December 2025

Collateral Mobility: A Proposal to Tokenize Encumbrance

This memo introduces an operational model for alleviating the need to over-collateralize—one that is engineered from the first principles of the market's requirement rather than reverse-engineered to fit into a construct defined by protocols used in decentralized finance. The model has not yet been attempted in any proof-of-concept or incorporated into any current T+0 proposed design.

1. Introduction

Collateral mobility is the structural engine of institutional finance. In practice, it is the reallocation of the right to use cash and securities as transaction insurance. This reallocation is essential wherever counterparty and settlement risk must be backed by recoverable value. But the markets that depend on collateral are fragmented by operational silos, and investors compensate by over-collateralizing: they hold cash and securities in stand-by reserves so they can pledge collateral when needed. Necessary as it is, this practice is capital-inefficient.

Over the past decade, the Bank for International Settlements has documented the same pattern of structural over-collateralization across every major jurisdiction. They found that the drivers are consistent: duplicated collateral requirements, fragmented settlement systems, intraday liquidity constraints and time-zone mismatches.[1]

2. Encumbrance makes it work

At the center of collateral mobility is the concept of encumbrance. Encumbrance is the legally enforceable allocation of rights over an asset—who may use it, under what obligation, with what priority, and subject to what release conditions. In legal and accounting terms, an encumbrance is a book-entry liability recorded against an asset when it is pledged.

For cash or securities to be encumbered, the liability must be recorded at the moment the asset is pledged. Before the asset can be reused, the conditions required to release that liability must be met. The event that releases the asset extinguishes the liability on the books. Only after that liability is extinguished can the asset be reused as collateral elsewhere.

Encumbrance is therefore not a transfer of cash or securities. It is a rights-state change: the authoritative record of who holds enforceable claims over an asset at any moment, expressed as a liability on an accounting system.

3. Liability Accounting is the Unlock

In practice, an encumbrance functions as a claim. Claims are complex accounting objects: they prioritize rights over an asset and define the obligations attached to it, including reuse permissions, substitution rights, release conditions, and lifecycle consequences.

To manage these objects, institutional collateral requires a ledger that can express the full accounting picture—who owes what, who holds rights over which assets, in what order, under what conditions, and with what lifecycle effects. Clearinghouses rely on this structure to manage margin obligations and liquidation rights. Depositories rely on it to maintain issuer record and lifecycle truth. Custodians rely on it to maintain an enforceable chain of custody. Settlement banks rely on it to provide finality signals that release or encumber cash collateral. Regulators rely on it to ensure PFMI compliance, legal enforceability, and jurisdiction-specific protections.

Author: Bill McCahey
+1 551-449-0343
bill@ledger4securities

4. Industry Drivers

The push toward T+0 reflects a desire to accelerate the market we already have and improve capital efficiency. And the market we already have runs on credit. Credit underwrites trading, enables netting, supports leverage, and determines how collateral is allocated and reused. It requires the recording of future payables and receivables, while settlement records obligations met and payments made. Any T+0 design must therefore treat credit as the organizing principle of settlement.

Collateral is not a separate innovation; it is a feature of credit, and credit is a feature of accounting. The industry already has a system capable of recording both sides of every obligation with legal finality: double-entry accounting. Bringing that structure directly into the settlement layer is what enables true T+0—not by discarding the logic of T+2 and T+1, but by executing the same bilateral accounting event with immediate finality.

5. A Bilateral Settlement Substrate Proposal

A settlement system capable of delivering true T+0 must be able to record both sides of every obligation in the same atomic state transition. This is the defining requirement of financial markets. On a double-entry accounting substrate, the accounting event becomes the settlement event: credit can be injected directly, collateral becomes an integrated feature rather than an external workaround, and finality emerges from the same bilateral structure that underwrites T+2 and T+1 today—executed with T+0 efficiency.

Implemented on blockchain technology, this substrate operates as a self-contained accounting environment dedicated to tokenization and designed to interoperate with external chains. It differs from other blockchains because it functions as a networked subledger for traditional finance systems.

The mechanism that enables this is the bilateral state transition. The settlement layer records the full accounting event at once: the asset leaving one balance sheet, the liability leaving another, the journal entries generated, and both obligations extinguished in a single authoritative update. In this model, legal settlement finality is achieved on-chain and then synchronized with off-chain systems, rather than the other way around.

This structure also allows credit to be expressed as a liability, making prefunding optional. Netting can occur directly on the accounting layer. Lifecycle events can be applied in sync on-chain and off-chain. Collateral can be allocated or released in real time.

The greatest utility of this architecture is its ability to coordinate the books of multiple institutions. The blockchain runs a distributed network of nodes, each maintaining a copy of the token-dedicated subledger. This Network Accounting Subledger (NAS) does not replace existing systems; it interacts with and synchronizes them. It ensures that each participant sees the same bilateral accounting event at the same moment, with the same legal effect, without introducing a new asset model or operational surface. Because the NAS runs on blockchain technology, it provides shared timing, ordering, and atomicity, enabling institutions to post the same double-entry event simultaneously while preserving their autonomy and regulatory boundaries.

6. Accounting as the Universal Interface

What makes this architecture immediately adoptable is that accounting is already the universal interface for financial-market systems. Every trading platform, risk engine, clearing system, collateral manager, and core banking platform ultimately resolves to a GAAP-compliant accounting endpoint, and they already interact with that endpoint through APIs today. A settlement layer built on double-entry accounting does not require institutions to change their systems, rewrite their workflows, or adopt a new conceptual model. It simply moves the accounting event—the event every system is already designed to produce and consume—into the atomic settlement layer dedicated to tokenized markets. This dramatically lowers the adoption bar.

