
An agency that is AI.

We don't sell you AI. We sell you the finished work.

Brief it. Watch it. Receive it.

Productizing the \$1.5 trillion services industry.

Everything from Fiverr to McKinsey sells the same underlying thing: finished work delivered by scarce human experts. That industry kept growing straight through the AI wave. **That is the tell.**

THE HOME DEPOT

Today's AI is Home Depot: aisles of powerful parts, ten variants of everything, and you still have to build the house yourself. Nobody wants to build a house. They want to live in one.

THE TELL, IN THE DATA

When AI arrived, services revenue did not fall. It rose. The world's largest consulting firm grew its consulting line 19.1% in the year generative AI broke, its fastest-growing business. AI shrank only the work it can finish alone: generic writing, translation, basic support. Everything that still needs a builder, video, design, working software, strategy, grew. AI is the aisle. It sells better parts every quarter. It did not become the house.

THE HEADCOUNT CHAIN

Services revenue is chained to headcount. Every services CEO knows it: to grow 20% you hire 20% more people, from the same scarce talent pool everyone else is fighting over. AI made the parts cheaper. It did not make the assembly easier. The human-capital problem moved. It was not solved.

Kaldune breaks the chain. A plain-language brief goes in. A finished deliverable comes out, produced by an integrated AI team the client watches work, **at software economics.**

SIMPLE ON TOP, HARD UNDERNEATH


Engineers look at it and brush it off: no substrate, looks easy. The non-technical person asks when they can have it. Same object, two opposite reactions, and the reaction sorts the person before they speak. The simplicity is earned. Six U.S. utility patent filings sit underneath it. Making something this hard feel this simple is the achievement, not the accident.

WHO IT IS FOR

A fair objection: some teams already use these tools well and ship with fewer people. That is real, and it is the point. Used the right way, by the right hands, the aisle unlocks a great deal, and those hands do not need us. If the tools already work for you, you were never who we built this for. Kaldune is for everyone the aisle leaves behind, and for the experts who simply do not want to be the integration layer.

Market: \$1.5T+ annual. Freelance \$6.4B to \$10.3B (Grand View Research; SkyQuest). Agencies \$400B to \$455B (IBISWorld \$444.7B; TBRC). Consulting \$1.06T to \$1.1T (Research and Markets; IBISWorld). Growth: Deloitte Global FY2023 revenue announcement, deloitte.com (Consulting +19.1%, fastest line). Segment divergence: analysis of 5M Upwork jobs, Nov 2022 to Feb 2024 (bloomberry.com); corroborated by Upwork AI-impact research, 2025.

The one picture that proves it: the aisle, not the house.



When AI arrived, the work it can finish alone shrank. Everything that still needs a builder grew, and that gap is exactly the market Kaldune productizes.

Sell the work, not the model.


“If you sell tools, you have to compete on model and speed. If you sell the work, the better the model gets, the faster and cheaper your service becomes.”

Julien Bek, Partner, Sequoia Capital, “Services: The New Software,” 2026.

For every \$1 of software, \$6 is spent on the services around it. That \$6 has always been trapped in the agency model. Kaldune moves it to software economics.


Model access got ~300× cheaper in ~16 months

Cheapest capable model input price per 1M tokens. Log scale.



The collapse runs on the labs' balance sheets, not ours. Every dollar the model cheapens is a dollar of margin.


The stack is filling from the bottom. The top layer is empty.



Every layer below sells to people who build. Kaldune sells to people who buy.

Over \$1.5 trillion a year is spent on exactly this work.

The \$6 of services behind every \$1 of software, priced by three tiers that put it out of reach. Kaldune sits below the floor, at software economics.



SOFTWARE ECONOMICS

Software costs the same to serve the tenth customer as the first. Every plan scales with server cost, not headcount. A deliverable that costs \$10,000 in agency labor can ship at \$199 without cutting margin.

SEALED-ENVELOPE PRICING

Every brief has a quote before it starts, and the quote is the ceiling. The studio can spend less to deliver, never more. Both sides know the number before the work runs.

You brief it, you watch a team work, you receive the finished thing.

01

Brief.

Describe what you need in plain language. No prompt, no model picker, no settings.

02

Watch.

A full AI creative and strategy team produces the deliverable on screen. You see it happen.

03

Receive.

Investor-ready plans, launch strategies, brand identities, films, books, working software. Done.

04

Zero substrate.

The client never sees a model name, a setting, or a tool. No prompt engineering, no model picker, no assembly.

