

COMMON QUESTIONS

Investor FAQ

The questions that come up in diligence, answered with the mechanism rather than the claim. Where something is designed but not built, it says so. Where a number is a model output rather than a result, it says that too.

- | | |
|---|-------------------------------|
| 1 The protocol — and why it is built this way | 6 Security and audit |
| 2 How a desk actually reaches it | 7 Regulatory posture |
| 3 The OTC business | 8 The round and the model |
| 4 Agents | 9 What would have to go wrong |
| 5 Compute | |

Hashlock Markets • File #10561489 October 2026 Stage: pre-seed

1. The protocol — and why it is built this way

What is Hashlock, in one line?

Hashlock is trustless atomic OTC settlement infrastructure. **Every trade either completes in full or fully refunds** — there is no counterparty, escrow agent or bridge to trust. The settlement guarantee is enforced by cryptography — paired hash time-locked contracts — not by a person, a balance sheet or a validator set.

What problem is actually being solved? Desks already trade.

They trade fine. What they cannot do is settle without trusting someone. Two counterparties who have agreed a price face a question with no good answer: who sends first. The three answers in use today are all bad. One side sends first and carries the other's credit for minutes or hours. A human escrow agent holds both sides, charges a spread that is rarely itemised, and becomes a single point of failure. Or the trade crosses chains through a bridge, which means a wrapped asset and a validator set in the path. **The cost of that third answer alone is documented at about \$2B in bridge hacks; the cost of trusting an intermediary with the assets is documented at about \$12.9B since 2022** — FTX, Genesis, Bybit. Hashlock removes the question instead of answering it.

Why hash time-locked contracts and not something newer?

Because the guarantee has to survive the counterparty being hostile, the network being congested, and Hashlock itself being absent. An HTLC gives exactly that and nothing more: a

secret that unlocks both legs and a deadline that returns them. There is no committee, no oracle, no attestation and no sequencer whose failure changes the outcome. Newer designs buy speed by reintroducing someone to trust — a relayer set, a bonded prover, a bridge validator. [For a settlement rail the conservative primitive is the feature](#), and it is the reason a desk can reason about the worst case without reading the code.

Walk through a trade. What actually happens?

Two parties agree a price. Then:

1. The initiator generates a **32-byte secret on their own machine** and sends only its SHA-256 hash to Hashlock. The secret never leaves that machine and Hashlock never holds it.
2. The initiator funds the **long leg** — the leg with the later deadline — into an HTLC locked to that hash.
3. The counterparty sees the funded leg on-chain and funds the **short leg**, locked to the same hash.
4. The initiator claims the leg they are receiving. [Claiming requires publishing the secret, so the act of taking their side of the trade is what makes the other side claimable.](#)
5. The counterparty — or the keeper on their behalf — reads the now-public secret from the chain and claims the other leg. The trade is settled.

If step 3 never happens, the initiator refunds at their deadline. If step 4 never happens, both refund at their own deadlines. There is no state in which one party holds both assets.

Why are the deadlines different on the two legs?

Because one party reveals a secret and the other learns it, and that asymmetry has to be paid for in time. The revealer must have a comfortable window left after revealing; if both legs expired together, a counterparty could wait until the last minute, force the reveal, and then race to claim while the initiator's own refund is already live. So the long leg always carries the later deadline, and the gap between the two legs is never less than **two hours**.

Chain	Leg timelock	Why
Bitcoin	24 hours	Block times are long and variable; the window has to survive a slow stretch without putting anyone at risk
Ethereum • Base • Robinhood Chain • Tron • Solana	4 hours	Fast, predictable finality

The long leg takes whichever is greater: its own chain's timelock, or the short leg plus two hours. On an ETH↔BTC trade that makes Bitcoin the long leg at 24 hours and the EVM leg the short one at 4. Each contract additionally enforces a **30-minute floor** on its own timelock, so no leg can be created with a window too short to act in.

So how long does a trade take in practice?

Deliberately not instant, and the plan does not pretend otherwise. A Bitcoin leg counts as funded after **three confirmations**, roughly thirty minutes, and the two-hour minimum gap is a safety property rather than latency to be optimised away. **What a desk gets is not speed — it is the removal of the window in which someone could run off with both sides.** For an OTC block that is the trade they want; for a retail swap it would not be, which is why the solver network in the roadmap exists as a separate product rather than as a change to this one.

