Cross-Chain Settlement Without Bridges

Bilateral Netting, Off-Chain Verification, and Native Chain Finality

Shane Calder
132 Engineering
April 2026

Decentralised Finance · Vol. 2 · Self-Custody Infrastructure Series
Companion to: *The Self-Custody Option* (Vol. 1) — https://self-custody.shanescalder.com

Abstract

Cross-chain value transfer is predominantly treated as a movement problem: assets are moved from one chain to another, typically through bridge infrastructure that locks tokens on the source chain and mints representations on the destination chain. This paper argues that the movement assumption is architecturally costly. Cumulative losses from bridge exploits have exceeded $1.3 billion in 2022 alone [6], with industry estimates of total losses since 2021 ranging from $2 billion to $2.8 billion depending on methodology and which incidents are classified as bridge-specific. These losses are attributable not solely to implementation failures but to a structural property of bridges: they create a persistent custody surface that scales its attack surface with total value locked.

We describe an alternative model: bilateral netting with native chain settlement. Two parties on different chains with opposing intent each create escrow objects on their native chain, linked by a shared cryptographic commitment (PREIMAGE-SHA-256). Settlement executes natively on each chain. No assets cross. No bridge holds custody. An off-chain verification layer records the agreement, monitors execution, and produces Merkle proofs for auditability without requiring consensus, tokens, or node infrastructure.

This model draws on established practice. Multilateral netting settles over $7 trillion daily in foreign exchange markets through CLS Bank, with netting efficiencies above 96% [7, 9]. Self-reported data from cross-chain intent protocols suggests that a substantial majority of cross-chain flows have natural counterparts [10], though independently verified figures are not yet available. The escrow primitives required for bilateral cross-chain settlement exist natively on several chains, with important caveats regarding asset-type support discussed in Section 3. This paper describes the architecture, analyses its properties relative to bridge infrastructure, acknowledges known limitations including the free option problem and fiat settlement trust assumptions, and identifies open gaps that represent opportunities for builders and researchers.

Keywords: *cross-chain settlement, bilateral netting, escrow primitives, bridge alternatives, off-chain verification, Merkle proofs, HTLC, self-custody, stateless settlement*

1. Introduction: The Bridge Assumption

Every major cross-chain settlement mechanism in production today moves assets between chains. The models differ in implementation but share a structural property: they create a custody surface that holds assets in transit.

1.1 Three Models of Cross-Chain Movement

Lock-and-mint bridges. Assets are locked in a smart contract on the source chain. The bridge mints a synthetic representation on the destination chain. The locked assets remain in the bridge contract for the duration—sometimes indefinitely. The Ronin Bridge held approximately $624 million in locked assets when it was exploited in March 2022 [1]. The Wormhole Bridge held an estimated $320–326 million when its signature verification was bypassed in February 2022 [2]. The BNB Bridge held approximately $568 million when a proof verifier bug was exploited in October 2022 [3]. In each case, the locked asset pool was the target.

Message relay networks. A network of validators or relayers observes events on the source chain and attests to them on the destination chain. Assets are released when sufficient attestations are received. The custody surface is the validator set—compromise enough validators and the attestation is forged. The Harmony Horizon Bridge required only two of five validator signatures, and its compromise resulted in approximately $100 million in losses [4].

Liquidity networks. Liquidity providers pre-fund the destination chain. The user deposits on the source chain and receives from the LP on the destination chain. The LP rebalances later. The custody surface is the liquidity pool and the rebalancing mechanism. While faster for the user, the systemic risk is transferred to the liquidity providers and the protocol managing rebalancing.

1.2 The Structural Problem

These three models differ in mechanism but converge on a shared vulnerability: a persistent custody surface that grows

with usage. Total value locked (TVL) in bridge contracts is often presented as a success metric. It is more accurately understood as a risk metric—every dollar locked is a dollar at risk for the duration it remains in the contract. Mandiant reported over $1.3 billion in bridge-related theft in 2022 alone [6]. CertiK reported that five bridge exploits in 2022 totalled $1.317 billion, representing 57% of total Web3 losses that year [5]. Industry aggregations suggest cumulative bridge losses since 2021 are in the range of $2–2.8 billion, though exact figures vary by which incidents are classified as bridge-specific.

