



Market Sizing

Compute for Equity — investor memo

Status: strategy memo. Figures are illustrative estimates triangulated from public sources; they are directional, not audited. Prepared for the data room, June 2026.

1. The market we are actually in

It is tempting to size "the GPU market" or "the AI market" and wave at a trillion-dollar number.

We will not do that, because it would not be honest about what we monetize. Compute for Equity is a

clearing layer. We do not sell GPUs and we do not invest a fund. We take a thin fee on the *flow* of value as compute and energy are exchanged for equity. So the right way to size us is to

size the **flow we can clear**, then apply a realistic take rate. We build the number from the bottom up and sanity-check it from the top down.

The three assets we make interoperable are each large and growing, but the relevant market is their

intersection: the subset of compute spend that an AI company would rather settle in equity than in

cash, brokered through a regulated venue. That intersection is small today because the instrument to

clear it barely exists — which is precisely the opportunity.

2. Top-down anchors (the context)

A few public reference points frame the opportunity:

▶ **AI infrastructure spend** is on a multi-hundred-billion-dollar annual trajectory, with leading

estimates placing the cumulative AI infrastructure financing gap in the hundreds of billions —

the figure we surface on the site, **~US\$490B**, sits in that range. Whatever the exact number, the direction is not in dispute: compute is the dominant cost of building AI, and it is being financed with increasingly exotic structures.

▶ **The GPU rental / neocloud market** (CoreWeave, Lambda, Together, and the hyperscalers' rental

lines) is already tens of billions of dollars annually and projected to grow toward the hundreds

of billions by the early 2030s.

▶ **GPU-backed debt** has gone from zero to a recognized asset class in under three years: CoreWeave

alone has raised facilities of \$2.3B, \$7.5B and \$8.5B collateralized by GPUs and their customer

contracts — the last achieving an investment-grade rating. The financialization of compute is

already happening; it is simply happening as *debt between two parties*, not as a *cleared, equity-linked, multi-sided market*.

▶ **Venture-as-compute** is live too: a16z's Oxygen program rents GPUs to portfolio companies below

market in exchange for equity, with a stash reported in the tens of thousands of chips. The behaviour we intermediate — compute traded for equity — is already happening bilaterally inside

single funds.

The conclusion from the top-down view: every adjacent structure (GPU debt, VC-as-compute, GPU

rental) validates demand for exchanging compute and capital — and every one of them is *bilateral*.

None is a neutral, regulated, multi-sided clearing venue. That gap is the white space.

3. Bottom-up TAM / SAM / SOM

We size the cleared flow directly.

TAM — total compute-for-equity flow that could exist.

Take global AI training-and-inference compute spend by venture-funded AI companies. A meaningful

share of that spend is, in principle, "equity-fundable" — i.e., the company would prefer to pay in

equity rather than burn cash, particularly for large, lumpy training runs. If global

venture-AI compute spend is on the order of **US\$80–120B/yr** and even **15–25%** of it is

equity-fundable, the addressable *flow* is roughly **US\$12–30B/yr** of compute that could be

settled as equity. That is the TAM for the value we clear — not the GPU market, the *convertible

slice* of it. As AI compute spend grows toward the hundreds of billions, this slice grows with it.

SAM — the slice we can serve given regulation and geography.

We serve professional clients, through an ADGM-regulated venue, anchored on energy-rich Gulf supply.

Realistically the early serviceable market is the flow that (a) involves a participant willing to transact through a regulated venue, and (b) touches the supply we can originate. We estimate the

serviceable flow at roughly **US\$2–5B/yr** within the first several years — the regulated, relationship-reachable subset of the convertible slice.

SOM — what we actually clear, near term.

Phase 1 is ^{80k}10 hand-brokered deals. At an illustrative ^{1.5M}**US\$1.5M per Compute** [Apply as a startup →](#)

\\$10–50M of cleared flow in the first phase — tiny as a market share, decisive as a proof.
As

Phase 2 opens the platform to a vetted cohort, we model cleared flow growing toward
\\$200–500M/yr,

i.e. roughly **5–15%** of the SAM. SOM is deliberately conservative because liquidity is hand-built,

not assumed.

4. Revenue model on the flow

We are capital-light by design; revenue scales with cleared volume, not with our own balance sheet.

The fee stack, illustratively:

- ▶ **Origination fee:** ~1–2% of the Compute-SAFE principal at structuring.
- ▶ **Clearing & settlement fee:** ~0.5–1% on settlement.
- ▶ **Custody fee:** a thin annual basis-point charge on assets under custody.
- ▶ **Secondary-market take rate (Phase 3):** ~0.25–0.5% per transaction once a secondary book exists.
- ▶ **Treasury float:** modest yield on segregated balances under conservative policy.

