

Kaldune



VALUATION DEFENSE · CONFIDENTIAL

The \$12M cap is conservative.

A pre-revenue, product-complete company raising up to \$2M on a SAFE at a \$12M post-money cap — and why, on three independent methods, that entry is fair to favorable.

Kaldune, a product of Studio58, Inc.
Hicham Abdessamad · Founder
Former Chairman & CEO, Hitachi America

MMXXVI
PREPARED FOR PROSPECTIVE
SAFE INVESTORS

A working company, priced below what it would cost to rebuild.

Kaldune is a pre-revenue, product-complete company raising up to \$2M on a SAFE at a \$12M post-money cap. At that cap, the entry is conservative on every reasonable method of valuing it.

What the buyer is buying is not a plan. It is a working system: **445 end-to-end client jobs already run, 409 delivered, 17,267 finished files produced** across at least nine kinds of work, all measured from the repository on 2026-08-20.

The category is the finished-work layer. A plain-language brief goes in; a finished deliverable comes out, produced by an integrated AI team the client watches work, at software economics rather than payroll economics. It sits directly on top of the largest services market in the world: more than \$1.5 trillion a year of spend on finished work delivered by scarce human experts.

The valuation argument rests on three legs, each developed in full in Section 3, each landing independently at or above \$12M.

THE EXPERT POSITION

The \$12M cap is **locked**. Money is already committed against it. This document does not argue for a higher number.

The argument is the opposite and stronger: three independent methods place the honest value of this company at or above \$12M, so a new investor entering at that cap is entering at a fair-to-favorable price set by a disciplined founder.

LEG ONE

COMPARABLE FINANCINGS

The AI builder-tools aisle one layer down is funded from the hundreds of millions to sixty billion dollars. Against that context a \$12M cap on the layer above is a rounding error, and the direct seed comparable clears it.

LEG TWO

BUILD / REPLACEMENT COST

~200,000 lines of shipped code, 79% of it orchestration; six filed U.S. utility patent applications underneath. Replacement cost of the code alone approaches \$12M, before IP, discipline, or demand.

LEG THREE

RISK-ADJUSTED SCORECARD

Two recognized pre-revenue frameworks (Scorecard and Berkus), inputs shown, place the defensible range with \$12M sitting at its conservative end.

THE HONEST CAVEAT, NEVER BURIED

This is a pre-revenue company: no ARR, no proven distribution engine yet. The open risk is go-to-market, precisely what the raise funds (Section 7). A serious investor prices it, and even priced, \$12M is conservative, because the product risk that usually kills a seed company is retired, measured, and filed.

A number sitting below three ceilings.

Before the methods, the sense of scale — because it disciplines everything after it. A conservative cap is not one that reaches upward; it is one that sits well below the market it addresses, the capital in the adjacent layer, and the cost to rebuild what already exists.

0.0008%

of the annual spend of the market it addresses. A \$12M cap is eight ten-thousandths of one percent of a \$1.5T+ services market — a seed entry into a company early to it, not a claim on it.

<1/1000

of Lovable's valuation. Lovable raised a \$400M Series C at \$13.3B one layer down; Kaldune's entire post-money cap is less than one-thousandth of it, on the layer above.

12–18_{mo}

a funded team's honest rebuild-and-harden clock. The fully-loaded cost of a team capable of it, over that window, sits in the same range as — or above — the entire post-money valuation.

AGAINST THE MARKET IT ADDRESSES

The finished-work layer sits on a services market of more than \$1.5 trillion a year. The cap is not a claim on that market; it is a seed entry into a company early to it. (The \$1.5T figure is Kaldune's own constructed rollup of consulting, agency, and freelance markets — framing, not a single-source citation; the tiers are sourced in Section 4.)

AGAINST THE CAPITAL ALREADY IN THE AISLE BELOW

Roughly \$1.5 billion went into the builder layers one floor down in a single week of August 2026. Kaldune occupies the layer above that aisle, positioned to capture the value of the stack beneath it as the token layer commoditizes.

AGAINST REPLACEMENT COST

The build is ~200,000 lines of production TypeScript across the engine, the identity gateway, and the studio client, with six filed U.S. utility patent applications on the non-obvious mechanisms.

A competent funded team's honest rebuild-and-harden timeline is 12–18 months (Section 3B), and the fully-loaded cost of a team capable of it, over that window, sits at or above the entire post-money valuation — before a single line of go-to-market.