The assumption underlying all three models is that cross-chain settlement requires cross-chain asset movement. Buterin [19] has argued that bridges introduce fundamental security limitations, noting that cross-chain applications face risks that do not exist in single-chain contexts. This paper argues that the movement assumption itself is unnecessary.

2. Netting: A Proven Settlement Mechanism

Bilateral and multilateral netting are not novel concepts. They are the foundation of the largest settlement systems in global finance. The application to cross-chain digital asset settlement is new; the principle is not.

2.1 CLS Bank: Netting at Scale

CLS Bank, designated a systemically important financial market utility (SIFMU) by the U.S. Federal Reserve and overseen by 23 central banks, settles foreign exchange transactions through multilateral netting across 18 currencies. In the first half of 2025, CLS reported an average daily settled value of $7.9 trillion, a 12% increase year-on-year [7]. On a single peak day in June 2024, CLS settled $19.1 trillion in payment instructions—requiring approximately $72 billion in actual funding [8]. CLS reports multilateral netting efficiencies above 96% on average [9]. The peak-day total funding reduction—approximately 99.6%—reflects the combined effect of netting plus additional liquidity-saving mechanisms (in/out swaps), not netting alone.

The principle is straightforward. If Bank A owes Bank B $100 million in EUR/USD and Bank B owes Bank A $95 million in GBP/USD, the net obligation is $5 million. Only the net amount moves. This approach to cross-product netting is well-established in traditional derivatives and securities financing infrastructure [20]. CLS extends it multilaterally across 76 settlement members and over 38,000 third-party participants [7].

2.2 Netting Applied to Cross-Chain Flows

Self-reported data from cross-chain intent protocols suggests that similar flow symmetry exists in digital asset markets. Everclear (formerly Connext), a cross-chain clearing protocol, has reported that over 80% of cross-chain transfer flows in a 24-hour period have natural counterparts and are nettable [10]. This figure originates from a research piece sponsored by Everclear and has not been independently verified. Nevertheless, the underlying logic is sound: for every user moving assets from Chain A to Chain B, there is often a user seeking the reverse. The question is whether matching infrastructure can capture this symmetry efficiently.

The insight is not that netting is possible—CLS has proven this at $7 trillion daily. It is that cross-chain demand may already be substantially bilaterally symmetric. The flows likely exist. The matching infrastructure designed specifically for netting does not.

An important terminological distinction: CLS performs netting at the *settlement layer*—gross obligations between counterparties are compressed to net obligations, and only the net amount settles. The protocol described in this paper (Section 4) does not compress obligations. Both parties escrow the full notional and both escrows release in full. The "netting" in this architecture operates at the *flow-matching layer*: instead of two parties each routing gross transfers through bridge infrastructure, their opposing intents are matched and they settle bilaterally on their native chains. This eliminates the bridge, not the gross settlement. The CLS analogy is instructive for the matching principle and the empirical observation that most flows are bilaterally symmetric, but the settlement mechanics differ. This distinction should be kept in mind throughout the paper.

2.3 Settlement Models Compared

3. Native Chain Escrow as a Settlement Primitive

Bilateral cross-chain settlement requires each party to lock assets on their native chain under conditions that enable atomic release. Several chains provide native or near-native primitives for this purpose, with important caveats.

3.1 XRPL Escrow

The XRP Ledger provides escrow as a first-class ledger object (EscrowCreate, EscrowFinish, EscrowCancel transactions). An escrow object can be created with a FinishAfter time condition, a CancelAfter time condition, and a PREIMAGE-SHA-256 crypto-condition per the IETF Crypto-Conditions specification [11]. The escrowed XRP is held in a ledger object—not in a smart contract, not in a pool, not in a third-party account. The creator's account remains the owner. The funds are released only when the cryptographic condition is fulfilled and the time condition is met. If neither condition is satisfied before CancelAfter, the escrow is refundable to the creator.

Token escrow (XLS-85): Prior to February 2026, XRPL native escrow supported only XRP. The activation of the XLS-85 amendment on February 12, 2026 [22] extended escrow to all Trustline-based tokens (IOUs) and Multi-Purpose Tokens (MPTs). Stablecoins such as RLUSD, tokenized

Table 1: Major cross-chain bridge exploits (2021–2023). Approximate loss figures compiled from CertiK [5], Mandiant/Google Cloud [6], and incident-specific post-mortems [1–4]. Figures vary across sources due to differences in token pricing at time of exploit versus time of reporting.

