

Kaldune

THE FINISHED-WORK LAYER · A PRODUCT OF STUDIO58, INC.

THE REVENUE PLAN · MMXXVI

Zero to ~\$25M in twenty-four months.

A driven revenue model for the finished-work layer, shown with its aggression on the page: the base case, the conservative reading, and the North Star, all one machine at different dial settings.

\$24.7M

MODELED ARR AT
MONTH 24

74.6%

ENTERPRISE SHARE
OF THE BOOK

175%

NET REVENUE
RETENTION, THE HERO
LEVER

BASE CASE CONFIRMED BY THE FOUNDER · \$50M MARKED AS THE STRETCH
EVERY ENGINE FIGURE MEASURED 2026·08·20 · COMPARABLES RE-CONFIRMED 2026·08·26

K

HOW TO READ THIS PLAN

One model, read three ways.

The endpoint is not a wish typed at the top of the page. It is a driven output of a small set of top-line dials, and the same machine reads at \$15M, \$24.7M, or \$50M depending only on where those dials are set. This plan shows the settings.

Kaldune sells finished work, not the tools that make it. That single fact is what this plan is built to monetize, and it is why the revenue math does not behave like any other AI company's. A plain-language brief goes in; a finished deliverable comes out, produced by an integrated AI team the client watches work; and the cost of producing it is compute, not payroll. Revenue can compound while the cost base does not have to.

The plan has one spine, and it is the customer architecture: three revenue motions that behave like three different businesses. The tourists fill the funnel, the plans hold the middle, and the enterprise base compounds into the moat. What follows walks that transformation month by month, states every assumption with its range, and marks the ambition as ambition rather than dressing it as a forecast.

Read the three numbers in one breath. **\$25M by month 24 is the target**, shown with its aggression on the page. **\$50M is the North Star**, marked for what it is and never drawn as a gate. **\$15M is the same machine read conservatively**. All of it is one model, and every figure in it traces to the model of record.

THE THREE READINGS

READING	M24 ARR
Conservative	~\$15M
Base case	\$24.7M
North Star	~\$50M

WHAT THE PLAN CONTAINS

- 1 • **The thesis.** Why finished work scales like software.
- 2 • **The customer architecture.** Three motions, one engine.
- 3 • **The ramp.** The first million is the hardest dollar.
- 4 • **The milestone gates.** Five gates, one North Star flagged.
- 5 • **The hero lever.** What net retention is worth.
- 6 • **Assumptions & sensitivity.** The aggression, shown.
- 7 • **The comparables.** The pattern is the argument.
- 8 • **Easy to miss.** Eight consequences of the model.

ARR means gross annual run-rate revenue, $MRR \times 12$ — the founder's explicit definition, no net-of-anything. Every Kaldune-internal figure is measured 2026-08-20; every public comparable was read on the owner's or reporting source's own page, re-confirmed 2026-08-26. Numbers trace to the model of record, Assumptions sheet.

01

A new category, priced like software, sold as *finished work.*

The ambition and the discipline, in one breath, because a serious reader will test both.

SECTION ONE · THE THESIS

Revenue that behaves like *software*, on a market built of people.

Everything from a freelance marketplace to a global consultancy sells the same underlying thing: a finished deliverable, produced by scarce human experts, priced by their hours. Kaldune sells the result and prices the compute. That is the whole plan, stated once.

The services market is more than one and a half trillion dollars a year, and it has never been available at software economics, because it has always been chained to people. A plain-language brief goes in. A finished deliverable comes out, produced by an integrated AI team the client watches work, and the cost of producing it is compute, not payroll.

The builder aisle one layer down sells power to people who want to build. Kaldune sits one layer up and sells the result to people who want to have it. The market does not yet have a clean name for that layer. The positioning brief names it the **finished-work layer**, and it is the empty top of the AI stack.

The reason this can scale like software when the incumbents cannot is not that Kaldune is cleverer with models. It is that Kaldune does not carry the headcount that defines their cost structure. A services firm that wants to grow twenty percent hires roughly twenty percent more people, drawn from the same scarce talent pool every competitor is fighting over. Kaldune's cost of goods is compute, and software costs nearly the same to serve the tenth client as the first.

