



CONFIDENTIAL • COMPUTE SETTLEMENT ONE-PAGER • SEPTEMBER 2026

Settlement infrastructure for the Spot Compute Market.

The delivery leg of the GPU trade, enforced in code.

Company: HashLock Corp. Jurisdiction: Delaware, USA File #10561489 Stage: pre-seed • \$1M

The gap

Compute is financializing fast. CME and Silicon Data list H100 and B200 rental index futures on NYMEX on 5 October 2026, pending regulatory review. ICE has announced its own GPU futures. Rental indices are live on Bloomberg.

All of it is cash-settled against an index. The physical leg has no plumbing. There is no neutral way to prove that GPU-hours were actually delivered to spec. Buyers prepay by the hour and eat idle capacity. Suppliers invoice on their own telemetry and ask to be believed. Every futures position that eventually wants to convert into real capacity runs into the same missing piece: a measurement and settlement layer neither side controls.

What we do

HashLock is building the delivery-versus-payment layer for compute. Buyer funds are locked in an HTLC smart contract — that rail is live on mainnet today. Independent metering will measure what was actually delivered against the agreed SLA, so that funds release on verified delivery and refund on non-delivery. Metered pay-as-you-go rather than prepaid blocks. Payment is enforced in code, not in an escrow agreement. Neither party can take the other's side of the trade.

Live today

- Atomic HTLC settlement live on mainnet across six chains: Bitcoin, Ethereum, Base, Robinhood Chain, Tron and Solana
- MCP server, MIT-licensed — 24 hosted settlement tools and 17 local ones, so AI agents transact programmatically
- Production codebase • 2 internal audit rounds complete
- Mainnet is live; the external audit with QuillAudits is complete (October 2026)

Planned (in build)

- Intent routing designed into the settlement flow; the solver network ships in stages on the HTLC rail that is live today, ahead of V2
- Neutral metering collector: provider-reported usage + marketplace-side GPU telemetry, with variance measured between the two
- SLA evaluation service: uptime, GPU identity and consumed hours, producing signed, machine-verifiable delivery evidence
- Settlement adapter: delivery evidence wired to HTLC release and refund triggers
- Provider adapter SDK + conformance suite — adding a supplier becomes a spec to comply with, not bespoke integration

Engineering partnership

AWS is backing the missing piece of our stack. Through GoML — an AWS Advanced Partner — AWS is committing in-kind engineering support of up to **\$250K of dedicated engineering resources** under a bespoke

SOW covering the metering and delivery-evidence components above — the metering collector, the SLA evaluation service and the provider adapter SDK. **The settlement adapter is outside the SOW** and is built in-house on the audited HTLC core. This is not a stock AWS Activate credit voucher; it is real engineering hours against defined deliverables, allocated because AWS validated the settlement thesis. HashLock retains full code and IP ownership, gives no equity, and takes no revenue share. GoML is not rebuilding our market core, which already runs — the entire engineering budget goes into the one component the market is missing.

Go-to-market

Two-sided, starting from the supply side. Neoclouds and GPU operators carry unsold inventory and are financed against contracted revenue — structurally short and need credible forward sales. Metered settlement monetises idle capacity without prepayment friction; verified delivery is what makes their contracts financeable.

On the demand side: AI teams and enterprises running post-trained open models, whose largest COGS line is compute and who cannot commit to 12-month lock-ins. Pilot discussions with GPU providers in progress — none signed yet. Second vertical: institutional crypto OTC, already runs on the same settlement rail and is the nearer-term revenue path.

Team

Barış Sözen • Co-Founder & CEO. Institutional treasury and FX across Akbank (1997—2011) and UBS AG Istanbul (2007—2008) (FX, rates, digital assets). Building settlement and MEV infrastructure since 2022. Owns buyer intuition, GTM, and the fundraise.

Aleksei Diakonov • Co-Founder & CTO. Blockchain / Protocol Engineer & technical co-founder in DeFi & EVM smart contract systems. Deep expertise in atomic cross-chain settlement + AMM math, precision-safe arithmetic, cross-chain messaging protocols. Owns the atomic-settlement stack, six chains, and the audit lane.

Sizing

Floor (DePIN cohort). Akash spent \$6.4M on H1 2025 (Messari). Four comparable neoclouds at similar volume $\approx$ ~\$50M/yr in verifiable compute-settlement demand today. 8 bps blended = starting revenue base — only three provider pilots to hit the floor. HashLock's edge over these: the Spot Compute Market settles atomically on the HashLock HTLC rail — no platform-custodied escrow.

At scale (CME/ICE physical delivery). GPU rental is on a \$50—80B run-rate for 2026 (Silicon Data index $\times$ CME open-interest proxy). 10% share $\times$ 5 bps = \$2.5—4M ARR at scale. Physical delivery of the Oct 5 futures unlocks the rest.

Raise & ask

Raising \$1M pre-seed at a \$10M post-money valuation cap (YC Post-Money SAFE). External audit with QuillAudits complete October 2026 — funded independently, not from this raise.

The ask: a conversation. If the delivery leg of the Spot Compute Market is a thesis you are tracking, we would like to show you the settlement architecture and where we think the standard gets set.