Bridge	Date	Approx. Loss	Attack Vector
Poly Network	Aug 2021	$612M	Access control exploit
Wormhole	Feb 2022	$320–326M	Signature verification bypass
Ronin (Axie Infinity)	Mar 2022	$624M	Validator key compromise (social engineering)
Harmony Horizon	Jun 2022	$100M	Private key compromise (2-of-5 multisig)
BNB Bridge	Oct 2022	$568M	Proof verifier bug
Nomad	Aug 2022	$190M	Trusted root misconfiguration
Multichain	Jul 2023	≈$130M	CEO-linked key compromise

Table 2: Comparison of settlement models. Bilateral netting eliminates the persistent custody surface but introduces a counterparty matching dependency. Netting efficiency is theoretical when a perfect match exists; real-world efficiency depends on flow volume and matching infrastructure.

Property	Bridge (Lock-and-Mint)	Multilateral Netting (CLS)	Bilateral (This Paper)
Asset movement	Full gross amount crosses	Net amount only	No movement—native settlement
Custody surface	Bridge contract holds TVL	CLS holds during window	Escrow on native chain only
Persistent risk	TVL at risk continuously	Risk during settlement window	No residual risk after settlement
Intermediary	Bridge protocol	CLS (regulated SIFMU)	None—bilateral agreement
Netting efficiency	0% (gross transfer)	96–99.6% (reported)	Dependent on match availability
Failure mode	Systemic (all users)	Systemic (all members)	Bilateral (two parties only)
Cold-start	None (always available)	Requires member network	Requires counterparty matching

real-world assets, and any issued token on XRPL can now be escrowed natively with the same PREIMAGE-SHA-256 conditions and time-locks available for XRP. Token issuers retain control over whether their assets can be escrowed through issuer-level flags, preserving compliance and governance structures. This amendment is significant for the bilateral settlement architecture described in this paper: settlement assets can now be escrowed directly, without intermediate AMM conversion, eliminating deploy-side slippage and simplifying the transaction sequence.

The cost is the base transaction fee (typically 10 drops, approximately $0.00001) plus a 0.2 XRP owner reserve for the duration the escrow object exists on the ledger [11]. During periods of elevated network activity, transaction fees may increase via the XRPL's fee escalation mechanism, though this has historically been rare.

3.2 Bitcoin HTLC

Bitcoin supports Hash Time-Locked Contracts (HTLCs) through its scripting language (OP_SHA256, OP_HASH160, OP_CHECKLOCKTIMEVERIFY, OP_CHECKSEQUENCEVERIFY). An HTLC locks bitcoin in a script that can be redeemed by providing the preimage to a hash or refunded to the sender after a timeout [12]. HTLCs are the mechanism underlying the Lightning Network and have been used for atomic cross-chain swaps since their description by Tier Nolan in 2013 and formalisation by Herlihy in 2018 [13, 14].

Bitcoin's HTLC differs from XRPL's escrow in implementation but provides equivalent functional properties: time-locked custody, hash-locked release, automatic refund, and no intermediary. The primary differences are transaction cost (higher and variable on Bitcoin), confirmation time (approximately 60 minutes for six confirmations versus 3–5 seconds on XRPL), and script expressiveness. Following the XLS-85 amendment, XRPL now supports locking both XRP and issued tokens natively; Bitcoin HTLCs lock BTC directly but do not extend to issued assets in a comparable way.

3.3 Ethereum and EVM Chains

Ethereum and EVM-compatible chains implement escrow through smart contracts. While functionally equivalent, the trust model differs: the user's assets are held in a contract address, not in a native ledger object owned by the user. The contract's code is the custody mechanism, and bugs in that code have been a primary attack vector in bridge exploits. Additionally, on EVM chains the preimage reveal transaction is visible in the public mempool before confirmation, creating a front-running risk: MEV searchers can extract the preimage and claim the counterparty's escrow before the original revealer's transaction confirms. This risk does not apply to XRPL (no public mempool) or Bitcoin (where the claim transaction structure differs).