That is the structural fact underneath every line of this plan: **revenue can compound while the cost base does not have to**. The plan that follows is the shape of that compounding, month by month, with the aggression shown rather than hidden.

If you sell tools, you compete on model and speed. If you sell the work, every price cut below you is margin.

WHY THE ECONOMICS RUN RIGHT · THE POSITIONING THESIS

Market sizing: consulting \$1.06T–\$1.1T (Research and Markets; IBISWorld), agencies \$400B–\$455B (IBISWorld; TBRC), freelance \$6.4B–\$10.3B (Grand View Research; SkyQuest). Consulting growth: Deloitte Global, FY2023 revenue announcement. The reachable market is not a slice of \$1.5T; it is that number plus everyone the price floor left out.

\$1.5T

Spent every year on finished work by scarce human experts — the market the finished-work layer sits on.

+19.1%

The largest consulting firm's fastest-growing line, in the year generative AI broke into the mainstream. Services grew straight through the wave.

Empty

The top layer of the AI stack, sitting directly on that market, with no named, funded player in it.

SECTION TWO · THE CUSTOMER ARCHITECTURE

Three motions that behave like *three different businesses*.

The plan has one spine, and it is the customer architecture. Kaldune earns revenue through three motions, and they do not blend into a single funnel. They run on one engine and each does a different job.

MOTION 1 — ENTERPRISE

The moat, and where the ARR base lives.

Land small, expand inside the account. A new logo arrives as a pilot department: about **8 seats at \$250**, roughly \$2,000 a month at entry. Then it grows over its life toward the founder's stated target of ~20 seats at ~\$5,000 a month.

That walk is the compounding engine of the entire plan. The enterprise motion is slow to start and relentless once it starts — the last motion to show revenue, and the first to matter.

MOTION 2 — PLANS

The self-serve wallet ladder, the steady middle.

The published subscriptions: **Studio \$199, Growth \$549, Agency \$1,490, Teams \$179/seat** (10-seat minimum). A plan is a wallet the client draws against, not a paywall; each tier grants monthly studio credit at par.

New subscribers arrive each month, most landing at Studio and climbing. The model reads the mix as one blended figure of ~\$465 a subscriber a month — a real recurring base, not the snowball.

MOTION 3 — PROSUMER

The tourists, the top of the funnel.

Pay-as-you-go from \$25 packs, revenue from the very first brief, and high monthly churn because tourists come and go by design. The model runs the base at **55% monthly churn**, which caps the active base even as new users climb.

Its job is not retention; its job is volume and conversion. A fraction climb into Plans every month, and Plans feed the enterprise top of the funnel. The widest mouth of the machine.

THE MATH, IN ONE LINE EACH

Enterprise MRR = last month × monthly NRR, plus new logos × seats × \$/seat. An installed-base recursion.

Plans MRR = active subscribers × blended \$/sub, net of 6% monthly churn.

Prosumer MRR = active users × ~\$45/user, drained by 55% churn each month.

The mix moves over the ramp, and the movement is the story of the next two sections. The early months are prosumer- and plans-heavy, because enterprise takes time to land and even longer to compound.

By month 24 the snowball has done its work and enterprise is the dominant share of the book, **74.6 percent of monthly recurring revenue**. The plan shows that shift at every month so the reader can watch it happen. Read in one line: the tourists fill the funnel, the plans hold the middle, and the enterprise base compounds into the moat.

SECTION TWO · THE MIX SHIFT

This is not an enterprise business. It *becomes* one.


The clearest way to see the whole plan is to watch the mix shift, because the shift is the argument in one moving picture. The book starts as tourists and plans; the moat is built underneath while the tourists fund the first year.

Early, the book is tourists and plans. By month 24 enterprise is 74.6 percent of monthly recurring revenue, plans and prosumer together making up the rest. This is a company funded in its first year by the tourists and the wallet ladder while the moat is still being built.

Read it as that transformation, not as a single line climbing to a number. **The number is where the transformation arrives, not what it is.**

~Month 13

Enterprise crosses 50% of the mix — the moment the compounding motion becomes the majority of the business.



