time, risk limits, inventory constraints, and client rules are often enforced after the fact instead of shaping the decision itself. Even the most experienced staff cannot achieve 100 per cent accuracy when processing these various touchpoints manually.

Over time, these inefficiencies scale faster than headcount. Many desks reach a point where they are operationally maxed out even as volumes grow. At that point, growth is no longer constrained by market opportunity, but by the ability of the operating model to support it.

That is why we built Spire 3.0, with seamless integration to the EquiLend Front-Office Automation engine.

Eliminating fragmentation at the source

Spire 3.0 represents a modern architectural approach designed to remove the disconnect between execution and lifecycle by operating as a real-time, continuously updated record of positions while simultaneously enabling front office decisioning.

With real-time connectivity between the EquiLend Front-Office Automation engine and lifecycle systems, every pricing, allocation, and execution decision is captured immediately and carried forward without transformation. There are no batch files, no reinterpretation, and no competing versions of the truth.

Trade data is written once and persists through the lifecycle. Inventory, availability, and restrictions are visible at the moment decisions are made, not after the fact. This eliminates most downstream breaks, fails, and more generally, the operational effort required to resolve them.

Equally important, it changes when and how value is created.

Real-time pricing, allocation, and strategy execution are no longer just inputs into lifecycle systems, they are the primary drivers of performance. When constraints, inventory, and client rules are known upfront, traders can optimise decisions in real time, improving revenue quality rather than simply increasing throughput.

Control is no longer applied after execution. It is embedded directly into the decision process.

This model addresses three structural gaps created by fragmented environments.

Operational friction is reduced because lifecycle events are captured once and reflected everywhere in real time. Data inconsistency disappears because execution and lifecycle operate on a shared data model. Control gaps close because limits and rules are enforced at the point of decision rather than identified later through exceptions.

The result is continuous data lineage from trade initiation through settlement, allowing firms to operate on a single, consistent view of positions while simultaneously optimising decisions as they are made.

Bringing front and back together in practice

The most effective operating models do not connect separate systems. They unify decisioning and lifecycle control into two tightly coupled layers operating on the same live data. Spire 3.0 and the EquiLend Front-Office Automation engine operate within this shared model.

One layer focuses on front office automation and decisioning, including pricing, allocation, and execution strategy. The other manages lifecycle processing, books and records, and control. Both operate on a shared, real-time data and control model.

Trades flow directly from execution into lifecycle processing without re-keying or reinterpretation. Inventory, eligibility, limits, and restrictions shape decisions in real time. Lifecycle events such as returns, corporate actions, and rerates are managed against a consistent position view.

Traders adjust strategy because constraints and inventory are visible upfront, not discovered later. Allocation decisions are made with full awareness of availability and client requirements. Automation does not just increase efficiency, it improves the quality and consistency of revenue generation.

At the same time, lifecycle updates continuously inform decisioning. Position changes, inventory movements, and constraints are reflected immediately, allowing the front office to adapt without delay.

Firms move from managing handoffs between systems to managing outcomes across a unified model.

Access to broader capabilities such as connectivity, market infrastructure, and analytics services can be embedded within the same environment, reducing context switching and further improving decision quality.

From operational constraint to scalable growth

The most meaningful outcomes of this model are scale, confidence, and improved decision quality.

Straight-through processing (STP) increases as trades move through the lifecycle without interruption. Operational risk decreases as exception handling is minimised and lifecycle processes are automated within a unified model.

At the same time, decision-making improves.

Firms can act on market opportunities with full visibility into inventory, constraints, and client requirements. Decisions are faster, but more importantly, they are more accurate and repeatable.

Confidence improves because trading, operations, and risk functions are all working from the same data, with the same understanding of positions and exposure.

Firms can scale volume without proportional increases in headcount because the operating model removes the need for manual intervention and reconciliation.

The shift is not just from manual to automated processes, it is from reactive workflows to proactive, data-driven decisioning.

A model for control and visibility

A unified model provides continuous visibility into inventory across strategies, clients, and constraints at the moment decisions are made.

Allocation logic and programme rules are applied consistently from execution through settlement.

Control is embedded directly into both decisioning and lifecycle processes.

Constraints are defined once and applied consistently, ensuring that execution aligns with operational, risk, and client requirements from the outset. Governance is no longer a downstream activity, it is part of how decisions are made.

This has a direct impact on decision quality.

Firms are able to reuse decisions more effectively, execute strategies more consistently, and deliver more predictable client outcomes.

Instead of correcting exceptions, they prevent them.

For lenders, this translates into better inventory utilisation, stronger control over programme rules, and more consistent revenue capture without increasing risk.

For borrowers, this means improved access to inventory, more consistent pricing, and greater confidence that availability, constraints, and counterparty requirements are reflected upfront, reducing failed trades and improving execution certainty.

The shift to platform-based operating models

Client expectations have shifted. The question is no longer whether a tool solves a single problem, it is how it fits into an end-to-end operating model.

There is a clear move toward platforms that support both real-time decisioning and lifecycle control within a unified architecture.

Flexibility still matters, but modular adoption must exist within a coherent framework where capabilities reinforce each other.

Disconnected tools no longer meet that requirement.

Spire 3.0 provides a foundation that supports incremental adoption while maintaining a consistent data and control model across execution and lifecycle. Spire's integration with the EquiLend Front-Office Automation engine ensures that front office decisioning and lifecycle control operate on the same real-time data model, reinforcing consistency across execution, inventory management, and downstream processing without introducing additional system complexity.

All of this is already in production.

The case for a unified lifecycle

Securities finance is moving toward a model where information continuity, real-time decisioning, and embedded control are baseline requirements.

Performance is no longer defined by how well individual components operate in isolation. It is defined by how effectively decisioning and lifecycle processes work together.

Spire 3.0, combined with the EquiLend Front-Office Automation engine, reflects how this architecture is being implemented in practice, creating a unified operating model where decision intelligence and lifecycle control evolve together.

The economics follow the architecture. Collapsing front and back office onto one platform delivers measurable hard and soft-dollar savings: firms eliminate the cost and complexity of managing multiple vendors and proprietary applications, and benefit from new features arriving fully integrated rather than requiring custom work to wire in.

The firms that treat execution and lifecycle as a single discipline now will be the ones setting the standard for what comes next. ■