THE READING

\$12M is not a number reaching upward from a pre-revenue company's ambitions. It sits well below three independent ceilings: the market it addresses, the capital in the adjacent layer, and the cost to rebuild what exists. That is the definition of a conservative cap, and the through-line of everything below.

Sources. Market tiers and comparable financings verified per Section 4 (2026-08-26); engine and build figures from the repository 2026-08-20. The ~\$1.5B and ~\$1.5T figures are the company's own framing.

03

The three-legged valuation.

No single method values a pre-revenue company well. The discipline is to triangulate: run three independent methods, show the inputs to each, and read where they converge. They converge above \$12M.

3 A

Comparable financings

3 B

Build / replacement cost


3 C

Risk-adjusted scorecard

The layer below is priced in the billions. This cap is a rounding error.

The finished-work layer sits on a hyper-funded aisle. Against that funding context a \$12M cap is a rounding error, while the direct seed-stage comparable set for a product-complete pre-revenue AI company also clears \$12M.

The point of this exhibit is not that Kaldune is worth what Lovable is worth. It is that the capital markets have **already priced enormous value into the layer directly below Kaldune**, which is the strongest available evidence that the demand Kaldune addresses is real and large. Every figure below was verified against its primary or reporting source, read on each source’s own page.



The pattern is the argument. Capital concentrated from the hundreds of millions to sixty billion in the adjacent builder layer; no named funded player in the layer above. Disclosed valuations in navy; undisclosed financings shown at total-raised in grey. Log scale — each gridline a 10× step.

Sources, read on each owner’s or reporting outlet’s own page (2026-08-26). Cursor/Anysphere: SEC Form 8-K (16 Jun 2026), Forbes/Quartz/TechCrunch — SpaceX **agreed to acquire** Anysphere (Cursor’s parent), \$60B all-stock, **signed not closed**, close targeted Q3 2026 (prior priced round \$29.3B, Nov 2025). Lovable: company blog, Series C \$400M at \$13.3B, 12 Aug 2026. Replit: PR Newswire via TechCrunch, Series D \$400M at \$9B, 11 Mar 2026. Bolt/StackBlitz: Series B \$105.5M (Jan 2025) verified; **~\$700M valuation reported, officially undisclosed** — last disclosed mark, may be stale. Lindy: ~\$50M total raised (2023), valuation undisclosed. Relevance AI: \$37M total raised (Series B, May 2025), valuation undisclosed.

Against the seed comp band, \$12M sits in the lower half.

The aisle exhibit proves the demand is real; it does not by itself set Kaldune's seed price, because those are later-stage companies. The direct comparable is the seed and pre-seed market for **product-complete, pre-revenue AI-native companies with a defensible IP position and a technical founder**.

In the 2025–2026 environment, that class of company routinely raises seed rounds on post-money caps well into eight figures, and higher where the product is already shipping and the founder pedigree is unusual. On this comparable, a \$12M cap sits squarely inside, and toward the lower half of, the band for a company that has already shipped **409 delivered jobs** and filed on its mechanisms.

The seed-cap band is stated qualitatively here rather than as a single printed figure: it is the one comparable most exposed to drift, and the funded-aisle pattern carries the leg without it. That is the conservatism, shown against the market's own comparable pricing.

WHAT SETS KALDUNE APART IN THE BAND

Not a prototype: 409 delivered jobs, six filed U.S. utility patent applications, a measured build. Most seed rounds in this band price a promise; this one prices a shipping system.

WHAT LEG ONE CONCLUDES

\$12M is small against the adjacent capital by three orders of magnitude, and mid-to-low against the direct seed comparable band. **Leg one places the floor of a fair value at \$12M, not the ceiling.**

The market bet correctly that AI is enormously valuable, and placed the bet one layer too low. The capital below is proof of the demand; the emptiness above is the room to own it.

THE FUNDED-AISLE READING · SECTION 3A

Sources. Comparable financings verified per Section 4. The direct-seed-comparable band is characterized qualitatively; no unverified single figure is printed for it, per the print-control discipline (Section 4).

What it would cost to rebuild this approaches the whole cap.

What a competent funded team would spend, and the time they would burn, to rebuild and re-harden this system to where it stands today sits in the same range as the entire post-money valuation — before assigning any value to the filed IP or the demonstrated demand.

