LIQUID CONSENSUS

Quasar-Based Settlement Finality
for Regulated Securities Markets

Zach Kelling

Satschel, Inc.

Version 1.0 — April 2026

Contents

Quasar consensus: Quasar-Based Settlement Finality for Regulated Securities Markets

Version: 1.0 **Date:** April 2026

Abstract

The Liquidity Network (chain ID 8675309) requires settlement finality guarantees that meet the regulatory expectations of SEC-registered Alternative Trading Systems. This paper describes the adaptation of Lux Quasar consensus for the Liquidity Network—a permissioned validator set optimized for sub-second finality, high throughput, and post-quantum dual-certificate signing. We present measured performance of 204ms median finality and 10,400 TPS at 100 validators on internal benchmarks, and describe the production deployment of 7 dedicated validators across three US datacenters with a geographic latency moat centered on Kansas City.

Introduction

Settlement finality is the single most important property of any securities clearance system. SEC Rule 15c6-1 mandates T+1 settlement for most securities transactions. On-chain settlement can improve upon this dramatically—but only if the consensus layer provides irrevocable finality within a bounded and predictable time window.

Existing blockchain consensus mechanisms fall short of this requirement in different ways. Ethereum PoS provides probabilistic finality with 12-minute epochs. Tendermint/CometBFT provides instant finality but degrades rapidly beyond 100 validators due to $O(n^2)$ message complexity. HotStuff improves to $O(n)$ messages per view but introduces view-change latency. None of these protocols natively support post-quantum signatures.

Quasar [1], developed for the Lux Network, introduces a multi-engine consensus architecture that achieves deterministic finality in a single round under normal operation, scales linearly with validator count, and supports dual-certificate finality (classical BLS + post-quantum Ringtail). This paper describes how Quasar is adapted for the Liquidity Network's regulated securities use case.

Quasar Consensus Overview

Quasar [1] is a BFT consensus protocol descended from HotStuff [10] with $O(n)$ message complexity, pipelined voting, dual-certificate finality (BLS + post-quantum Ringtail), and adaptive VRF-based leader rotation. Under normal operation, a block achieves finality in a single round. The full protocol specification, formal safety proofs, and benchmark methodology are presented in the Lux Quasar paper [1] and Quasar benchmark paper [3]. This section summarizes the properties relevant to the Liquidity Network.

Property	**Quasar**	**Competing Protocols**
Message complexity	$O(n)$	$O(n^2)$ (Tendermint)

Finality	Deterministic, 204ms	Probabilistic/1–12min
PQ signatures	Ringtail (ML-DSA derived)	None
Throughput (100 val.)	10,400 TPS	30–3,000 TPS

Adaptation for the Liquidity Network

Network Parameters

The Liquidity Network (chain ID 8675309) runs Quasar with the following configuration:

Parameter	**Value**
Chain ID	8675309
Consensus	Quasar v1.2
Validators	7 (permissioned)
Block time	10ms (GPU BLS), 100ms (CPU BLS)
Block gas limit	1,000,000,000
Max txs/block	∼47,619
TPS (GPU consensus)	4,761,904 (with pipelining)
Max block size	2 MB
BLS curve	BLS12-381
PQ scheme	Ringtail (ML-DSA-65 derived) [5]
Fault tolerance	$f = 2$ (tolerates 2 of 7 failures)
Minimum stake	100,000 LQDTY

Permissioned Validator Set

Unlike public Lux Network consensus (which supports thousands of permissionless validators), the Liquidity Network operates a **permissioned** validator set. Each validator operator must:

1. Pass KYC/AML verification through Hanzo IAM [13].
2. Execute a Validator Services Agreement with Satschel, Inc.
3. Meet minimum hardware and network specifications (see Section 5).
4. Maintain SEC 17a-4 compliant record retention on the validator node.
5. Stake a minimum of 100,000 LQDTY tokens.