Blended, this lands around a **2–4% effective take rate on cleared primary flow**, plus recurring

custody and secondary revenue that compounds as the book grows. Apply that to the SOM:

- ▶ **Phase 1** (~\\$10–50M cleared): roughly **\\$0.3–1.5M** of fee revenue — enough to validate unit

economics, not to be a business yet.

- ▶ **Phase 2** (~\\$200–500M cleared): roughly **\\$6–18M/yr** — a real, capital-light revenue line.
- ▶ **Phase 3** (SAM penetration deepening + secondary): the model scales toward

\\$30–60M/yr as cleared flow approaches the upper SAM and secondary volume turns over the book

multiple times.

I am a

[Apply as a startup →](#)

These are illustrative and tied to the financial model in the data room (`model.json`), which carries the Bear/Base/Bull build. The point is structural, not the decimal: **a 2–4% take on a multi-billion-dollar cleared flow is a venture-scale business, achieved without us ever owning a GPU or raising a fund.**

5. Why the take rate is defensible

Thin take rates are only defensible when the venue provides something the parties cannot easily

replicate bilaterally. We provide four: a **neutral price** (we own no compute, so our Oracle quote

is trustable), a **standard instrument** (the Compute-SAFE, so deals are repeatable and tradable),

regulated custody (so institutional and sovereign capital can participate at all), and

liquidity (a secondary market no single counterparty pair can offer). Each justifies a slice of the take. Strip any one of them and you are back to a private bilateral deal — which is exactly the

status quo we improve on.

6. Sensitivities and what would change the number

▶ **AI compute spend trajectory.** If AI infrastructure spend grows faster than consensus (plausible

given current capex), every tier scales up. If it stalls, the convertible slice shrinks but does not vanish — dilution-sensitivity is structural to startups in any market.

▶ **Equity-fundable share.** Our 15–25% assumption is the swing variable. The a16z Oxygen precedent

suggests real appetite; wider adoption pushes the share up.

▶ **Regulatory timeline.** A faster path to a full FSP pulls Phase 2 revenue forward; a slower one

pushes it out. The RegLab on-ramp caps the downside by letting Phase 1 revenue start regardless.

▶ **Take-rate compression.** As the asset class matures, fees compress — but custody and secondary

revenue (recurring, book-driven) grow to offset primary compression. The mix shifts from transactional to recurring as we scale, which is the healthy direction.

7. Comparable venues and why clearing layers earn premium multiples

The right comparables for valuation are not GPU clouds or VC funds — they are **clearing houses,

exchanges, and marketplace infrastructure.** Public market operators (ICE, CME, Nasdaq) and

financial-infrastructure businesses trade at premium multiples precisely because they are capital-light, take-rate businesses with network effects and regulatory moats — the exact shape we

are building. A GPU cloud is a capital-intensive, depreciating-asset business that trades on revenue multiples weighed down by capex; a fund earns carry but does not compound as infrastructure. A *venue* compounds: each marginal deal costs us almost nothing to clear, and each

adds to the liquidity that attracts the next. That is why a clearing layer with a 2–4% take on a growing multi-billion-dollar flow is worth more, per dollar of revenue, than the GPU business underneath it. We monetize the *flow*, not the *iron* — and the flow is where the durable margin is.

A useful sanity check: GPU-backed debt went from \(\$0 to tens of billions in under three years as a

bilateral market. A *cleared, multi-sided, equity-linked* market for the same underlying — with a

Apply as a startup →

secondary — is a structurally larger and stickier prize, because it serves more participants and trades more than once. The debt market's velocity is the proof of appetite; the clearing layer captures more of it.

8. The honest bottom line

We are not claiming a trillion-dollar TAM. We are claiming something more underwritable: a **multi-billion-dollar annual flow of compute that AI companies would rather settle in equity**, a **2-4% take** on clearing it, and a **regulated, supply-anchored position** that lets us capture a growing share of it without owning the underlying assets. The adjacent markets — GPU debt, VC-as-compute, GPU rental — already prove the behaviour at scale. We are building the neutral venue that turns that behaviour into a cleared market. Start at \ \$10–50M of flow hand-built in Phase 1, compound toward hundreds of millions as the platform opens, and the take rate does the rest.

Human TODOs before external circulation: reconcile every figure here against the latest `model.json`; cite specific third-party market reports for the top-down anchors; have finance sign off on the take-rate assumptions.