WHAT EXISTS, MEASURED FROM THE REPOSITORY 2026-08-20

Shipped production TypeScript	~200,000 lines
— engine / identity gateway / studio client	114,024 / 17,236 / 68,169
Orchestration vs. integration wiring (engine)	79% / 21%
Equipment adapters	54
Money-engine modules	37
Engine / orchestration modules	103
Verification modules / named gates	19 / 27
Test files (named suites)	545 (366)
Commits over 49 active days	2,099 (~43/day)

THE REPLACEMENT-COST ARITHMETIC, SHOWN

The rebuild-and-harden window for a funded team is **12 to 18 months**, and the clock does not start at day one — it starts when the team's own first hundred real briefs begin breaking things already fixed here.

A team capable of building this is not one engineer; it is a small senior team across orchestration, the money engine, the verification machinery, the adapter treadmill, and the durable-room infrastructure, plus the founder-equivalent judgment layer.

A senior team of that shape, over a 12-to-18-month window, is a fully-loaded spend in the **low-to-mid seven figures before a line of go-to-market**. That figure alone sits at or near the entire \$12M cap — and it buys only the *code*.

WHAT LEG TWO CONCLUDES

Replacement cost of the code alone approaches \$12M; the retired-work discipline, the six filed applications, and 409 real deliveries are additive on top of it (next page). **Leg two places fair value at or above \$12M on replacement grounds before demand or comps are counted.**

Sources. All build figures measured from the repository 2026-08-20 (moat facts); dated because the system grows fast enough that any printed count is already conservative. The fully-loaded senior-engineer cost is stated as a defensible market range, not a single sourced figure; the leg holds on the measured code alone.

Three things a fork does not get — and each adds to the floor.

3,900

retirement events — finished pieces of work produced, verified, and deliberately withheld for not clearing the bar.

THE RETIRED-WORK DISCIPLINE

The single fact a skeptic cannot argue with: a capability tool has no such number, because the operator decides what is good enough and never logs the takes they rejected. A fork gets the code; it does not get the calibrated judgment that produced 3,900 rejections, because that judgment lives in the constitution and the gate thresholds, not in any vendor's model.

6

filed U.S. utility patent applications on the non-obvious mechanisms.

THE FILED IP POSITION

Six filed applications, not granted patents (inventor Hicham Abdessamad; assignee Studio58, Inc.; first filed 2026-08-04). U.S. applications publish at roughly 18 months, so these will not appear in a public search until early 2028. Their value is real but bounded: they raise switching and copying cost and give an acquirer a cleaner story — not an issued monopoly.

409

delivered jobs, 17,267 files, 122,108 logged model calls, median 209 calls per finished deliverable.

THE DEMONSTRATED DEMAND

A demo proves the machine runs once; 409 deliveries prove it holds to a standard across hundreds of real jobs, which is a property that cannot be bought or accelerated. The printed number is the conservative one — this instance's ledgers only; the production registry is separate and higher.

A SUPPORTING FRAME · FOUNDER OPPORTUNITY COST

The founder previously ran Hitachi America as Chairman and CEO. His stated prior compensation is on the order of ~\$4M/yr — *the founder's own assertion, presented here as opportunity cost, never as an audited fact about Kaldune*. One year of that opportunity cost, foregone to build this alone, is a supporting data point that the human capital poured into the asset already approaches the post-money cap. It is a supporting frame only; the measured code, the discipline, the IP, and the demand carry the leg.

Sources. Retirement, delivery, and model-call figures from the room ledgers, measured 2026-08-20 (moat facts). Patent filing metadata is the company's own; the applications' **status — filed, not granted** — is held throughout this document. The ~\$4M/yr figure is a client-private founder assertion.

Two recognized frameworks, inputs shown, bracket \$12M.

Pre-revenue frameworks are judgment made explicit, not precision. Their value is that they force the inputs into the open. Two are standard for this stage; both are shown, neither is treated as an oracle.

METHOD ONE · SCORECARD (PAYNE)

Starts from a regional median pre-money for comparable pre-revenue companies, then adjusts by weighted factors. The base is illustrative; the leg leans on the direction of the factor reads, not a printed number.