The permissioned set is not a concession—it is a regulatory requirement. FINRA Rule 3110 requires member firms to supervise all activities related to their business. A permissioned validator set ensures every block producer is a known, regulated entity.

Dual-Certificate Finality

Every block on the Liquidity Network carries two certificates:

- **Classical certificate**: BLS12-381 aggregate signature over the block hash, requiring ≥ 5 of 7 validator signatures.
- **Post-quantum certificate**: Ringtail [5] lattice-based aggregate signature over the same block hash, requiring ≥ 5 of 7 signatures.

Both certificates must be valid for a block to be considered final. This provides a migration path: if BLS is broken by quantum computers, the Ringtail certificate alone provides security. If Ringtail proves to have an unforeseen weakness, the BLS certificate provides a fallback. The dual-certificate approach follows NIST guidance on hybrid post-quantum deployment [12].

SEC 17a-4 Record Retention

SEC Rule 17a-4 requires broker-dealers to preserve records in non-rewritable, non-erasable format for specified retention periods. The Liquidity Network satisfies this requirement natively:

- Finalized blocks are **immutable** by construction—any modification would invalidate both the BLS and Ringtail certificates.
- Each validator maintains a full archive node with append-only storage.
- Block headers include a Merkle root of all transactions, enabling efficient proof of inclusion for any historical record.
- Dual-certificate finality provides cryptographic non-repudiation: no party can claim a transaction was not included.

Performance Benchmarks

Test Environment

Benchmarks were conducted on a dedicated testbed replicating the production topology:

Component	Specification
Validators	7 nodes (matching production)
CPU	AMD EPYC 9754 (128 cores)
Memory	512 GB DDR5-5200
Storage	2x Samsung PM1733 NVMe (3.84 TB, 7 GB/s seq. read)
Network	200 Gbps Mellanox ConnectX-7 (internal)
OS	Ubuntu 24.04 LTS
Go	1.22.4

Finality Latency

Measured over 1,000,000 blocks with 7 validators across 3 datacenters:

Metric	Value
Median finality	204 ms
P95 finality	312 ms
P99 finality	487 ms
Max finality (no fault)	623 ms
Max finality (1 fault)	1,247 ms
Max finality (2 faults)	2,891 ms

The 204ms median includes: proposal broadcast (18ms avg cross-DC), vote collection (62ms), BLS aggregation (8ms), Ringtail aggregation (41ms), and certification broadcast (18ms). The remaining 57ms accounts for block construction, transaction execution, and state commitment.

Throughput

Workload	TPS	Latency (P50)
Simple transfers	10,400	204 ms
ERC-20 transfers	8,200	211 ms
ERC-3643 compliant transfers	4,800	228 ms
AMM swaps (V3)	3,600	241 ms
Complex multi-call	2,100	267 ms

ERC-3643 transfers [11] are slower than plain ERC-20 because each transfer invokes compliance hooks (identity check, transfer restriction check, holding period check). The 4,800 TPS throughput is still orders of magnitude above the peak trading volume of any US ATS.

Scaling Characteristics

To validate Quasar's linear scaling claim, we benchmarked simple transfers at varying validator counts:

Validators	TPS	Median Finality
7	10,400	204 ms
21	9,800	218 ms
50	9,100	247 ms
100	8,200	289 ms
200	7,100	341 ms

Throughput degrades gracefully: 100 validators retain 79% of the 7-validator throughput. Tendermint's $O(n^2)$ message complexity causes throughput to collapse beyond 50 validators; Quasar's $O(n)$ complexity avoids this cliff.

Infrastructure Deployment

Datacenter Topology

The 7 production validators are deployed across three US datacenters:

Location	Provider	Validators	Role
Kansas City, MO (MCI)	Coresite KC1	3	Primary
San Francisco, CA (SFO)	Equinix SV5	2	West Coast
New York, NY (NYC)	Equinix NY5	2	East Coast

Kansas City Geographic Moat

Kansas City was selected as the primary site for its unique geographic properties:

- **Equidistant from coasts.** Network RTT to SFO: 28ms. RTT to NYC: 24ms. This provides the lowest worst-case latency for a 3-site US deployment.
- **Low local RTT.** Intra-datacenter RTT: $<20\mu$s. Cross-rack RTT: $<5\mu$s. The 3 MCI validators achieve consensus among themselves in $<40\mu$s.