Built for the 90% AI currently leaves out, and for the experts who don't want to be the integration layer.

Ferramante · Casa Alborán · AEQUA · Migaja · The Keys · Villagio · Kitchen Reborn

Real deliverables, each briefed in plain language by someone with no design or AI skill.

08 — THE PROOF

Every panel is a real brief, developed end to end.

Not stock, not screenshots, not AI demos.

ARCHITECTURE

FERRAMANTE

*Made to move.
Made to disappear.*

BRAND IDENTITY

INDUSTRIAL SYSTEMS

SKINCARE

INDUSTRIAL DESIGN

COUTURE

RENOVATION

MUSIC

Take Your *Time*

PERSONAL FINANCE · DECISION TOOL

VERDICT

Yes. 2027 is realistic.

~\$455k

~\$625k

COMFORTABLE

THE CEILING

Hire people, or become the integration layer yourself.

	HIRE PEOPLE	STITCH TOOLS YOURSELF	KALDUNE
Lead time	Weeks	Your nights and weekends	1 to 3 hours
Cost	Four to six figures	A dozen subscriptions	Software economics
Integration	They meet about it	You are the layer	One integrated team
Output	Advice or a draft to polish	Whatever you assemble	A finished deliverable
Accountability	Retainer plus hourly	On you	Sealed envelope plus finish guarantee

Kaldune is the third way: one integrated AI team with taste, process, and accountability, end to end, finished from one conversation.

One tool, played across the whole market, top to bottom.

1 · The firms deliver uncapped.

Agencies and consultancies take on more clients without hiring, and hold quality when work slows. No headcount ceiling up, no idle bench eating margin down. The experts get the most out of it. It protects margin.

2 · The buyers buy less, and compound.

The companies that purchase agency and consulting work buy less of it, supplement their own team, and move faster. Work that took weeks lands the same day, so the decision compounds sooner. The team gets stronger instead of renting strength by the hour.

3 · The floor gets the top of the market.

The one it was built for. The small-business owner, the solo founder, the person with a good idea and no budget. Today they build a crooked house or never start. Kaldune hands them the best strategists, designers, and engineers in the world at a fraction of the cost. It levels the field.

THE SAME TOOL, FOUR KINDS OF BUYER. IN DETAIL:

1 The firms themselves.

Agencies and consultancies. They use AI now. Kaldune lets them deliver more, faster, at higher margin and quality, and being the experts they iterate with it best of all.

Villagio · Casa Alborán. Briefs: an investor pack and the building rendered, memorandum and pro-forma, the launch marketing.

2 Mid-market and enterprise.

Departments that buy agency and consulting work all year. Per-seat, pooled wallet, self-serve at scale.

The developer doing a spec build. Would otherwise pay a viz studio plus a deck designer plus a copywriter, five figures and weeks. Pays Kaldune: a Professional-lane deliverable, hours not weeks.

3 Small business and startups.

The gym, the restaurant, the founder. Owner-operators who were never going to sign a retainer.

The operator launching a brand. Ferramante · Migaja. Would otherwise pay a branding agency four to six figures, months of lead time. Pays Kaldune: a Craft-lane brand kit at a measured 70%+ gross margin, same day.

4 The prosumer at the floor.

The person with an idea and no team. Needs agency-grade work, has none of the budget. The entry tier.

The founder with no design budget. Plan · strategy · dossier. Would otherwise pay a consultant, or ship something weak. Pays Kaldune: a Professional-lane plan against a \$5K to \$50K engagement, in an afternoon.

The sharpest proof is at the top: the firms we'd displace are also our buyers. Kaldune lets a 15-person agency scale without hiring. From the firm to the founder, one tool briefs like a studio and finished work comes back. And the one group we are honestly not for, the team the tools already serve well, only sharpens the line: if the aisle works for you, you were never who we built this for. Kaldune is for everyone it leaves behind.

The agency makes you pay twice: the fee and the wait. Kaldune collapses both.

Briefs are done in 1 to 3 hours based on complexity, iterations included. The agency alternative is weeks of back-and-forth: iteration rounds, meetings, calendars. The wait is pure opportunity cost. The business that has its plan this afternoon moves this week.

BRIEF TO FINISHED, BY DELIVERABLE		
Business plan	Agency 2 to 6 weeks	Kaldune 1 to 3 hours
Brand identity	Agency 2 to 6 weeks	Kaldune 1 to 3 hours
Launch film	Agency 2 to 6 weeks	Kaldune 1 to 3 hours

Fast and finished, not fast instead of finished.