FACTOR	WT.	READ
Management team	~30%	Strong +
Size of opportunity	~25%	Strong +
Product / technology	~15%	Strong +
Competitive environment	~10%	Positive
Sales / marketing / partnerships	~10%	Negative
Need for further investment	~5%	Positive
Other (traction / engagement)	~5%	Neutral–+

Six of seven factors read positive to strongly positive and one clearly negative (distribution). The multiplier lands materially above 1.0, placing the Scorecard pre-money into the low-to-mid eight figures. **The honest output is a range that brackets \$12M and sits above it at the center.**

METHOD TWO · BERKUS

Assigns up to a fixed value to each of five risk-reduction milestones. Shown with the conventional cap and Kaldune’s reads.

ELEMENT (RISK REDUCED)	READ
Sound idea (basic value)	Full
Prototype (technology risk)	Full — 409 deliveries
Quality team (execution risk)	Full
Strategic relationships (market risk)	Full to partial
Product rollout / sales (financial risk)	Partial / low

Four elements read full and one reads low. The strict conventional per-element cap was designed for an earlier funding era and underprices a product-complete company in a hot AI market, which is why Berkus is presented as the **conservative floor of the range, not its center**. Scaled to the AI seed environment, the same four-of-five-strong read supports a materially higher figure and rejoins the Scorecard and comps range.

WHAT LEG THREE CONCLUDES


Scorecard brackets \$12M and centers above it. Berkus at its strict conventional cap sits below (the deliberately conservative floor); scaled to this market it rejoins the range. Read together, the risk-adjusted range runs from a conservative low to a center in the mid-teens of millions, with **\$12M sitting at the conservative end of the credible range**.

Method note. Both frameworks are shown to demonstrate the argument does not cherry-pick methods. Base figures are illustrative rather than printed as unverified specifics; the frameworks’ value is the transparency of the factor reads, not a false-precise estimate.

Three independent methods, one reading.

No single method is asked to carry the argument; their convergence is the argument. Each measures something different, and each places fair value at or above \$12M.

METHOD	WHAT IT MEASURES	WHERE IT PLACES FAIR VALUE	READS \$12M AS
Comparable financings (3A)	Market's own pricing of the adjacent layer and the seed comp band	Floor at \$12M; adjacent capital 3 orders larger	Conservative
Build / replacement cost (3B)	Cost and time to rebuild what exists, plus IP and demand	At or above \$12M on code alone; additive from IP and miles	Conservative
Risk-adjusted scorecard (3C)	Weighted pre-revenue frameworks, inputs shown	Range brackets \$12M, centers above it	Conservative end



The convergence, drawn. The credible range and where each method lands. The \$12M cap sits below the center of the range on every reading — a conservative entry, not a stretch.

Every figure verified, or it does not print.

The rule is absolute: a figure that does not verify does not ship. An undisclosed figure prints as undisclosed. Nothing is estimated to fill a gap. Each figure below was followed to the source that owns or reports it and read on that source’s own page.

PLAYER	ROUND	VALUATION	DATE	VERIFIED STATUS & THE RULE THAT MUST HOLD
Cursor (Anysphere)	SpaceX agreed to acquire, all-stock	\$60B	signed 16 Jun 2026	VERIFIED. Say “agreed to acquire / acquiring,” never “acquired” — signed, not closed (close targeted Q3 2026). Prior priced round \$29.3B (Nov 2025).
Lovable	Series C, \$400M	\$13.3B	12 Aug 2026	VERIFIED (primary, company blog). Matches the deck exactly.
Replit	Series D, \$400M	\$9B	11 Mar 2026	VERIFIED. Matches the deck exactly.
Bolt (StackBlitz)	Series B, \$105.5M	~\$700M	Jan 2025	Round VERIFIED ; valuation <i>reported / officially undisclosed</i> — keep the tilde and the “(reported)” qualifier; last disclosed mark, may be stale.
v0 (Vercel)	part of Vercel	n/a	—	Part of Vercel, no standalone valuation.
Lindy	~\$50M raised total	undisclosed	2023	VERIFIED total raised; valuation <i>undisclosed</i> . Grouped with Relevance as “valuations not disclosed.”
Relevance AI	Series B, \$24M (\$37M total)	undisclosed	May 2025	VERIFIED total raised; valuation <i>undisclosed</i> .
The finished-work layer	—	open	—	Kaldune’s layer. The aisle below is funded from the hundreds of millions to sixty billion; the layer above is unfunded.