What is deployed, chain by chain?

Chain	What is deployed	Note
Bitcoin	No contract — a P2WSH HTLC script per leg	3 confirmations
Ethereum	HTLCFactory 0xa9B6a64837259034E949787B5d17Bfa46304ff71	Chain id 1
Base	The same factory address	Chain id 8453
Robinhood Chain	The same factory address	Chain id 4663
Tron	SharedHTLC pool contract	TRC-20, incl. USDT-Tron
Solana	HTLC program, one PDA per escrow	SPL

On the EVM chains each escrow is a minimal-proxy clone from a single factory, which is why the same address works on Ethereum, Base and Robinhood Chain and why deploying an escrow costs a fraction of a full contract deployment.

What is the keeper, and does it weaken the guarantee?

The keeper is a service that submits claim and refund transactions on behalf of users on the EVM chains, Tron and Solana, so a desk does not have to sit watching a timelock. **It can only perform moves the contract would already allow** — it cannot redirect funds, cannot change a recipient and cannot claim without the secret the chain has already made public. If the keeper is down, any party can submit the same transaction themselves.

On **Bitcoin** it **cannot act at all**, because spending a P2WSH output requires the owner's own signature. The Bitcoin leg is signed by the user. This is stated plainly because it matters for any fully automated flow.

Can Hashlock stop, reverse or seize a trade in flight?

No, and that is structural rather than a policy. Once a leg is funded, the only transitions the contract permits are claim-with-the-secret and refund-after-the-deadline. There is no admin key

that moves funds, no pause that redirects them and no discretionary intervention. The nine states a swap can occupy — `agreed`, `initiator_funded`, `counterparty_funded`, `initiator_claimed`, `counterparty_claimed`, `refunding`, `refunded`, `expired`, `failed` — contain no branch in which Hashlock decides an outcome.

What exactly does not come back on a refund?

Two things, and a desk should hear both before the pilot rather than after. The **protocol fee**, which the taker pays alongside their funding, is not refunded if the deal fails — it pays for the settlement attempt, not the settlement. And on **Solana**, the funder's account rent stays with the escrow account, because those accounts are never closed. Everything else returns to the wallet it came from.

How is the fee actually charged?

The taker pays it, on top of the leg they fund, in that leg's own asset. There is no invoice, no separate settlement in a different token, and no balance held anywhere. On a Bitcoin leg the fee is paid on the Bitcoin side, with a dust waiver below 546 satoshis so the protocol never creates an unspendable output.

Does the trade have to cross chains?

No. Same-chain settlement — $\text{ETH} \leftrightarrow \text{ETH}$, for instance — runs under the same atomic guarantee, and it is a real use case: two desks on the same chain still face the who-sends-first problem. The one restriction is that **TRON $\leftrightarrow$ TRON is not offered**; one side of a Tron pair must sit on another chain.

Which assets can be settled?

Native assets and standard tokens on each supported chain — BTC on Bitcoin, ETH and ERC-20s on the EVM chains, TRC-20 including USDT-Tron, SPL on Solana. Tokens launched on **Bankr**, **Virtuals** and **Clanker** clear the checks automatically because the launchpad already enforces them; any other token is submitted by contract address and listed only if it passes the same safety and liquidity checks. Nothing is listed without passing them. The practical significance is that **78% of OTC flow is now stablecoins**, which are natively settleable on this rail with no conversion step.

2. How a desk actually reaches it

What are the ways in?

Three, and they are the same protocol underneath — not three products.

- **The web application.** Request, quote, negotiate, lock, claim, with the settlement state machine visible throughout.
- **The `/v1` REST API**, with webhooks for state changes and a maker WebSocket feed for desks that want to quote programmatically.

- The MCP server at <https://hashlock.markets/mcp>, or locally over stdio, so an agent runtime can call the same primitive.

Where does Telegram fit? Earlier material called it the main interface.

It is a **negotiation channel**, not a settlement surface, and the earlier framing overstated it. Desks genuinely do agree prices in chat, and meeting them there removes a real adoption barrier — but the funds move through the contracts, reached by one of the three interfaces above. The distinction matters in diligence: **no part of the settlement guarantee depends on Telegram**, and a desk that never touches it loses nothing.

What does a desk have to change to use this?