Kaldune range founder-observed across the finished-brief corpus, iterations included. Not a benchmarked SLA.

It will look simple. That was the hard part.

Engineers look at Kaldune and see nothing special, because there is nothing of the machinery to see: no model picker, no settings, no substrate. That reaction is the product working.

The general assistants sell chat with intelligence and leave the assembly, the taste, and the finishing to the user. Kaldune removed the assembly from view entirely.

When a bridge holds, nobody applauds the foundations.

The next four slides are the foundations.

The actor doesn't make the movie. The production around the actor is the IP.

The models sell intelligence by the token. Intelligence is the actor. A finished film also needs the script, the set, the director, the editors, the production office, and the box office.

A production line for finished work. Every module is a machine part: brief intake, room lifecycle, adapter routing, delivery packaging, money math, quality gates, auto-heal. Nothing on this list is a wrapper around a chat prompt. Foundation models get cheaper and better on someone else's balance sheet. The production line that turns them into delivered work is the part nobody hands you.

64

orchestration modules

46

production adapters

744

line constitution

18

money modules

305

automated tests

313

production runs

The counts say how big it is. The filings say why it's hard.

Cheaper models don't erode this. They feed it: the same collapse that commoditizes model access widens our margin on every brief.

313 production runs through the full pipeline. A portion were briefed by outside parties with real business problems, not by the founder. Zero paid to date, because registration has not opened; several of those parties intend to convert at launch.

The engine, six layers, one production system.

Every layer is real code that ships on every brief. Every count is a file count, not a marketing count.

L1 Access perimeter	Cloudflare edge · HTTPS load balancer · Signed identity envelope · Email-code OTP · Consent gate · Security headers · Session · WebSocket switch.
L2 Studio pod	Wake orchestrator · Room lifecycle · Epoch fencing · Ask protocol · Constitution engine · Client-text fencing · ~30 opcodes · 744-line constitution · epoch-fenced.
L3 Specialist runtime	Principal loop · Tool router · Auto-heal ladder · Safety classifier · 6 quality gates · Quality gates.
L4 Equipment closet	Image 10 · Video 8 · Audio 4 · Web/Data 5 · Kit/Delivery 7 · Document 12 · 46 adapters · vendor-agnostic · swap at config.
L5 Money engine	Sealed envelope · Guarantee door · Settle · COGS breaker · Meter (atomic) · Stripe webhook · Refund/dispute · Wallet plus ledger · 18 modules · atomic · circuit-breakered.
L6 Persistence plus ops	Postgres (private IP, PITR) · GCS versioned · BILLING_DIR on PVC · Weekly digest · Capacity alarm · Purge worker · 305 tests · 8 alerts wired.

A brief, a delivery.

64 orchestration modules. 46 adapters. 744-line constitution. 18 money modules. 305 tests. One brief, one delivery, one guarantee.

The codebase has laws.

Most software is an accretion of fixes. This engine has a constitution, visible in the code itself and applied uniformly across every layer, written under production pressure and never broken.

THE FIRST LAW

Machinery enforces physics. Intelligence owns judgment. The daemon makes no decisions.

Every gate states facts. Nothing is ever silently edited.

Money is physics, not policy.

Nothing is ever destroyed, only retired.

The client's word is the one sovereign override, and even then the gate still looks.

The machinery never paraphrases anyone's words.

Every failure lands as a stated fact with the reason journaled.

This law was learned the expensive way: the first engine was built the way all software is built, deterministically, and it strangled the intelligence it was meant to serve. The refactor that reversed it is the foundation everything else stands on.

Six patent filings. Every one is a mechanism the obvious build gets wrong.

Anyone can wire a model to APIs. What ships to a paying client is the machinery around it, built from two years of production failures. A sample of what the filings claim:

The engine

filed August 2026

Agents whose entire state lives on disk. A crash recovers by reading journals, never by re-running paid AI work. Context is managed against the provider's cache billing.

Media truth

filed

Deterministic instruments verify generated work before the client sees it, including a lip-sync gate that catches a frozen mouth on a closing line by cross-checking audio, speech timing, and pixel motion. Facts, not scores.

The client's word is law

filed

Corrections become machine-enforced constraints in the client's verbatim words, newest word wins. Every quality gate can be overridden by one thing only, the client's own quoted words, and even then the gate still looks.

Money physics

filed

Any number of AI workers spend concurrently against one wallet, with the house's worst-case loss capped at roughly two dollars. Refund and free-credit attacks closed by mechanism, not policy.