THE PATTERN IS WHAT CARRIES THE ARGUMENT


Two valuations are officially undisclosed and print as such; \$60B is a signed but not-yet-closed acquisition price and prints in the present progressive; ~\$700M is reported, not disclosed. None of that weakens the exhibit, because the load-bearing claim is not any single number. It is the pattern: capital concentrated from the hundreds of millions to sixty billion in the adjacent builder layer, and no named funded player in the layer above.

THE THREE POINTS A DILIGENT READER CATCHES

- 1. Cursor/SpaceX: “acquiring / agreed to acquire,” never “acquired” (signed, not closed).
- 2. Freelance high end is **Growth Market Reports**, not SkyQuest (market tiers).
- 3. Bolt’s ~\$700M is reported/undisclosed; the tilde and qualifier stay. All three are held throughout.

The market, as verified ranges — never collapsed to a single point.

All three tiers are stated as honest ranges spanning named research houses, which is the correct, defensible way to state a market size where sources draw the category boundary differently. The \$1.5T aggregate is labeled for what it is: Kaldune’s own constructed rollup, not a single-source citation.



Verified ranges, log scale. Consulting \$1.06T–\$1.1T (Research and Markets; IBISWorld). Agencies \$400B–\$445B (IBISWorld; TBRC) — the ceiling trimmed to a fully-source-backed endpoint. Freelance \$6.4B–\$10.3B (Grand View Research; Growth Market Reports — the corrected attribution).

**ON THE AGENCIES
CEILING**

Trimmed to **\$400B–\$445B** so both endpoints are confirmable on a named research-house page (IBISWorld \$444.7B; TBRC). A fully-sourced ceiling costs the argument nothing and removes a catchable soft spot.

**ON THE FREELANCE
ATTRIBUTION**

The \$10.3B endpoint traces to **Growth Market Reports**, not SkyQuest (whose own figure is ~\$6–7B). Both endpoints are real and sourced; the deck’s citation label was wrong and is corrected here.

ON THE \$1.5T ROLLUP

A framing figure, not a citation — directionally sound (consulting ~\$1.06–1.1T plus agencies ~\$0.4T clears \$1.5T before freelance and adjacent services). It carries no single-source footnote, by design.

ONE FIGURE THE DECK CARRIES THAT THIS DOCUMENT DOES NOT

The deck’s “+19.1% fastest-growing consulting line” (a FY2023 revenue release) was **not confirmed against the source’s own release** in this pass, so it is not printed here as a confident sourced figure. The comps and the measured moat carry the whole case without it: consulting revenue grew through the AI wave, a directional claim the freelance-divergence data supports, but no unverified specific is asserted.

Sources, read on each research house’s own page (2026-08-26). Consulting: researchandmarkets.com (\$1.0638T); ibisworld.com (\$1.1T). Agencies: ibisworld.com (\$444.7B); thebusinessresearchcompany.com. Freelance: grandviewresearch.com (\$6.4B); growthmarketreports.com (\$10.33B). The ~\$1.5T aggregate is the sum of those three markets, Kaldune’s framing.

05

The moat, measured.

Subtraction is expensive, and it shows in the numbers. Every figure that follows was measured from the repository on 2026-08-20. The dating is part of the argument: this system moves fast enough that any printed number is already conservative — the deck's own engine slide had gone stale in sixteen days.

The wiring is the smallest part of the building.

21%

of the engine is integration wiring — the adapters that call external tools.

79%

decides, sequences, verifies, meters, and recovers — the part that does not come in the box.

3.6:1

orchestration to wiring inside the engine; closer to 7:1 across the whole product.

“Just give me the wiring” asks for the fifth that was never the hard fifth.

MOAT · EXHIBIT A

The system throws away finished work, and keeps the receipts.

3,900

pieces of work produced, verified, and deliberately withheld from the client for not clearing the bar.

19

dedicated verification modules, across three gate stations: before, during, and after generation.

27

distinct named gate, verdict, and wall functions invoked across the engine.

EXHIBIT B

Nobody demos the work they threw away. A capability tool has no such number, because the operator decides what is good enough and never logs the takes they rejected. **This is the one fact a skeptic cannot argue with.**

Measured from the repository, 2026-08-20. Line-weight, gate, and retirement figures from the moat-facts crawl; retirement events from the room ledgers.

The build, the miles, and the filed applications.