Less than the pitch for most infrastructure. They keep negotiating where they negotiate. They keep their own wallets — Hashlock never takes custody, and the Portfolio view reads balances from the chain rather than holding them. What changes is the settlement step: instead of wiring to an escrow agent and waiting, they lock into a contract and claim. The switching cost is a workflow habit, not an integration project, which is why the pilot is scoped at **ten settlements over four weeks** rather than as a build.

3. The OTC business

Who are the first customers, and why them?

The beachhead is MENA crypto-OTC, and the primary wedge is **Bitcoin mining pools**. Three reasons, in order of how much they matter. They sell mined BTC continuously, so the counterparty question recurs weekly rather than occasionally. Their flow is natively cross-chain — BTC out, stablecoin in — which is the exact trade this rail was built for. And one pool relationship reaches many miners, so the acquisition arithmetic works for a founder-led motion. From there: exchange OTC desks, crypto hedge funds and Gulf family offices, on the same product with no change of motion.

Why MENA?

Because the flow is there and the incumbent trust model is weakest there. The UAE alone received **\$56B+ in on-chain value in 2024–25**, up 33%, driven by large institutional transactions (Chainalysis). Gulf desks clear size bilaterally, often through relationships rather than regulated venues, which is precisely the setting where a human escrow agent is the norm and where removing him is worth paying for.

Why would a desk switch? Their escrow agent has never failed them.

That is the honest objection and it is the one the sales conversation is built around. The answer is not “your agent will fail”. It is the discovery question: what does the current arrangement cost you, in the spread your escrow agent takes and in the days your capital sits idle waiting? **That cost is almost never itemised on the desk's own side** — it is embedded in a quote. Against it

sits a published fee under 10 bps and a settlement that cannot half-complete. The pitch is a line item replacing an invisible one, with the tail risk removed as a side effect.

How is the fee structured?

A single volume-tiered settlement fee on notional — sub-10 bps, 8 bps today — stepping down as a counterparty's settled volume rises.

Tier	Rate	Audience
Entry	~8—10 bps	Early pilots, smaller desks
Desk scale	~5—6 bps	Active desks and mining pools settling weekly
Institutional floor	3 bps	The lowest rate offered — the schedule does not go below it

The financial plan is modelled at a flat 8 bps blended rate for the whole window; the volumes that would reach the 3 bps floor sit far beyond this plan, so the tiering is upside to the customer and neutral to the model. Exact schedule shared in diligence.

One revenue line sounds fragile. Why not more?

Because every additional line would require something the model does not have — custody, inventory, a balance sheet, or a licence. A fee on settled notional is the only mechanic that scales with usage and costs nothing to carry. **No custody means no capital requirement, no inventory risk and structurally high gross margin**; it is also what keeps the regulatory surface narrow. The discipline is deliberate: one protocol, one mechanic, one number to grow.

What is the moat?

Not one feature — the set a competitor has to assemble at once. Native dual-HTLC settlement across six chains with no bridge in the path; a real Bitcoin leg rather than a wrapped derivative; the audit posture; and distribution that meets desks inside the tools they already use, with the same primitive exposed to agents. Each piece is replicable alone. **Assembled, audited and live, the combination is not forkable in under 18 months** — and the cost of being wrong about that is borne by whoever ships a settlement rail without an audit.

Who is the competition, really?

The status quo, in three forms: bilateral trust where one side simply sends first, human escrow agents and guarantors operating out of Telegram OTC groups, and bridges plus custodial escrow on the cross-chain side. Bilateral-trust OTC desks — Cumberland, Galaxy and their peers — are incumbents in the flow, but they are counterparties rather than settlement rails; the rail under their client trades is exactly the position Hashlock is built for.

How many customers does the plan actually need?

Sixty-four active settling clients and \$265M in monthly settled volume at Month 18. The plan buys them rather than courting them one at a time: **\$650,000 — 65% of the round — is a**

[customer-acquisition budget](#), deployed over twelve months from January 2027 at an assumed \$1,000 per qualified prospect and a 10% conversion, which is a \$10,000 paid CAC against a \$64.8K first-year LTV. Roughly 86% of the year-one cohort is bought; the rest comes from warm intros. Neither of those two assumptions has been measured yet, which is the risk named in §9.

4. Agents

Is the agent story live, or is it a slide?

Partly live, and the distinction is worth being exact about. **Shipped:** the MCP server in production, the `/v1` API with webhooks and a maker feed, scoped API keys, and published SDKs (`@hashlock-tech/sdk` on npm, with a Python package on PyPI) — an agent can settle on Hashlock today. **Not shipped, and not claimed:** agent-driven revenue. No figure in this plan depends on it.