Revenue mix by segment across the ramp. Enterprise is 0% of the book in months 1–2, ~24% by month 6, crosses 50% around month 13, and reaches 74.6% by month 24. Each bar sums to 100% of that month's MRR.

Source: the model of record, monthly mix output. Base-case dials.

ENTERPRISE, MONTH 24

74.6%

The moat, built from small lands compounding.

PLANS, MONTH 24

17.5%

The steady middle, the wallet ladder.

PROSUMER, MONTH 24

7.9%


The tourists, still filling the funnel.

A driven output of the three-motion model at the locked base-case dials; softening a driver changes both the endpoint and the month enterprise crosses 50%.

SECTION THREE · THE RAMP

The first million is the *hardest dollar* in the plan.

The honest way to read this ramp is to refuse the hockey stick. The model does not let the first six months look easy, because they are not.



Total ARR by month, 24 months, on an honest zero-based axis. Milestones flagged where first crossed: \$1M at month 6, \$10M at month 17, \$15M at month 20. The curve reaches \$24.7M at month 24; \$20M is crossed at month 22. The flat early months are the grind, not a drafting choice.

Source: the model of record, consolidated ramp. Base-case dials.

Enterprise logos land at near zero through the grind: zero in month one, zero in month two, then a trickle of one, two, three, four across months three through six. The snowball has no base to compound yet, so revenue is carried almost entirely by early prosumer volume and a thin line of plans.

The first million dollars of annual run-rate revenue is crossed at month six, and it is the hardest dollar in the whole plan. Part of that grind is physics, not slowness: SOC 2 is the gate enterprise procurement waits on, and it opens partway through year one.

THE S-CURVE, NOT THE STICK

Logo velocity accelerates toward ~49 new logos a month at exit — about 564 enterprise logos in total.

But the mid-ramp bend is mostly expansion, not new logos. The logos already landed grow month after month. That is Section 5.

Every dollar after the first million is easier than the one before it; every dollar before it is the company proving the machine can sell at all — so test the grind, not the peak. Enterprise logo ramp (months 1–6): 0, 0, 1, 2, 3, 4. Enterprise is 0% of the mix in months 1–2 and ~24% by month 6.

SECTION FOUR · THE MILESTONE GATES

Five gates on the base curve, and one North Star *flagged for what it is.*

The spacing between the gates tells the truth about the plan. The years are not symmetric — and asymmetry, acceleration late, is exactly what a compounding business looks like.


THE FIVE GATES ON THE BASE LINE

MILESTONE	MONTH FIRST CROSSED
First \$1M ARR	Month 6
\$10M ARR	Month 17
\$15M ARR	Month 20
\$20M ARR	Month 22
~\$25M ARR (\$24.7M)	Month 24

THE NORTH STAR, PRICED
HONESTLY

\$50M by month 24 is not on this line. Reaching it asks two hard levers to run hot at once: **logo velocity +35%** (65+ logos/mo) **and net NRR walking to 195%/yr**, sustained for two years from a standing start.

The model shows it is reachable. The plan refuses to dress it as a forecast — a North Star quietly promised is a number a diligence room uses against the whole document.



Months to cross each successive gate. Zero to one takes six months; one to ten takes eleven; the last three gates take two to three months each, because by then the snowball is carrying the weight.

Source: the model of record, milestone crossings.

The distance from zero to one takes six months. From one to ten takes eleven more. From ten to twenty-five, only seven. Acceleration late is the proof it is compounding rather than being pushed.

\$15M, reached around month 20 on the base curve, is both a milestone the base case passes and the natural endpoint of the same model at softer dials (net NRR ~145%, softer velocity). It is the conservative reading, and it is not a different plan — it is the same machine at a different setting, which is why it belongs on the page beside the base case rather than hidden from it.

\$25M is the target. \$50M is the ambition, marked. \$15M is the same machine read conservatively. The discipline is the valuation defense's own: state the aggressive figure as aggressive, show the machine that produces it, and let the reader price it.

05

Net revenue retention is the *snowball*.