EXHIBIT C · THE BUILD, IN NUMBERS

Shipped production TypeScript	~200,000 lines
Equipment adapters (28 API hosts, 30+ models)	54
Money-engine modules (atomic accounting)	37
Engine / orchestration modules	103
Test files (named suites)	545 (366)
Constitution	791 lines
Encoded operating judgment	~104,000 words
Commits / active days	2,099 / 49

EXHIBIT D · THE MILES, AT PRODUCTION SCALE

End-to-end rooms run	445
Rooms reaching delivery	409
Delivered files (9+ kinds of work)	17,267
Model calls logged	122,108
Median calls per finished deliverable	209

A demo proves it runs once; 409 deliveries prove it holds. This instance's ledgers only — the production registry is separate and higher, so the printed number is the conservative one.

EXHIBIT E · THE SIX FILED APPLICATIONS — THE STATUS HELD HARD

Six **filed U.S. utility patent applications**.
Inventor: Hicham Abdessamad. Assignee: Studio58, Inc. First filed 2026-08-04 as Application 64/125,281.

These are **filed applications, never “patents,” never “granted,” never “patented.”** As applications filed 2026-08-04, they publish at roughly 18 months, so they will not appear in a public search until around early 2028. That is stated as a fact and never leaned on as though the claims were examined or allowed. The filings publish the mechanisms; the calibration behind every threshold stays trade secret.

- 01 The crash-recoverable on-disk engine
- 02 Deterministic media-truth verification
- 03 The client's-word-is-law constraint mechanism
- 04 Capped concurrent-spend money physics
- 05 Cross-server fleet custody
- 06 Safe delivery of sandboxed apps

Measured from the repository, 2026-08-20. Patent filing metadata is the company's own; the applications' status — filed, not granted — is the discipline this document holds, not a figure to verify externally.

The raise accelerates a working company. It does not keep a fragile one alive.

\$2M

raise, up to, on a post-money SAFE. The founder's stated terms.

\$12M

post-money cap, locked.

\$650K

already committed against the round.

THE CAP-TABLE POSTURE

A SAFE round: no priced-round overhang, no preferred stack, no investor consent required on a sale. This is the clean cap table the acquisition story rests on (Section 8), and it is a genuine feature for a seed investor — the entry is simple, the exit path unencumbered.

TWO TRANCHES, TWO DIFFERENT THINGS

The product risk is already retired and self-funded to product-complete (Section 5), so the two tranches buy different things. Neither funds survival, because there is nothing to keep alive that is not already running.

THE ONE-LINE VERSION

The committed money proves the model; the full raise scales it.

THE \$650K ALREADY COMMITTED

Funds the immediate next steps: the first paid distribution tests (the founder-led social and hand-placed access-code channels, where the code system is already built and shipped), the working-capital cushion to begin absorbing enterprise slow-pay, and the SOC 2 process that unlocks enterprise procurement. It gets the company from product-complete to revenue-generating with instrumented distribution.

THE FULL \$2M UNLOCKS SCALE

Funds the vendor-commitment volume discounts, the light hiring and contractor bench, the advisor equity, and above all the working-capital float that lets enterprise revenue grow without the cash gap choking it. The full raise is the difference between testing distribution and scaling the channels that pay back.

Terms. Raise size, structure, and the committed figure are the founder's stated terms (client-private), to be confirmed on the final SAFE. The cap is locked; money is committed against it.

The raise buys customers and the cash to serve them — not runway.

Method note: the figures below are an illustrative allocation of the \$2M in rounded buckets with stated rationale, not exact commitments. A seed company does not know its channel mix to the dollar before it runs the tests; round buckets with clear logic are more defensible in diligence than fake-exact line items.

BUCKET	SHARE	WHAT IT FUNDS	RATIONALE
Working capital / collection float	~30–35%	The cash gap between prepaying vendors for compute and collecting on enterprise net-30/60 invoices	The headline line , and the answer to “why raise if unit economics are positive.” Margins are positive per deliverable (64–85%+ measured GM), so this funds the <i>timing</i> difference, not losses, and grows as enterprise revenue succeeds.
Go-to-market / distribution	~30–35%	Paid channels deployed against a measured cost-per-customer and payback bar; channels that clear the bar get more, the rest get cut	The test-then-scale discipline: nothing scales before cost-per-paying-customer and payback are known, and every user is revenue from the first brief, so payback is measured from day one.
Vendor / provider commitments	~10–15%	Committed spend to unlock volume discounts on compute and model access	Directly improves already-positive unit economics; a lever on gross margin, not a cost center.
SOC 2 certification	~5%	The audit and remediation to reach SOC 2	Enterprise-unlock spend. Procurement at many enterprises will not sign without it; funding it converts seeded trials into signable contracts.
People: founder + advisors + light hiring	~10–15%	~\$100K/yr founder salary; advisor equity/retainer; a first key hire and 1099 contractors as needed	Deliberately lean. The ~\$100K founder salary is disciplined and below-market for the role. Hiring stays light and contractor-weighted; no invented salaries for unnamed future hires.
Reserve / contingency	~5%	Unallocated buffer	Signals no bridge round is planned: the company intends to reach its milestone on this raise and holds its own cushion rather than assuming it can return to the market.

