QuasarSTM 4.0 Activation — 2026-02-14

Audience: Lux validators, application developers, ecosystem partners **Activation:** 2026-02-14, 17:00 UTC **Spec:** LP-135 (production spec) and LP-010-quasar-stm-4 (paper)

TL;DR

QuasarSTM 4.0 is live on Lux mainnet as of 2026-02-14, 17:00 UTC. It is the production version of the GPU-native ordered MVCC execution fabric that 3.0 (2025-12-25) shipped as the substrate. Every gap 3.0 left open is now closed:

- **Real cryptography.** BLS12-381, Pulsar Module-LWE, and Groth16 verifiers run on-device. The HMAC-keccak placeholders that 3.0 shipped at launch are gone from production.
- **Full EVM.** All 175 opcodes including the entire `CALL`/`CREATE` family, `LOG`n, `EXTCODE*`, `RETURNDATA*`, `TLOAD`/`TSTORE`, `MCOPY`, `BLOBHASH`/`BLOBBASEFEE`. Journaled SSTORE with EIP-2929 cold/warm and EIP-3529 refunds.
- **Production validation.** Tier 1 lane-clock fast path, Tier 3 semantic reducer commit at horizon time, ConflictSpec ABI, NEMO LaneClass, dynamic per-lane versioning mode.
- **Production execution.** Fiber checkpoints per `CALL`-frame, commit horizon over Nova prefixes / Nebula causal cuts, MVCC GC bound to the horizon, Motor-style VersionBlock layout (8 inline versions per key).
- **Production GPU pipeline.** 16 service drains (the 12 from 3.0 plus AdaptiveSchedule, BridgeAttest, MarketAuction, FiberCheckpoint — A/B/M/F).
- **Production rigor.** CSMV-style commit-server scaffold, formally specified CPU reference, cross-backend determinism harness gated in CI at >= 80% line coverage on Apple Metal and NVIDIA CUDA.

The wire format is unchanged. The public host API (`QuasarGPUEngine`) is unchanged. The 3.0 invariants — determinism, ordered serializability, repair monotonicity, horizon safety — are preserved verbatim.

What changed from 3.0

Topic	3.0 (2025-12-25)	4.0 (2026-02-14)
EVM coverage	118 opcodes (no `CALL`, no `CREATE`, no `LOG`, no `EXTCODE*`)	175 opcodes — full
SSTORE	linear write-through	journaled with EIP-
BLS verifier	HMAC-keccak placeholder	real BLS12-381 pair
Pulsar verifier	HMAC-keccak placeholder	real Module-LWE sh
Groth16 verifier	HMAC-keccak placeholder	real Groth16 over B

Topic	3.0 (2025-12-25)	4.0 (2026-02-14)
Tier 1 lane-clock	sketch	production fast path
Tier 3 semantic reducers	enum stub	production reducer c
ConflictSpec ABI	not implemented	production (Static /
LaneClass	not implemented	NEMO classes (Own
Versioning mode	always-on multi-version	dynamic per-lane (S
Fiber checkpoints	not implemented	production, per CAI
Commit horizon	sketch	production (Nova pr
MVCC GC	not implemented	production, horizon-
Motor VersionBlock	not implemented	production (8 inline
GPU services	12	16 (12 + A/B/M/F)
CSMV commit server	not implemented	scaffold landed, sing
Formal CPU reference	not implemented	small executable sem
Cross-backend determinism	per-backend tests	single harness, byte-
Coverage	informal	>= 80% line coverag

The 4.0 deliverable train shipped over six weeks (2026-01-15 through 2026-02-07 spec freeze), as v0.41 through v0.49 milestones in `cevm`:

Milestone	Theme
v0.41	CUDA backend mirror — persistent CTAs, `__threadfence`, layout-byte-identical with Metal
v0.42	Cert-subject hardening, `KnownTotalOrder`, separated mode roots
v0.43	ConflictSpec ABI + LaneClass + dynamic VersioningMode + AdaptiveSchedule drain
v0.44	Real BLS12-381 pairing kernel
v0.45	Real Pulsar Module-LWE + real Groth16 verifiers
v0.46	Full EVM coverage + journaled SSTORE + EIP-2929/3529 + Tier 3 reducers
v0.47	Predictive scheduling + Chiron-style hint roots
v0.48	Motor VersionBlock + A/B/M/F drains + fiber checkpoints
v0.49	CSMV commit-server scaffold + formal CPU reference + cross-backend determinism harness

---

Performance

Apple M1 Max, 1024-tx end-to-end stress on the regulated-DEX workload mix (50% per-account orders, 20% cancels, 15% fee accumulation, 10% balance deltas, 5% admin):

Metric	3.0	4.0
Wallclock (median)	150 ms	**108 ms**
Wave ticks	8	**6**
Repair amplification	~0.97	~**0.42**

Metric	3.0	4.0
EVM coverage	118/175	**175/175**
BLS verify (per cert)	9 us (HMAC-keccak)	46 us (real BLS12-381)
Pulsar verify (per share)	7 us (placeholder)	38 us (real Module-LWE)
Groth16 verify (per cert)	8 us (placeholder)	62 us (real Groth16)

The 28% wallclock improvement is not because 4.0 makes the same operations cheaper — real cryptography is 4-7x more expensive than the HMAC-keccak placeholders. The wins come from Tier 1 lane-clock fast-pathing >80% of commits without an MVCC chain walk and from Tier 3 reducers absorbing the contended writes (fee accumulators, order-book level updates, balance deltas) that 3.0 turned into same-key SSTORE conflicts.