3.4 Escrow Primitive Comparison

4. Cross-Chain Bilateral Settlement via Shared Preimage

The mechanism for bilateral cross-chain settlement is a direct application of the HTLC-based atomic swap protocol described by Nolan [13] and formalised by Herlihy [14]. The hash-lock, time-lock asymmetry, preimage reveal sequence, and refund path are established mechanisms. The contribution of this paper is not the protocol itself but its framing as settlement infrastructure for matched cross-chain intent, the off-chain verification layer (Section 6), the comparative risk analysis relative to bridge architecture (Section 7), and the integration of native chain escrow primitives including XRPL's XLS-85 token escrow. The distinction from a standard atomic swap is contextual rather than mechanical: in a swap, two parties exchange assets of different types; in bilateral settlement, two parties with opposing cross-chain intent settle natively, with no assets crossing chains.

4.1 Protocol Description

Both chains must support the same hash function for the shared preimage—in this protocol, SHA-256. If one chain uses a different hash construction (e.g., RIPEMD-160(SHA-256) as in Bitcoin's OP_HASH160), the escrow on that chain must be configured to use the compatible function (OP_SHA256) to ensure the same preimage satisfies both locks.

4.2 Time-Lock Asymmetry

The requirement that $T_2 < T_1$ introduces a calibration problem when the two chains have different finality characteristics. The buffer between T_2 and T_1 must be at minimum: (a) the time required for Party B to observe the preimage reveal on Chain 1, plus (b) the time required to construct, sign, and confirm the claim transaction on Chain 2, plus (c) a safety margin for network congestion or temporary chain unavailability.

For an XRPL-to-Bitcoin settlement, Chain 1 (XRPL) confirms in 3–5 seconds while Chain 2 (Bitcoin) requires approximately 60 minutes for six confirmations. A conservative parameterisation would set the buffer at 2–4 hours. Formal derivation of minimum safe buffers per chain pair is an open research problem (see Section 10.2).

4.3 Atomicity Properties

The protocol provides conditional atomicity: if both parties cooperate, both settlements execute. If either party defects, both escrows refund after their respective timeouts. No party loses assets regardless of counterparty behaviour. The failure mode is a failed trade, not a lost asset—stronger than bridge-mediated transfers, which can lose assets if the bridge fails mid-transfer [15].

Algorithm 1: Bilateral Cross-Chain Settlement

Input: Agreed terms: notional, duration, fee, conditions
Output: Native settlement on both chains, or refund

Setup: Party A generates random secret s;
Compute $h \leftarrow$ SHA-256(s);
Share h with Party B;
Lock: Party A creates escrow on Chain 1:;
  beneficiary $\leftarrow$ Party B;
  hash-lock $\leftarrow h$, timeout $\leftarrow T_1$;
Party B creates escrow on Chain 2:;
  beneficiary $\leftarrow$ Party A;
  hash-lock $\leftarrow h$, timeout $\leftarrow T_2$ where $T_2 < T_1$;
Settlement: Party A reveals s by finishing escrow on Chain 1;
Party B observes s on-chain;
Party B finishes escrow on Chain 2 using s;
Refund: if s *not revealed before* T_2 **then**
  Party B's escrow refunds automatically;
  Party A's escrow refunds after T_1;
  No assets lost by either party;
end

4.4 The Free Option Problem

A known limitation of any HTLC-based protocol is the free option problem, identified by Robinson [16], analysed in the context of fair exchange by Dziembowski et al. [21], and discussed further in subsequent literature. Between the Lock and Settlement phases, Party A holds a free American-style option: they can observe market conditions and choose whether to reveal the preimage (completing the settlement) or allow the timeout to expire (cancelling it). If conditions have moved unfavourably for Party A, they are incentivised to abandon—at no cost beyond the time value of locked capital.

This is a genuine limitation. It is partially mitigated by: (a) shorter timeout windows, which reduce the observation period; (b) upfront fees or premium payments that are non-refundable even on timeout, introducing a cost to abandonment; and (c) reputation consequences in the verification layer, where a pattern of forced timeouts is visible in the counterparty's settlement history. However, none of these fully eliminate the optionality. Any implementation of this protocol must account for the free option cost when pricing settlements, and counterparties should factor it into their risk assessment.