THE MILESTONE ANCHOR

The allocation is scoped to reach one destination: **the founder’s stated target of \$25M–\$50M ARR**, his aspiration and the milestone the raise is sized against, *never a studio forecast or endorsed number*. Because margins are positive per deliverable, revenue is gated not by burning capital on the work but by finding channels that pay back and carrying the float as they scale, which is what the top two buckets fund.

RUNWAY & THE NO-FURTHER-CAPITAL THESIS

The raise is scoped to **8–24 months**, burn dominated by working-capital float and go-to-market spend. **The round does not fund compute**, and margins improve as model prices fall.

A company that holds its own contingency and does not fund its cost of goods from the raise does not need a bridge: the discipline a \$12M cap already signals.

Note. Allocation shares are illustrative rounded buckets, not exact commitments. The \$25M–\$50M ARR target is the founder’s stated aspiration (client-private), never a studio forecast.

Named before the reader names them — each with its mitigation in the same line of sight.

The strongest investor document names what could bite. A serious investor prices these risks; this section shows they have been seen clearly, and where the raise funds a fix, exactly how.

RISK 1 · KEY-PERSON

A sole founder, no employees on record.

All six patent applications name him as sole inventor. If he is unavailable, the company is exposed. This is the first and largest risk, stated plainly.

MITIGATION

The judgment that would otherwise live only in his head is already documented and machine-enforced: a 791-line constitution inside ~104,000 words of encoded guidance, 19 verification modules, 27 named gates, 545 test files. The system runs its standard without him watching each brief, and the raise funds advisor equity and a first key hire to reduce single-point dependency.

RISK 2 · THE ARR CLAIM AS A DOUBLE-EDGED SWORD

\$25M–\$50M ARR on \$2M can read as a brag.

An investor is right to test a raise scoped to that target.

MITIGATION — THE BRAG BECOMES A PLAN

The target is the founder's, the milestone the raise is scoped to reach, never a studio forecast. Unit economics are positive per deliverable (64–85%+ measured GM), so revenue is not gated by burning capital on the work; it is gated by finding channels that pay back and carrying the collection float as they scale, both sized in the use of funds. The claim is defensible precisely because the money is not funding the cost of goods.

RISK 3 · CONCENTRATION & SLOW-PAY, COMPOUNDING

Enterprise revenue on net-30/60 terms.

A small number of large accounts may carry a disproportionate share, and the cash gap widens as those accounts succeed — more revenue on longer terms means more cash tied up. Concentration and slow-pay compound each other.

MITIGATION

The working-capital float is the headline use-of-funds line for exactly this reason, and the reserve absorbs the lumpiness of a concentrated, slow-paying book. As the customer base broadens beyond the first large accounts, concentration falls; until then, the float and reserve are the designed buffer.

RISK 4 · PLATFORM / SUBSTRATE DEPENDENCY

Kaldune runs on providers it does not own.

Provider price changes, rate limits, terms shifts, or a provider standing up a competing finished-work layer are all real exposures.

MITIGATION

Two structural defenses. Because Kaldune sells the finished work at a multiple of measured cost, every provider price cut flows to margin, so provider pricing is a tailwind. And the moat is the 79% of the engine that decides, verifies, meters, and recovers, provider-independent, not the 21% wiring a swap touches. A provider that tried to build this layer faces the same 12-to-18-month miles.

Pendency, enterprise-readiness, and the category.

RISK 5 · PATENT PENDENCY

Six filed applications, not granted patents.