Why do agents need this particular primitive?

Two properties that no other settlement arrangement offers them. First, [an agent cannot instruct a human escrow agent](#) — the incumbent trust model has a person in it, and a person is not an API. Second, an agent cannot safely accept a trade that half-completes, because recovering from a partial failure requires judgement, negotiation and often a phone call. An atomic complete-or-refund outcome is retry-safe: the failure mode is a clean refund the agent can simply attempt again. Those two properties are not conveniences for a machine counterparty; they are preconditions.

Why does MCP specifically matter?

Because it removes the integration step. An agent runtime adds the Hashlock MCP server as a connector, signs in, and the settlement primitive is callable with no custom adapter to maintain and no bespoke integration per framework. A desk that would rather write code than wire an agent takes the SDK instead. [Distribution on this surface costs nothing per additional user](#), which is why it is pursued now even though it carries no revenue in this round: directory listings, framework documentation and the docs published as a single file an LLM can read.

If the market is \$1.7T by 2030, why is it not the plan?

Because the projection is Citi's and the revenue would be ours to prove, and those are different things. Agent-to-agent commerce at that scale has no settlement incumbent, which is the opportunity — but it also has no desks writing cheques in 2027, which is the problem. **Institutional OTC funds the rail; agents extend it without a rebuild.** An investor should underwrite this round on the OTC numbers and treat the agent surface as a free option that is already partly exercised.

What would have to be true for agents to become the main line?

Three things, none of which Hashlock controls: agents holding and spending value in their own right at scale rather than executing on a human's behalf; a settlement expectation that tolerates the two-hour cross-chain window, or enough same-chain and fast-chain flow that it rarely binds; and at least one framework treating payment as infrastructure rather than as a plugin. [We would rather name those conditions than forecast through them.](#)

Is there an agent limitation worth knowing about?

Yes: the Bitcoin leg. Because spending a P2WSH output needs the owner's signature, a fully unattended agent flow cannot sign a Bitcoin claim the way the keeper does on the other chains. For agent traffic the EVM, Tron and Solana rails are the automatable ones, and Base and Robinhood Chain — where the AI-agent platforms already operate — are the first targets in the solver roadmap for exactly that reason.

5. Compute

What is the compute vertical?

A second market for the same settlement problem. GPU capacity is bought and sold forward and spot, and the buyer faces the mirror image of the OTC question: pay first and hope the capacity is delivered, or deliver first and hope to be paid. The **Compute Power Exchange** applies the settled-or-refunded primitive to that trade. Its architecture and specification are **complete**; the settlement core is **in design**.

So is it a product or a plan?

A plan with funded engineering behind it, and the room should be read that way. **Signed**: an engineering SOW with GoML / AWS worth \$250K of dedicated hours against defined deliverables, with Hashlock retaining code and IP ownership. **In progress**: a GPU telemetry module under that SOW. **Not built**: the exchange itself, which is scheduled into V2 in 2027 and gated on an external audit before any mainnet deploy.

Why would it be cheap to build, given the OTC rail exists?

Because the audited HTLC core is reused unchanged. [The only new audit surface is the delivery-proof stack](#) — the part that establishes whether the compute was actually delivered — which is why the compute module enters the V2 audit scope rather than requiring a separate protocol. The settlement half of the problem is already solved and already reviewed.

Who would the counterparties be?

Decentralised GPU networks and compute marketplaces — Akash, Aethir and io.net are the named integration targets. To be exact about status: **these are prospective and none is engaged**. They appear in the plan as the shape of the first cohort, not as a pipeline.

Why is compute kept in a separate room?

Because it is a different underwriting decision and mixing them does a disservice to both. The OTC room asks an investor to back a live rail with a fee already charged. The compute room asks them to back a V2 build against a signed engineering contract. An investor should be able to say yes to one and not the other, which is only possible if the materials are separated — so they are.

And the solver network — where does that sit?

It is the next build on the OTC rail rather than a third vertical, and it launches in three stages, in order. Stage one puts signed orders and an open solver bot on **Base and Robinhood Chain**, where the AI-agent platforms operate. Stage two extends to Bitcoin, Tron and Solana over the HTLC settlement that is live today, with an auction and solver deposits. Stage three adds instant solver payment on EVM-to-EVM routes, with optimistic repayment settled against a proof and a challenge bond. Every stage is reviewed before mainnet, and any contract change passes an external audit scoped and funded at the time. **The OTC desk is unchanged by all of it** — the solver network serves ordinary swaps where a user wants a price in seconds, not a negotiation.