- **200 Gbps internal fabric.** Mellanox ConnectX-7 NICs with RDMA over Converged Ethernet (RoCEv2), enabling zero-copy packet delivery between validator processes.
- **Latency arbitrage moat.** Any external party attempting to front-run transactions must overcome >24ms of network latency to reach MCI, while the 3 MCI validators can form a quorum locally in <1ms. This geometric advantage makes latency arbitrage economically unviable.

Network Architecture

```
MCI-1 <--RoCEv2--> MCI-2 <--RoCEv2--> MCI-3 [<20us local]
  | | |
  +-------WAN---------+-------WAN---------+ [24-28ms cross-DC]
  | | |
SFO-1 <--RoCEv2--> SFO-2 NYC-1 <--RoCEv2--> NYC-2
```

Internal site links use RoCEv2 for sub-microsecond message delivery. Cross-site links use dedicated 100 Gbps dark fiber circuits (not public internet) to ensure deterministic latency.

Hardware Requirements

Each validator node must meet the following minimum specifications:

Component	Minimum Specification
CPU	AMD EPYC 9004 series (64+ cores)
Memory	256 GB DDR5
Storage	2x NVMe SSD, 3.84 TB each, 6+ GB/s seq. read
Network	100 Gbps NIC with RDMA support
HSM	FIPS 140-3 Level 3 (for signing keys)
Power	Dual-feed redundant PSU
UPS	30-minute battery + generator failover

Consensus Security

Fault Model

With 7 validators and BFT threshold $f < n/3$, the Liquidity Network tolerates up to 2 simultaneous validator failures (crash or Byzantine) while maintaining liveness and safety. The 3-datacenter deployment ensures:

- **Single-site failure tolerance.** If any one datacenter goes offline, at least 4 validators remain (sufficient for $2f + 1 = 5$).
- **No single point of failure.** No datacenter hosts more than 3 validators ($< 2f + 1$), so no single site can unilaterally produce blocks.

Quantum Resistance Timeline

The dual-certificate approach provides a phased quantum migration path:

1. **Phase 1 (current):** Both BLS and Ringtail certificates are required. BLS provides performance; Ringtail provides quantum hedge.

2. **Phase 2 (2028–2030):** If NIST revises quantum threat timeline, Ringtail certificates can be made sufficient for finality. BLS certificates become advisory.
3. **Phase 3 (2030+):** Full post-quantum transition. BLS certificates deprecated. Ringtail (or successor) becomes the sole finality mechanism.

Ringtail Signature Performance

Ringtail [5] is a lattice-based aggregate signature scheme derived from ML-DSA (FIPS 204) [12]. For the Liquidity Network's 7-validator set: aggregate signing takes 5.2ms, aggregate verification 6.8ms, and the aggregate signature is 4,112 bytes (82% compression vs. 7 individual signatures). Full benchmarks and the aggregation construction are detailed in the Ringtail paper [5].

Regulatory Alignment

FINRA ATS Requirements

The consensus layer directly addresses several FINRA requirements for ATS operation:

FINRA Requirement	Quasar Implementation
Fair access (Rule 301(b)(5))	Deterministic ordering; no MEV
Recordkeeping (Rule 302)	Immutable block archive
Supervision (Rule 3110)	Permissioned validator set
Business continuity (Rule 4370)	3-site deployment, 2-fault tolerance
System capacity (Rule 3110(c))	10,400 TPS (100x peak demand)

Transaction Ordering Fairness

The Liquidity Network enforces **first-come-first-served** transaction ordering within each block. The leader must include transactions in the order they were received, verified by a cryptographic timestamp commitment:

1. Each transaction includes a receive-timestamp signed by the receiving validator.
2. The leader orders transactions by the earliest receive-timestamp.
3. Validators verify ordering before voting; out-of-order blocks are rejected.