The canonical Block-STM hot-key test (16 same-key contending transactions producing 120 conflicts -> 120 repairs -> 16 commits) is unchanged in 4.0 and is the determinism harness's headline regression. Same input -> byte-identical roots on Metal and CUDA.

NVIDIA H100 cert costs for reference:

- BLS12-381 aggregate verify: 31 us
- Pulsar share verify: 24 us
- Groth16 verify: 42 us

Migration

Validators

1. Upgrade `cevm` to a release built from `cevm` >= v0.49.
2. Re-roll the Quasar consensus engine (LP-020) with `execution_version=4` support.
3. Verify the cert verifier table at engine creation. On production networks the verifier table must be the production table; the `EVM_DEV_HMAC_VERIFIER` flag is rejected at startup.
4. Run `cevm-genesis-replay` against the latest 3.0 round to confirm byte-identical roots before publishing.
5. Publish on the validator gossip channel and join 4.0 round production.

There is no on-disk state migration. 4.0 reads 3.0 round artifacts as historical 3.0 rounds. New rounds emitted post-activation are 4.0 rounds.

Validators that remain on 3.0 after 17:00 UTC on 2026-02-14 are rejected by the consensus engine and must upgrade to participate.

Application developers

1. Recompile contracts against the 4.0 EVM (175 opcodes). Contracts that compiled against 3.0 (118 opcodes) are forward-compatible.
2. Adopt the Solidity `lane_class` attribute on hot state variables to opt into the AdaptiveSchedule fast path. Without the attribute, the predictor and historical sources still apply; the attribute adds a higher-priority `Static` source.
3. For DEX-style hot lanes (fee accumulators, order-book levels), adopt the `Commutative` class to route writes through Tier 3 reducers. The repair-amplification reduction comes from this change.
4. Re-run integration tests against the cross-backend determinism harness; contracts that produce identical roots on Metal and CUDA are 4.0-clean.

Ecosystem partners

- Bridge operators: BridgeAttest is a new GPU drain with a stable ABI; existing bridge messages route through it without re-signing.
- DEX operators: MarketAuction runs batched matching on-device. The `AuctionMatch` reducer entry replaces per-trade SSTORE conflicts; existing order/cancel/fill flows are forward-compatible.
- Wallet integrators: no API changes. The wire-format is unchanged.

Backward compatibility

Component	3.0 -> 4.0 compatibility
Wire format (Quasar consensus)	unchanged; LP-020 cert-lane structure preserved
`QuasarRoundDescriptor` / `QuasarRoundResult`	unchanged byte-for-byte; `hint_roots` ride in the existing r
Public host API (`QuasarGPUEngine`)	unchanged
`ServiceId` enum	extended (`Count` 12 -> 16); old IDs unchanged
Layout types	`VersionBlock.MAX_INLINE_VERSIONS` 4 -> 8 (forward-com
HMAC-keccak verifier path	gated behind `EVM_DEV_HMAC_VERIFIER=1` for legacy test ve
Validators running 3.0 after 2026-02-14	rejected at the consensus engine — must upgrade

There is no on-disk state migration. Pre-activation rounds remain valid as historical 3.0 rounds; the activation height is the partition.

Known limitations (out of scope for 4.0)

The following are explicitly *not* in 4.0 and are tracked under QuasarSTM 5.0 research:

- Multi-GPU MVCC over RDMA. 4.0 ships the Motor-style VersionBlock layout (which is RDMA-friendly) and the CSMV commit-server scaffold (which is the multi-GPU pattern), but the production deployment is single-GPU per validator.
- Operation-window validation for pathological long transactions.
- Formal proof of kernel refinement against the abstract semantics. 4.0 ships the abstract semantics in `lib/consensus/quasar/spec/` (executable C++ + Lean 4 stubs) and the differential fuzzer; the refinement proof is 5.0.
- Distributed CSMV commit server (4.0 ships single-GPU).
- 800G InfiniBand / GPUDirect cell topologies.

These are not on the 4.0 production path. They will land in 5.0 on a schedule TBD post-activation review.

References

- **LP-135** — QuasarSTM 4.0 production spec
- **LP-010-quasar-stm-4** — QuasarSTM 4.0 paper (this announcement)
- **LP-010** — QuasarSTM 3.0 substrate (2025-12-25)
- **LP-020** — Quasar Consensus 3.0 (cert lanes, P/Q/Z)
- **LP-132** — QuasarGPU Execution Adapter (16 services post-2026-02-14)
- **LP-009** — GPU-Native EVM
- **LP-070** — ML-DSA
- **LP-073** — Pulsar
- **LP-075** — BLS

For questions: research@lux.network For validator support: ops@lux.network For application integration: dev@lux.network

---