6. Security and audit

Is the code audited?

Two internal audit rounds are complete, layered with a five-tool hardening stack — Slither, Halmos, Echidna, Stryker and fast-check — and a runtime invariant monitor. The **external audit with QuillAudits is complete, October 2026**, and independent sign-off is in hand. Total external audit spend is \$10K, with a further \$10K committed to a bug bounty pool — \$20K of security spend, already incurred.

Why does the audit matter more here than for most products?

Because the contracts are **immutable once deployed**. There is no upgrade path on V1 by design — an upgradeable settlement contract has an admin key, and an admin key is the thing the product exists to remove. The consequence is that improvements ship to the next contract generation rather than to the deployed one, and that V2's mainnet deploy is gated on external audit-firm sign-off before anything moves.

What is the honest residual risk?

Smart-contract risk never reaches zero, and the plan does not claim it. What is true is that the blast radius is bounded by design: no pooled treasury to drain, no custody balance, and each escrow holding only its own trade. **A failure costs a trade, not a book.**

7. Regulatory posture

What is the regulatory position?

Non-custodial by construction, and that does most of the work. Hashlock never holds customer funds; the escrow pays only the leg's fixed recipient or refunds the funder; timelock refunds are unilateral and need no one's permission. [There is no client money to safeguard](#), which is what narrows the surface — money-transmission and custody regimes largely attach to holding other people's assets. Full reasoning is in the Regulatory Strategy Memo.

What is deferred, and why is that acceptable?

Dubai business development, a full regulatory build-out, a Cayman foundation (only relevant if a token appears) and money-transmitter analysis (only relevant if the model ever becomes custodial) are all deferred to the next priced round. Each is conditional on something that has not happened. Spending against those conditions now would be spending against a hypothesis.

Who screens the counterparty?

Partly Hashlock, contractually; the rest sits with the desk. The Terms of Use require every user to represent that they are not a sanctioned person — OFAC SDN, UN, EU and UK HMT — and not located in a sanctioned jurisdiction, and they reserve Hashlock's right to screen and block addresses. [There is no automated screening inside the settlement flow, and this FAQ will not claim one](#). So the control is a representation the counterparty makes and a right Hashlock holds, not a gate a transaction passes through.

Beyond that, a desk using Hashlock remains responsible for its own counterparty diligence, exactly as it is when it settles bilaterally today. KYC/KYT tooling is in the plan from January 2028; automated sanctions screening is a next-round decision, and the point at which it becomes necessary rather than optional is a judgement the next round funds.

Is this a securities matter in any form?

No. There is no token in this plan, no yield, no pooled investment and no promise of return from the efforts of others attached to the protocol. The round itself is a YC Post-Money SAFE — a securities instrument, sold as one, to investors qualified as such.

8. The round and the model

What is Hashlock raising, and for what?

A \$1M pre-seed at a \$10M post-money cap (YC Post-Money SAFE), single SAFE close January 2027 via Clerky. The \$1M funds an 18-month plan: acquire the first cohort of paying desks, operate the protocol live on mainnet, land and prove the first cohort of paying desks, and build a data room for the next priced round with a target close before Month 24. The Use of Funds document is canonical.

What is not funded by this \$1M?

Dubai BD, live KYC onboarding gates, a Cayman foundation and MTL analysis are sequenced to the next round, as is every salaried hire except one — BD, compliance, DevRel and the solver researcher all wait. Funded inside this round: both co-founders (CEO \$48K/yr, CTO contractor \$90K/yr), one protocol engineer at \$144K/yr from January 2028, D&O insurance from January 2027, KYC/KYT tooling from January 2028, and the full regulatory memo at \$8K in January 2027. [The round is sized to buy customers for a rail that already works, not to scale headcount.](#)

What does the cash trajectory look like?