4.5 The Sore Loser Problem

Xue and Herlihy [17] describe a related attack: a party may prefer to force a timeout—even at the cost of their own

Table 3: Comparison of escrow primitives. Following the XLS-85 amendment (February 2026), XRPL provides native self-custody escrow for all issued tokens with no contract risk. Bitcoin HTLCs lock BTC directly. EVM contracts offer flexibility but introduce code-level risk. Note: confirmation times represent time to first inclusion, not finality. Bitcoin achieves strong probabilistic finality after six confirmations (≈60 min). Post-merge Ethereum achieves Casper FFG finality in approximately two epochs (≈12.8 min), not after a single slot.

Property	**XRPL Escrow**	**Bitcoin HTLC**	**EVM Contract**
Implementation	Native ledger object	Script-level	Deployed contract
Custody model	Creator retains ownership	Locked in script output	Held by contract address
Hash-lock	PREIMAGE-SHA-256	OP_SHA256 / OP_HASH160	Solidity function
Time-lock	FinishAfter / CancelAfter	CLTV / CSV	block.timestamp
Auto-refund	CancelAfter (native)	Timeout path in script	Requires function call
Asset support	XRP + all issued tokens (XLS-85)	BTC (native)	Any ERC-20 token
Confirmation	3–5 seconds	≈60 min (6 conf.)	12–15 sec (1 conf.)
Tx cost	≈10 drops ($0.00001)	Variable ($0.50–$5+)	Variable ($1–$50+)
Contract risk	None (native primitive)	Minimal	Depends on code quality

time and locked capital—to prevent the counterparty from profiting. In bilateral settlement, this manifests as Party A refusing to reveal the preimage when the settlement would benefit Party B. The cost to Party A is the time-value of locked capital. This attack is bounded—the attacker pays a real cost—but it is not zero-cost. Mitigation strategies include shorter timeouts, reputation penalties, and the credible threat of future non-cooperation. Formal game-theoretic analysis in the bilateral settlement context is an area for future work.

5. Threat Model

The bilateral settlement protocol assumes the following:

Chain liveness. Both chains are live and processing transactions for the duration of the settlement. If either chain halts, timeouts may not execute as expected. This is a dependency shared with bridges and all cross-chain protocols.

Preimage secrecy. Party A's secret s remains unknown to Party B until Party A reveals it on-chain. Compromise of s before the Lock phase completes would allow Party B to claim without Party A's consent. Standard key management practices apply.

Chain observability. Party B can observe the preimage reveal on Chain 1 in sufficient time to claim on Chain 2 before T_2 expires. A sustained denial-of-service attack on Party B's ability to observe Chain 1 could prevent them from claiming, resulting in a timeout on both sides (no loss, but a failed settlement).

Hash function security. SHA-256 is computationally infeasible to invert or find collisions for. A break in SHA-256 would undermine the protocol, but would also undermine Bitcoin and most existing cryptographic infrastructure.

Rational adversary (limited). The protocol does not assume fully rational behaviour. The free option and sore loser problems (Sections 4.4, 4.5) represent deviations from cooperative behaviour that the protocol tolerates—both result in timeout and refund, not loss. The protocol guarantees *safety* (no party loses assets) under any adversary behaviour. It guarantees *liveness* (settlement completes) only under cooperative behaviour.

6. The Verification Layer

The gap between two independent escrows with a shared preimage and a working settlement system is verification. Both parties need assurance that terms were agreed, escrows were created correctly, and the history is auditable.

6.1 Requirements

A verification layer must provide: (1) a tamper-evident record of agreed terms before either escrow is created; (2) confirmation that both escrows match the agreed terms; (3) a record of settlement execution or refund on both chains; (4) the ability for either party to prove any fact about the settlement to a third party without revealing full contract data; and (5) optional anchoring to a public chain for universal verifiability.

6.2 Architecture

Record creation. Each significant event—agreement, escrow creation, preimage reveal, settlement confirmation, or refund—is captured as a record entry with a timestamp, event type, transaction identifiers, and a cryptographic link to the previous entry.

Block formation. Entries are grouped into blocks. Each block is sealed with a Merkle root. The root of each block is chained to the previous block's root, producing an append-only, tamper-evident structure.

Selective disclosure. Any single entry can be proven to exist within a block using a Merkle proof—a logarithmic-sized authentication path from the entry to the block's root.