First filed 2026-08-04. They have not been examined; claims may be narrowed or rejected; and because U.S. applications publish at ~18 months, they will not be publicly visible until roughly early 2028. Anyone treating this as an issued, defensible monopoly today would be wrong.

MITIGATION — THE HONEST FRAMING IS THE MITIGATION

The filings are presented throughout this document exactly as what they are — filed, pending, not granted — so no diligence finding contradicts the doc. Their value is real but bounded: they raise switching and copying cost and give an acquirer a cleaner IP story. The defensibility argument rests on the measured moat, the retired-work discipline, and the miles, with the pending IP as one supporting layer.

RISK 6 · ENTERPRISE-READINESS GAP

SOC 2 is not yet in hand.

The company is seeding enterprise trials, but enterprise procurement often will not sign without SOC 2. A trial that cannot convert to a signed contract is not yet revenue.

MITIGATION — NAMED WITH ITS FIX

SOC 2 certification is an explicit enterprise-unlock line in the use of funds. The gap is named honestly and the fix is funded in the same breath: the raise pays for the audit and remediation that convert seeded trials into signable contracts. Until it lands, the gap is real; the raise closes it.

RISK 7 · CATEGORY & TIMING

A crowded, fast-moving space.

“AI agency” is crowded, the funding climate could flip, and a well-capitalized entrant could target the same layer.

MITIGATION

The posture does not depend on the funding climate: product-complete, self-funded to that point, positive-margin per deliverable — a downturn that starves demo-stage competitors leaves a shipping, revenue-capable company standing. On the entrant threat, the defense is the miles: a funded team copies the demo in a month, but their clock to match a standard held across 409 deliveries runs 12 to 18 months. The crowd is one layer down, in the builder aisle; the finished-work layer has no named funded player.

THE READING OF THE RISK SECTION

None of these risks is hidden and none is fatal. The two that most affect the valuation — key-person and distribution — are precisely what the raise addresses, which is why a \$12M cap that prices them in is still conservative: **the product risk that usually kills a seed company is already retired, and the rest are the ordinary, fundable risks of a company that has built the hard part and now needs to sell it.**

\$12M sits at the conservative end of a triangulated range.

*Three independent methods, inputs shown, converge on the same reading.
No single method carries the argument; their convergence is the argument.*

METHOD	FAIR-VALUE PLACEMENT	\$12M READS AS
Comparable financings (3A)	Floor at \$12M; adjacent-layer capital 3 orders of magnitude larger; direct seed comp band mid-to-low	Lower half of the band
Build / replacement cost (3B)	Code replacement alone approaches \$12M; IP, retired-work discipline, and 409 deliveries additive on top	At or above, before demand
Risk-adjusted scorecard (3C)	Scorecard brackets \$12M and centers above it; Berkus (scaled) rejoins the range	Conservative end

THE CREDIBLE RANGE

Read together, the methods place defensible fair value in a band from the mid-single-digit millions at the deliberately conservative Berkus floor to the mid-teens of millions at the comps-and-scorecard center. **\$12M sits at the conservative end of that range.**

WHY THE CONSERVATIVE END IS THE RIGHT PLACE TO BE

The cap is locked because money is already committed at it, and this document has not argued to move it. The argument is stronger when the cap is defended as fair-to-favorable than when it is stretched.

A NEW INVESTOR ENTERING AT \$12M IS ENTERING...

- below the market’s pricing of the adjacent layer by three orders of magnitude, and at or below the cost to rebuild what exists;
- inside the lower half of the direct seed comparable band;
- into a company whose product risk is retired, measured 2026-08-20, filed in six pending applications;
- behind a founder disciplined enough to price his own company conservatively and take a ~\$100K salary to build it.

THE ONE-SENTENCE CONCLUSION

The \$12M cap is not a number a pre-revenue company reached for. It sits below three independent ceilings and at the conservative end of a triangulated range, which makes the entry fair to favorable, and makes the founder’s discipline, not his ambition, the thing the investor is underwriting.

Kaldune



THE FINISHED-WORK LAYER

Empty, on a \$1.5T market. It's where we built.

Every layer below sells building power to people who build. The top layer sells the finished work — and the \$12M cap on it is a conservative entry, defended on three independent methods, each landing at or above it.

Kaldune, a product of Studio58, Inc.
Hicham Abdessamad · Founder
Former Chairman & CEO, Hitachi America

CONFIDENTIAL
PROSPECTIVE SAFE INVESTORS
MMXXVI