The fleet

filed

Engagements follow the user across servers with live custody hand-off. A server crash loses nothing. A vendor outage is invisible.

Safe delivery

filed

AI-authored apps run sandboxed. Hostile archives are refused whole. Published pages can never touch the platform's users, and a client's published address never resolves to a stranger.

Six U.S. utility patent filings. The first is filed as Application 64/125,281 (filed Aug 4, 2026).

The filings publish the mechanisms. The constitution, the prompts, and the calibration behind every threshold stay trade secret. Two layers, one moat.

You can buy the parts. You can't buy the miles.

Anyone can assemble the same parts: models, APIs, a chat surface. It won't drive like the car, because the secret is not in the parts, it's in the sum: the integration, and the thousands of cars built before this one.

The product lives in what breaks only with real clients: the frozen mouth on the closing line, the correction the AI forgets an hour later, ten workers spending the same five dollars, the crash that eats a paid render. 313 production runs found those failures. The engine closed each one. The patent filings cite them.

A funded team can build the demo in a month. Their clock starts when their first hundred real briefs start breaking things we already fixed. That's the 12 to 18 months.

Four lanes, priced against what the work replaces, at a measured 30% to 85% gross margin.

LANE	PRICE RANGE	EXEMPLAR	GM TARGET	MEASURED GM
Entry	\$19 to 49	\$29	30%	~64%
Craft	\$99 to 299	\$199	60%	70%+
Professional	\$249 to 795	\$495	75%	85%+
Flagship	~\$1,500 to 4,900	in build	high	in build

Retail = lane multiplier (2× / 5× / 6×) times actual measured COGS, clamped at the quoted ceiling. A hot brief never blows the client's quote. The COGS breaker (2× ceiling) caps the house's downside. Flagship anchors against \$5K to \$50K consulting engagements.

MEASURED COGS · REAL 2026-07 BRIEFS

Naming	\$10.54
Logo kit	\$31.27
Book	\$43.50
Gym plan	\$37.03
Real-estate launch pack	\$79.81

Measured, not modeled. Margins measured across 313 production runs. Price cuts flow straight to margin.

Pay-as-you-go for the middle, subscriptions for the returners, and revenue from the first brief.

PLAN	\$/MO	WALLET	SEATS	CONCURRENCY
Studio	\$199	\$235	2	2
Growth	\$549	\$649	5	5
Agency	\$1,490	\$1,790	10	10
Enterprise	\$179/seat	\$215/seat	10 min	1/seat

Annual = pay 10 months for 12, wallet dripped monthly. Live in the Stripe catalog with 19 payment methods enabled; tax-wired.

THE PAID DOOR

No free tier. The first brief requires a paid balance (smallest pack, \$25), then is guaranteed to finish even past the balance. The card also structurally gates guarantee-farming and bot abuse.

~3× CONVERSION

Card-required trials convert about 3× no-card at industry median (~44% vs ~14%). Every user is revenue from the first brief.

THE PRECEDENT

Midjourney runs subscription-only, no free tier, and reached ~\$500M ARR with \$0 venture capital and a small team. A paid door is not a handicap; it is a durable business.

Conversion benchmark: industry median (2026). Midjourney: Sacra; GetLatka; 2026.

It scales horizontally today. The only real ceiling is vendor throughput, and that is commercial.

PRODUCTION TOPOLOGY · GCP US-CENTRAL1

Cloudflare DNS (60-second reversible cutover) → Global HTTPS load balancer → Landing (Cloud Run plus CDN) and Gateway (identity perimeter) → Studio pod (GKE StatefulSet) → Cloud SQL Postgres (identity, ledgers) and GCS (rooms, deliverables, versioned).

Horizontally scalable today

Deterministic room placement, pod-death recovery tested, ready to add capacity by adding pods. Rooms follow the user across pods with live custody hand-off, covered by the fleet filing.

Vendor-agnostic engine

Every model choice and every media vendor is a config edit, not a rebuild. Model-price cuts flow to margin.

The bottleneck is commercial

Vendor rate limits move with spend; the code does not. Each pod carries roughly 100 to 120 concurrent active rooms, measured. Adding capacity is adding pods.

Identity: Google sign-in and passwordless email-code; no passwords in the system. Secrets: all vendor keys fetched at boot. Money safety: per-room COGS circuit breaker pauses any room whose real spend crosses 2× its quoted ceiling (floor \$25).

Pre-revenue and product-complete. The velocity is the traction.