Optional anchoring. The Merkle root of any block can be anchored to a public blockchain as a single transaction—

a memo field on XRPL, an OP_RETURN on Bitcoin, or calldata on Ethereum.

6.3 Properties and Limitations

This architecture provides immutability, verifiability, privacy (underlying data is off-chain), and efficiency (no consensus overhead). The construction draws on well-established verifiable data structures, including certificate transparency logs and authenticated append-only data structures. The specific contribution is not the data structure itself but its application to bilateral settlement audit trails—providing selective disclosure of settlement events without requiring consensus infrastructure or exposing private contract terms.

A limitation is that the verification layer is operated by the settlement participants. If both parties collude to fabricate records, the verification layer cannot prevent this. For settlements where third-party auditability is critical, an independent observer could provide additional assurance, at the cost of introducing a third participant.

7. Stateless Settlement vs Permanent Infrastructure

A bridge is permanent infrastructure. It exists between settlements—holding assets, requiring maintenance, presenting an attack surface continuously. Bilateral settlement is stateless—each settlement is independent, and nothing persists after completion.

7.1 Risk Profile Comparison

Bridge risk is systemic and cumulative. Let $R_b(t)$ represent aggregate bridge risk at time t:

$$R_b(t) \propto \text{TVL}(t) \times P(\text{exploit}) \tag{1}$$

where $\text{TVL}(t)$ is the total value locked and $P(\text{exploit})$ is the probability of a successful exploit per unit time. This function is non-zero continuously and grows with usage. A single successful exploit affects all users whose assets are locked.

Bilateral settlement risk is bounded and isolated:

$$R_h(t) = \sum_i V_i \times P(\text{counterparty\_failure}_i) \tag{2}$$

where the sum is over settlements currently in the lock phase. Importantly, V_i represents the time-value-of-capital exposure during the lock phase, not loss-of-principal risk—since the protocol guarantees refund on failure, the direct financial cost of a failed settlement is the opportunity cost of locked capital for the timeout duration, plus any adverse market movement during that period. This function is bounded and drops to zero between settlements. Adjacent settlements are cryptographically independent.

These are informal risk characterisations, not formal security proofs. They are intended to highlight the structural difference: bridge risk is pooled and persistent; bilateral risk is isolated and transient.

8. Fiat Integration as a Fourth Settlement Rail

The bilateral model extends to fiat settlement, but with significant caveats. The crypto leg uses escrow primitives from Section 3. The fiat leg settles through traditional banking rails. The verification layer records both legs.

8.1 The Confirmation Honesty Problem

The fiat leg introduces a trust assumption the crypto leg does not have. Party A must honestly confirm receipt of the bank transfer before revealing the preimage. If Party A claims non-receipt—whether truthfully or not—the settlement times out and Party A retains both the escrowed crypto and the fiat payment.

This is a fundamental asymmetry. Possible mitigations include: (a) a fiat attestation service where the receiving bank confirms the transfer independently; (b) payment rails that provide cryptographic receipts; (c) escrow agents for the fiat leg; or (d) reputation consequences in the verification layer. None fully resolve the problem. Fiat integration remains an area where the trust model is weaker than the pure crypto-to-crypto case, and implementations should be transparent about this limitation.

8.2 Additional Fiat Asymmetries

Beyond the confirmation problem, fiat settlement introduces slower execution (hours to days), reversibility (chargebacks), opacity (not independently verifiable), and jurisdictional compliance requirements (KYC, AML). The bilateral model does not eliminate these constraints. It provides a framework in which they can be satisfied by each party independently, without a central intermediary processing both sides.

9. Pools of Willingness: Intent-Based Counterparty Discovery

Bilateral netting requires counterparties with opposing intent. In mature markets, flow volume ensures counterpart availability. In early-stage cross-chain markets, counterpart discovery is the cold-start problem. CLS Bank's netting efficiency is achieved within a regulated membership of vetted financial institutions with legal settlement obligations, regulatory oversight from 23 central banks, known participant identities, and deep liquid fiat currency markets. An open intent pool for cross-chain settlement would lack this institutional scaffolding, making the incentive design for participation and reliable settlement behaviour a distinct challenge. The CLS analogy is instructive for the netting principle but should not be read as implying that CLS-scale outcomes are

directly transferable to a pseudonymous, volatile-asset environment.