Author: Bill McCahey
+1 551-449-0343
bill@ledger4securities

Because the settlement layer becomes an accounting layer, every existing financial-market system can connect to it natively. Netting engines, risk engines, trading systems, finance-desk systems, collateral platforms, treasury systems, and reporting systems already speak accounting through APIs. They already produce debits, credits, payables, receivables, and lifecycle events as their authoritative outputs. A settlement substrate built on double-entry accounting simply receives those outputs directly. Nothing needs to be re-engineered, no new asset model needs to be adopted, and no new messaging surface needs to be invented. The systems that run the market today—front office, middle office, back office, and regulatory—plug in immediately because accounting is the language they already use to 'run the bank'.

Taken together, these elements point to a simple but consequential shift: the path to T+0 does not require a new market structure, a new asset model, or a new category of financial institution. It requires a settlement substrate that operates on the same bilateral accounting events that underwrite today's markets, but executes them with immediate finality. By grounding settlement in double-entry accounting, credit becomes native, collateral becomes integrated, and the entire market stack can connect without reinvention. The industry's goal is clear: T+0 for the market it already has. And the architecture capable of delivering it is equally clear: a settlement layer built on accounting, not prefunding; obligations, not representations; bilateral truth, not unilateral state—implemented on blockchain technology without inheriting its unilateral constraints.

7. Conclusion

The industry's pursuit of T+0 has been constrained not by technology, but by the assumptions inherited from existing blockchain designs. These systems can represent assets or claims, but they cannot express the bilateral accounting events that define real settlement, support credit, or integrate collateral. Financial markets already possess the mechanism that resolves these requirements with legal certainty: double-entry accounting.
By relocating that accounting event into the settlement layer, the market gains a substrate that expresses obligations natively, supports credit directly, integrates collateral automatically, and delivers finality as an intrinsic property of the system. This is not a departure from existing market structure; it is its natural evolution. A settlement layer built on accounting provides the industry with what it has been seeking all along—T+0 for the market it already has, delivered with the safety, clarity, and institutional coherence it cannot afford to lose.

Postscript

A. Synthetic Encumbrance

The industry has attempted to use T+0 constructs and a range of blockchain-based mechanisms to address over-collateralization. These designs rely on lockups, vaults, wrappers, staking constructs, and tokenized representations of collateral. All of them freeze cash or securities rather than encumber them. Frozen assets become references for a placeholder that moves around on a blockchain, but the placeholder lacks the institutional properties that make the underlying asset recoverable—a legal requirement for any enforceable collateral arrangement.

Because these constructs cannot express encumbrance natively, they operate as forms of synthetic encumbrance. The blockchain layer simulates the rights state, but the authoritative lifecycle, custody, and settlement truth remain off-chain. To complete a transaction legally, each operation must call back to the same accounting systems that support T+2 and T+1 settlement, extinguish the liability there, and only then allow the collateral to be reused.

This model works in theory and has been validated through several real-world proof-of-concept implementations. But it carries structural limitations: execution risk, reconciliation overhead, audit exposure, and operational fragility. These limitations arise because the blockchain layer is not an extension of the accounting system. As a result, these approaches must operate asynchronously and rely on orchestrated code paths and abstracted liability metadata to reproduce settlement events in the off-chain systems that remain authoritative. This indirection is necessary because the blockchain is limited to a non-authoritative representation of intent.

Author: Bill McCahey
+1 551-449-0343
bill@ledger4securities

B. The Root Cause

The reason these constructs operate as synthetic encumbrances is architectural. They are workarounds created by the treatment of liabilities on blockchains. All mainstream blockchains inherit the foundational Bitcoin protocol model, where only balances are written to the ledger. Commonly called asset-only ledgers, these systems were designed to support an alternative monetary network that does not rely on intermediaries.

Eliminating intermediaries was one of the first principles Satoshi Nakamoto implemented when he appropriated blockchain technology to introduce the Bitcoin protocol. Blockchains themselves date to 1991, introduced by Haber and Stornetta, but Satoshi's 2008 design repurposed the concept to create a monetary system that removed intermediaries, which he viewed as the source of most failures in modern finance. His most fundamental design choice was to eliminate the recording of liabilities on the ledger—changing the nature of a claim and, more importantly, the nature of an obligation.

This design forces gross settlement between parties. Counterparties must be fully prefunded before they can transact. No obligation going in, none coming out. If the transaction does not meet these qualifications, it does not occur. Once final, it cannot be reversed.

This is the antithesis of today's regulated markets, where credit lives and liquidity is deep. And here is the remarkable part: the asset-only ledger has never been redesigned since its inception. The only major evolution is Ethereum—but with Ethereum, accounting must be reconstructed on a thin compute layer (the EVM or an L2) where "smart contracts" contain data and fragments of code.

These functions cannot be reduced to asset transfers on a Bitcoin-style protocol. They can be approximated through smart contracts, but doing so requires bespoke layers of orchestrated components working together to reproduce the accounting systems and business rules we already have. The result is a fragile surface of variable parts subject to execution risk. Even then, they are not the accounting functions themselves, and they remain dependent on integration with off-chain systems.

Author: Bill McCahey
+1 551-449-0343
bill@ledger4securities