Plan year	Net revenue	Spend	EBITDA	Customers (end)	Settled volume	Cash (end)
Y1 • Sep 2026 — Aug 2027	\$306K	\$655K	— \$349K	46	\$479M	\$724K
Y2 • Sep 2027 — Aug 2028	\$2.12M	\$607K	+ \$1.51M	62	\$3.31B	\$2.23M

Cash at Month 18 (Feb 2028) is **\$1.14M** — more than the raise, against \$1.02M of cumulative spend — with \$265M of monthly settled volume, 64 paying customers and roughly \$2.03M of ARR run-rate, net of the 20% referral rebate. [EBITDA turns positive in month 12.](#)

Why does a live fee sit beside a model that starts revenue in January 2027?

Because the model has not been re-dated since mainnet went live, and saying so is more useful than quietly reconciling it. The protocol charges its fee today; the financial plan's revenue curve still begins in January 2027 and runs Year 1 from September 2026. [Treat the model's dates as plan checkpoints rather than calendar dates](#) until it is restated. The figures themselves are unchanged and the restatement moves timing, not totals.

Is there any revenue today?

The fee mechanism is live and charged on settled notional. There is no cohort of paying desks yet — that is what this round funds, and the go-to-market document states plainly that the customer lane opened when the external audit completed in October. An investor should underwrite the first cohort as work to be done, not as work already done.

Who is the team?

Two co-founders throughout, plus one protocol engineer from January 2028 and commission-based sales — headcount 2 → 3 inside the window. **Barış Sözen** (Co-Founder & CEO, 82.5%) — institutional treasury at Akbank 1997–2011, Dealer to Position Manager in Treasury and FX, with a UBS AG Istanbul fixed-income and derivatives interlude 2007–2008; independent algorithmic trading since 2016; low-latency MEV and statistical-arbitrage systems since 2022.

Aleksei Diakonov (Co-Founder & CTO, 17.5%) — protocol engineer, author of the atomic cross-chain settlement stack, owner of all six chains and the audit lane. See the Founder Bio and Hiring Plan.

What partnerships are signed?

Two. **Motion Trade** — a strategic advisory and BD mandate, executed 10 September 2026, covering counterparty and investor outreach. **GoML / AWS** — the \$250K engineering SOW described in §5. Both are signed; neither is revenue.

9. What would have to go wrong

The risks an investor should price, stated before they are asked about.

The customer lane starts from zero

Outreach to settlement counterparties was gated on the external audit and the audit completed in October. There is a signed BD mandate and a defined funnel, but **no desk is in pilot today**. If the first twenty conversations convert materially worse than the assumed rates, the Month 18 figure moves and the next round moves with it. This is the single largest open item in the plan.

Acquisition efficiency is the constraint, and it is unmeasured

Two thirds of the round is a customer-acquisition budget resting on two numbers that have never been tested: **\$1,000 per qualified prospect and 10% conversion**. At twice the cost per prospect, or half the conversion, the \$650K buys 33 desks instead of 65 and every Month 18 figure halves with it. **The company stays solvent in that case — the cash trajectory through month 12 is dominated by the raise, not by revenue — but the seed-ready story does not survive it.**

The mitigation is sequencing rather than optimism: the first \$50K of spend measures both numbers, the model is re-run on realized values, and only then is the remaining \$600K committed. An investor should expect to see those measured numbers before the budget is fully deployed, and should treat their absence as the thing to ask about.

Key-person concentration on the engineering side

One CTO owns the settlement stack across all six chains and the audit lane. The mitigating facts are real — the contracts are deployed, immutable and externally audited, so the system does not depend on continued development to keep working — but the next generation of contracts does depend on him.

The two-hour window is a product constraint

It is a safety property and it will not be removed. It narrows the addressable set to trades where settlement certainty matters more than settlement speed. That is the OTC block market by definition, and it is not the retail swap market — which is why the solver network is a separate product rather than a tuning of this one.

Sanctions screening is contractual, not automated

The Terms of Use carry the representation and the reserved right to block, but **no automated screening runs inside the settlement flow**. For a non-custodial rail that is a defensible position, and it is the same position the desk occupies when it settles bilaterally — but an institutional counterparty will ask, and the honest answer is that the control is contractual today and automated screening is a next-round decision.

Model timing is unreconciled

As set out above, the financial model has not been re-dated against a live fee. Totals are unaffected; the month-by-month cash trajectory is not. It is on the list to restate.

Canonical sources for anything above: the Use of Funds and Financial Plan for every figure, the Technical Whitepaper for protocol mechanics, the Go-to-Market Strategy for the funnel, and the Regulatory Strategy Memo for the compliance reasoning.