The proposed mechanism is an intent pool: a registry of parties who have declared willingness to settle on a particular chain pair. This is not a liquidity pool—no assets are deposited, no protocol holds custody. It is a registry of intent: chain pair, direction, notional range, duration range, and fee tolerance.

9.1 Reputation from Verifiable History

Counterparty evaluation requires verifiable performance data, not subjective ratings. The verification layer provides this natively: every completed settlement produces an auditable record of whether the settlement executed on time, whether terms matched, and whether the counterparty honoured conditions. Reputation in this model is a queryable history. Credit decisions become evidence-based rather than trust-based. This does not eliminate counterparty risk, but it makes that risk assessable from auditable data.

9.2 Liquidity Fragmentation

Bridges pool liquidity—any user can access the full pool. Bilateral settlement fragments liquidity across individual pairs. In early markets with thin flow, this fragmentation could make bilateral settlement impractical even when the architecture is sound. Standardised contract sizes or tiered notional ranges could mitigate this, at the cost of reduced flexibility.

10. Open Gaps and Opportunities

The following gaps represent unsolved problems. Each is also an opportunity.

10.0.1 Matching Infrastructure

The cross-chain equivalent of CLS Bank's matching layer does not yet exist. Building it is both a technical and a commercial opportunity.

10.0.2 Time-Lock Calibration

A formal model for cross-chain time-lock parameterisation across chains with heterogeneous finality is needed. This is an academic research problem with immediate practical application.

10.0.3 Dispute Resolution

When the verification layer's record and on-chain state disagree, a resolution mechanism is needed. This bridges cryptography and law.

10.0.4 Regulatory Classification

Is bilateral cross-chain netting a financial service, a money transmission activity, or a clearing arrangement? Recent SEC guidance on tokenized securities [18] suggests regulators are engaging with adjacent questions.

10.0.5 Fiat Settlement Assurance

An attestation layer for fiat settlement—where banks confirm transfer completion in a cryptographically verifiable format—is the hardest open problem in fiat integration.

10.0.6 Cross-Chain Reputation Portability

A Merkle-anchored reputation aggregation protocol portable across chain pairs is needed for efficient counterparty discovery.

10.0.7 Insurance and Risk Products

Every gap above is an insurable risk. Parametric insurance products priced against verifiable settlement data represent a natural financial products layer. This connects to the broader question of structured products from escrow-backed instruments, which is the subject of a companion paper.

11. Conclusion

Cross-chain settlement does not require cross-chain asset movement. It requires matched intent, native chain escrow, a shared cryptographic commitment, and a verification layer. The netting principle is proven at $7 trillion daily in foreign exchange. The escrow primitives exist on multiple chains, with caveats regarding asset-type support.

The bridge model persists not because it is the correct architecture but because it was the first architecture. The cost of that framing—significant losses, persistent custody surfaces, systemic failure modes—is documented. The bilateral model offers different properties: no persistent custody surface, no systemic failure mode, and risk bounded per settlement rather than cumulative.

These properties are not strictly superior in all dimensions. Bilateral matching is harder than routing through a bridge. The cold-start problem is real. The free option and sore loser problems introduce costs that must be priced. Fiat integration introduces trust assumptions the pure crypto model avoids. This paper has attempted to present both the strengths and the limitations honestly. It should also be noted that this is a theoretical architecture: no simulation, prototype, or analysis of historical cross-chain flow data has been conducted. The claims about matching efficiency and risk reduction, while grounded in established principles and analogous systems, await empirical validation.

The open gaps described in Section 10 are the demand signal for a new class of settlement infrastructure—matching

layers, calibration standards, dispute protocols, reputation systems, and insurance products. The primitives are ready. The architecture is described. What remains is to build, to test, and to measure whether the theoretical properties hold in practice.

This paper is the second in the Self-Custody Infrastructure Series. The first paper, **The Self-Custody Option** *(Vol. 1), describes escrow-settled, oracle-free bilateral derivatives on a single chain. The third paper will address the financial products layer—self-secured lending, verifiable reserves, and structured products—that may emerge from this settlement infrastructure.*

Research conducted and architecture developed by Shane Calder, April 2026. Developed in collaborative reasoning with Anthropic Claude. AI tools were used throughout the research, drafting, and review process. All claims, references, and architectural decisions were verified by the author.