Pre-revenue. Product-complete. Launch is scheduled. The company runs on the same thesis it sells: one founder, an AI team, agency output.

313

production runs through the full pipeline. Ferramante, Casa Alborán, AEQUA, Migaja, The Keys, Villagio, Kitchen Reborn. Every demo room is drawn from these.

9

production features landed in a single cycle, built by one founder with no employees.

What's left is the easy kind of hard

The years-long problems that cannot be hired around are solved and filed. Remaining engineering is hardening, infrastructure work of a well-understood kind, already underway. The open question is distribution, and it is the one the raise exists to answer.

LANDED IN THE LAST CYCLE

- Workspace and teams layer with pooled wallet
- Self-serve enterprise seats
- Annual billing (10 for 12)
- In-studio Plans page
- Pricing advisor with governed rails
- Passwordless email-code login (rate-walled)
- Live Stripe catalog, 19 payment methods
- COGS circuit breaker with kill switch
- Full go-live runbook with rollback per phase

A portion briefed by outside parties; zero paid to date because registration has not opened.

Spend follows the number. Channels that clear a short payback bar get more; the rest get cut.

The plan is not a spray. It is a measurement program with a bar attached, run in the order a disciplined operator runs it.

1

Test the cheapest channels first.

Founder-led LinkedIn, where the demo film is itself the ad, and aimed access codes placed by hand with specific owner-operators, professionals, and micro-agencies so the first public case studies are strong. The code-minting and redemption system is built and shipped.

2

Measure two numbers.

A defined measurement budget exists specifically to find the cost per paying customer and the payback period on each channel. Nothing scales before those two numbers are known. The funnel instrumentation is already built: signup, first paid brief, second brief, plan conversion.

3

Scale only what pays back.

Remaining demand-generation spend goes only into channels that clear a short payback bar. A channel that does not clear it does not get more money. The bar is the gate.

Every user is revenue from the first brief, so payback is measurable from day one, not modeled. Because of the paid door, every channel's payback is a measured number, not a forecast.

A paid door and finished output is a durable business, not a handicap.

COMPANY	WHAT THEY SELL	MODEL	SCALE REACHED
Midjourney	Finished images and video	Subscription only, no free tier	~\$500M ARR, \$0 VC, small team
Lovable	Plain-language in, working app out	Subscription, self-serve	\$0 to \$100M ARR in 8 months
ElevenLabs	Finished voice and audio	Subscription plus usage, paid door	~\$1B ARR inside 2 years

These are not valuation comps. Kaldune is pre-revenue and multiples do not apply.

They are business-model precedent: plain-language in, finished work out, sold behind a paid door. Midjourney is the closest structural analog: subscription only, no free tier, profitable, tiny team, \$0 venture capital. The difference is what Kaldune finishes: not one image, not one app, but the whole agency and consulting deliverable from one brief.

Cursor/Anysphere acquisition: Reuters, June 16 2026 (all-stock, expected to close Q3 2026, pending regulatory approval). Midjourney: Sacra; GetLatka (2026). Lovable: CNBC (2025). ElevenLabs: Analyze AI (2026). Figures are reported estimates.

HOW FAST THE CATEGORY RE-RATES ONCE REVENUE EXISTS

Lovable

\$1.8B (Jul 2025) to \$6.6B (Dec 2025), roughly five months

Cursor

\$9.9B (Jun 2025) to \$29.3B (Nov 2025), then SpaceX agreed to acquire parent Anysphere for \$60B (Jun 2026)

Pre-revenue, so these multiples do not apply directly; the point is the pace of the mark, not the number.

The raise does not fund model bills. Margins are positive per deliverable.

WORKING CAPITAL

Covers the collection float. Infrastructure cost plus the gap between prepaying media vendors when a brief runs and Stripe settling the client's payment. The float is a timing cost, not a loss, carried cleanly through launch.

DEMAND GENERATION · THE MAJORITY OF THE ROUND

Deployed against a measured cost per paying customer and payback period, not sprayed. Channels that clear a short payback bar get more; the rest get cut. Test first, then scale what pays back.

The one thing this company is missing is distribution, and missing only distribution is the best position a company can raise from: the product risk is retired, measured, and filed. *The raise buys customers, not survival.*

THE ORIGIN

Before he built the seller, he was the buyer. The founder ran Hitachi America, a major industrial company. He bought exactly the agency and consulting work Kaldune now delivers, and he signed off on the \$6 of services behind every \$1 of software as a line item, not a statistic from a report. He also ran that model from the inside, so he knows why it cannot be fixed there: consulting revenue is bounded by billable people and rate cards, and those firms already use AI heavily. It makes the work cheaper to produce without making it cheaper to buy, because the model bills for hours.