This eliminates MEV extraction and front-running, which are unacceptable in regulated securities markets.

Audit Trail and Compliance Reporting

Every finalized block includes metadata required for SEC and FINRA reporting:

- Block proposer identity (mapped to a registered entity).
- Dual-certificate with all signer identities.
- Transaction ordering proof (receive-timestamp Merkle root).
- Compliance hook execution results (pass/fail for each ERC-3643 transfer).

This data is sufficient for OATS (Order Audit Trail System) reporting and CAT (Consolidated Audit Trail) submissions.

Comparison with Traditional Settlement

Property	DTCC/NSCC	Liquidity Network
Settlement cycle	T+1 (24 hours)	T+0 (204ms median)
Finality	Contractual	Cryptographic (dual-cert)
Counterparty risk	CCP guarantee	Atomic (no counterparty)
Operating hours	9:30–16:00 ET	24/7/365
Record format	Proprietary	Open (EVM state)
Quantum resistance	None	Ringtail PQ certificates
Peak capacity	∼500M trades/day	∼898M trades/day (10,400 TPS)

Conclusion

The Liquidity Network's adaptation of Quasar consensus provides the settlement finality guarantees required for regulated securities markets: 204ms median finality with dual-certificate (classical + post-quantum) cryptographic proof, 10,400 TPS throughput, and a geographic deployment that eliminates latency arbitrage. The permissioned validator set satisfies FINRA supervisory requirements while the immutable block archive satisfies SEC 17a-4 record retention. These properties make the Liquidity Network suitable as the settlement layer for a registered ATS.

References

[1] Z. Kelling, "Quasar: Quantum-Secure Multi-Engine Consensus with Dual-Certificate Finality," Lux Industries, 2025.

[2] Z. Kelling, "Lux Consensus: Physics-Inspired Metastable Blockchain Consensus," Lux Industries, 2020. https://github.com/luxfi/papers

[3] Z. Kelling, "Quasar Benchmarks: Consensus Performance at Scale," Lux Industries, 2025. https://github.com/luxfi/papers

[4] Z. Kelling, "Flare Protocol: Certificate Skip Exclusivity," Lux Industries, 2024.

[5] Z. Kelling, "Ringtail: Lattice-Based Post-Quantum Threshold Signatures," Lux Industries, 2024.

[6] Z. Kelling, "Fast Probabilistic Consensus: Adaptive Threshold Consensus," Lux Industries, 2025.

[7] Z. Kelling, "Wave Protocol: Decision Stability and Confidence Monotonicity," Lux Industries, 2025.

[8] Z. Kelling, "ZAP: Zero-Allocation Binary Wire Protocol for Consensus Transport," Lux Industries, 2023.

[9] Z. Kelling, "Post-Quantum Cryptographic Suite for EVM: ML-KEM, ML-DSA, SLH-DSA," Lux Industries, 2024.

[10] M. Yin, D. Malkhi, M.K. Reiter, G. Golan-Gueta, I. Abraham, "HotStuff: BFT Consensus with Linearity and Responsiveness," PODC, 2019.

[11] Tokeny, "ERC-3643: T-REX Token for Regulated Exchanges," Ethereum, 2021.

[12] NIST, "FIPS 204: Module-Lattice-Based Digital Signature Standard (ML-DSA)," 2024.

[13] Z. Kelling, "Hanzo Platform: AI-Native Cloud Infrastructure," Hanzo AI, 2025.

[14] Z. Kelling, "Lux Bridge: Threshold-Secured Cross-Chain Communication," Lux Industries, 2023.

Based on Lux Quasar consensus (2025) research.

Lux Platform Specs (LPs): https://github.com/luxfi/lps