Free to read and share. Not to reproduce without written permission.
© 2026 Shane Calder · 132 Engineering · 132eng.dev

References

[1] Chainalysis, "Axie Infinity's Ronin Bridge Hack: DPRK Hack Seizure," *Chainalysis Blog*, 2022. Loss of 173,600 ETH and 25.5M USDC (≈$624M). Attributed to DPRK Lazarus Group.

[2] Wormhole, "Wormhole Incident Report," Feb. 2022. ≈120,000 wETH ($320–326M) exploited via signature verification bypass. Losses covered by Jump Crypto.

[3] BNB Chain, "BNB Chain Bridge Incident Report," Oct. 2022. ≈2M BNB tokens ($568M) created via proof verifier bug. Chain halted.

[4] Harmony, "Horizon Bridge Hack Post-Mortem," Jun. 2022. ≈$100M stolen via private key compromise of 2-of-5 multisig.

[5] CertiK, "Cross-Chain Vulnerabilities and Bridge Exploits in 2022," CertiK Research, Aug. 2022. Five bridge attacks totalling $1.317B (57% of total Web3 losses).

[6] Mandiant/Google Cloud, "Decentralized Robbery: Dissecting the Nomad Bridge Hack," Google Cloud Blog, Mar. 2024. Reports $1.3B+ in bridge theft in 2022.

[7] CLS Group, "CLS Welcomes U.S. Bank to CLSSettlement," Press Release, Oct. 2025. Average daily settled value $7.9T in H1 2025.

[8] CLS Group / Finance Magnates, "CLS Settles Record $19.1T in FX Payment Instructions," Jun. 2024. ≈$72B funding required (0.38% of gross).

[9] Euromoney, "World's Best FX Post-Trade Solution 2025: CLS," Awards Report, Sep. 2025. Netting efficiencies above 96%; total funding reduction (incl. in/out swaps) above 99%.

[10] The Block Research, "How Everclear Is Building a Decentralized Clearing Layer," Feb. 2025. 80%+ of cross-chain flows nettable. Note: sponsored research.

[11] Ripple, "XRPL Escrow Documentation," xrpl.org. EscrowCreate/Finish/Cancel. PREIMAGE-SHA-256 crypto-conditions. XRP-only prior to XLS-85; issued tokens now supported.

[12] Bitcoin Wiki, "Hash Time-Locked Contracts," bitcoin.it/wiki. OP_SHA256, CLTV, CSV operations.

[13] T. Nolan, "Alt chains and atomic transfers," BitcoinTalk Forum, May 2013. https://bitcointalk.org/index.php?topic=193281.msg2224949.

[14] M. Herlihy, "Atomic Cross-Chain Swaps," *Proc. ACM PODC '18*, pp. 245–254, 2018. DOI: 10.1145/3212734.3212736. arXiv: 1801.09515.

[15] M. Herlihy, B. Liskov, and L. Shrira, "Cross-Chain Deals and Adversarial Commerce," *Proc. VLDB Endowment*, vol. 13, no. 2, pp. 100–113, 2019. DOI: 10.14778/3364324.3364326.

[16] D. Robinson, "HTLCs Considered Harmful," Stanford Blockchain Conf., 2019. Analyses the free option problem in HTLC protocols.

[17] Y. Xue and M. Herlihy, "Hedging Against Sore Loser Attacks in Cross-Chain Transactions," *Proc. ACM PODC '21*, pp. 155–164, 2021.

[18] SEC Div. of Corp. Finance et al., "Statement on Tokenized Securities," Jan. 28, 2026. Taxonomy for tokenized securities incl. synthetic structures.

[19] V. Buterin, "Why I'm Optimistic About Multi-Chain but Pessimistic About Cross-Chain," Reddit, Jan. 2022.

[20] ISDA, "Accounting for Cross-Product Netting," Dec. 2023. Netting frameworks under master netting agreements.

[21] S. Dziembowski, L. Eckey, and S. Faust, "FairSwap: How to Fairly Exchange Digital Goods," *Proc. ACM CCS*, 2018.

[22] RippleX, "Token Escrow (XLS-85) Is Now Live on XRPL Mainnet," Feb. 12, 2026. Escrow for all Trustline tokens and MPTs. 88%+ validator support.