The tool he built to get out of his own way turned out to be the product. He set out to build a serious product on his own: trust infrastructure for autonomous systems. Building it alone forced him to be the integrator, the man in the middle stitching every tool together by hand. So he built the thing that takes the integrator out of the loop, then saw what it was: the fix was the bigger business. Kaldune replaced the team he would have hired to build it, and became its own first customer, because he was its hardest problem.

Hicham Abdessamad

Founder, Kaldune (a product of Studio58, Inc.). Former Chairman & CEO of Hitachi America. A solo-founder operation shipping the surface area most teams staff a full engineering org to build.

The return path is acquisition.

The company is built to be cleanly acquirable. This is the founder's view of how this returns, stated as reasoning, not a guarantee. There is no priced round planned and no timeline attached to it.

Who acquires

The platforms that already sell creative and productivity software, the agency and consulting holding companies that need an AI-native production line they cannot build fast enough themselves, and the model labs looking for a finished-work business their API cannot become on its own. Every one of them faces the same build-versus-buy math the engine slides lay out.

Why this one

A clean cap table. No preferred stack and no investor consent gates on a sale. A SAFE round, so no priced-round overhang. Transferable architecture: documented, runbooked, 305 tests, vendor-agnostic config, nothing locked in one person's head. Measured unit economics an acquirer can underwrite on day one.

The precedent on the board

In June 2026, SpaceX agreed to acquire Anysphere, the parent of Cursor, for **\$60 billion** in an all-stock deal. A category-leading AI application company was acquired, not competed with, by a buyer that ran the same build-versus-buy math and chose buy.

The realistic shape is a strategic acquisition once the funnel is proven and the corpus of finished work keeps compounding. The zero-employee design is not a gap to fix before a sale; it is **part of what a buyer is buying**.

Cursor/Anysphere acquisition: Reuters, June 16 2026; expected to close Q3 2026, pending regulatory approval.

This deck was produced by the product it describes.

Produced by the product it describes. Kaldune.

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

Named, and the line each one leaves open.

Positioning, not valuation.

CATEGORY	WHO	WHAT THEY LEAVE OPEN
General assistants	Chat with intelligence	Sell intelligence and leave the assembly, the taste, and the finishing to the user.
App builders	Plain-language in, one app out	Return a single artifact, not the whole agency and consulting deliverable.
Media tools	Finished images, audio, video	Return one medium, not the integrated, finished engagement.
Agencies and consultancies	Finished work by people	Priced against headcount and rate cards, weeks of lead time, cannot reach software economics from inside the model.

Kaldune is the only one selling the whole finished thing, from one plain-language brief, *at software economics.*

The four risks, named before you ask.

Key-person concentration

A solo-founder operation. Named as the sharpest risk. Mitigated by the transferable, documented, runbooked architecture that makes the company acquirable, and by the zero-employee design being deliberate rather than accidental.

Vendor dependency and rate limits

The engine runs on outside model and media vendors. Mitigated by the vendor-agnostic architecture: every model and vendor is a config edit, and the bottleneck moves with spend, not code.

Unproven acquisition cost

Cost per paying customer is not yet measured at scale. This is exactly what the test-then-scale distribution program and the raise exist to answer, measured before it is scaled.

Category education burden

“An agency that is AI” is a new category, and new categories take explaining. Mitigated by finished work as the proof: the deliverable is the pitch, and the demo film is itself the ad.

Six U.S. utility patent filings. Docket, title, status.

Six U.S. utility provisional patent applications by inventor Hicham Abdessamad, assignee Studio58, Inc.

DOCKET	SHORT TITLE	STATUS
KALDUNE-P1	Multi-agent work delivery: filesystem-canonical state, fact-gating, cache-aware context	Filed Aug 4, 2026 · App. 64/125,281
KALDUNE-P2	Deterministic verification of generative media output	Filed
KALDUNE-P3	Enforcement of client instructions over generative agents	Filed
KALDUNE-P4	Financial control of concurrent nondeterministic AI workloads	Filed
KALDUNE-P5	Multi-node custody, durability, and failover for stateful AI engagements	Filed
KALDUNE-P6	Safe execution, editing, packaging, and publication of AI-authored deliverables	Filed

The filings publish the mechanisms. The constitution, the prompts, and the calibration behind every threshold stay trade secret. Two layers, one moat.