If a reader takes one number away from this plan, it should be this one.


SECTION FIVE · THE HERO LEVER

A zero-logo month *still grows* enterprise MRR 4.7%.

That is not a projection about sales effort. It is arithmetic. At 175% net revenue retention a year, the monthly compounding rate is 1.047 — so the installed base expands on its own every month whether or not a single new account signs.

+4.7%

Enterprise MRR growth in a month with **zero** new logos, at 175% net NRR. The snowball, quantified.



The most important chart in the document. One cohort's own walk: 8 seats at land (~\$2,000/mo) to 20+ seats at maturity (~\$5,000/mo). This one curve explains the entire snowball.

Source: the model of record, single-cohort expansion at 175% net NRR.

The instinct with a plan like this is to sell it on logo count. The model says the opposite. At 175 percent net retention, the accounts already landed out-produce the accounts still to be added.

A single pilot department that lands at eight seats and walks to twenty-plus is worth more over the plan than several new logos that sign and never expand. This is why the whole strategic case leads with retention rather than acquisition. Net revenue retention compounds the enterprise base roughly threefold over twenty-four months, and that compounding, not the logo velocity, is what bends the ramp upward in its second year.

The sensitivity math makes the point unarguable. Enterprise revenue scales roughly **linearly** with logo velocity for a fixed retention shape. It scales roughly **exponentially** with retention over twenty-four months. So a twenty-point swing in annual net retention moves the month-24 endpoint more than a proportional change in velocity does.

THE HERO IS ALSO THE HONEST RISK

The whole \$25M reading rests on the snowball actually running at 175%. If it runs at 145%, this is a \$15M plan.

So the go-to-market's single most valuable activity is not filling the top of the funnel. It is making the accounts already inside it expand.

WHY THE EXPANSION HEADROOM IS STRUCTURALLY AVAILABLE HERE

Land is a fraction of the account

A logo enters at ~8 seats against a department that will hold 20+. The room to expand is built into how small the land is by design — the account is nowhere near saturated when it signs.

The product spreads by use

Finished work is consumed across a team, not gated to one seat. As more of the department briefs the studio, seats and spend rise on their own — the expansion is pulled, not pushed.

Net already absorbs the churn

The 175% is net of the accounts that leave and contract. The gross expansion behind it runs higher, so the plan quotes the conservative framing, not the flattering one.

SECTION SIX · THE ASSUMPTIONS

The base case sits at the aggressive end of *every range at once.*

A diligence room will find this if the plan does not say it, so the plan says it first. Credible and aggressive does not mean splitting the difference. It means putting the aggression on the page where the reader can see it and price it.

THE DRIVER SET · BASE VALUE, RANGE, AT THE AGGRESSIVE SPINE

DRIVER	BASE	RANGE
New enterprise logos/mo (exit)	49	25–65
Net revenue retention (annual)	175%	135–195%
Starting seats per logo	8	4–12
\$/seat/mo (enterprise)	\$250	\$200–300
New plan subscribers/mo (exit)	~85	40–130
Blended \$/plan subscriber/mo	\$465	\$350–560
New PAYG users/mo (exit)	~2,000	600–3,500
Avg spend/user/mo (PAYG)	\$45	\$25–70
Monthly PAYG churn	55%	40–70%
Blended gross margin	~78%	64–85%+

The \$25M base case runs net retention at 175 percent, a strong logo-velocity ramp, and a real pilot-department land size, all hot together for twenty-four months. That is the honest cost of \$25M by month 24.

Each driver carries a base value and an honest range. Where a false-precise number would be a lie, the range is the answer, and the simulator exposes every one as a lever so the reader can move them and watch the endpoint recompute.

GROSS MARGIN, MEASURED BY LANE


Entry ~64%, Craft 70%+, Professional 85%+. Blended by revenue mix, rising as the mix shifts to higher-tier work and as model prices fall — every price cut flows straight to margin.

Enterprise logo churn is modeled inside net NRR as negative expansion, not a separate line, so the retention figure is net. 175% net NRR is high against public benchmarks (best-in-class enterprise SaaS ~120–140%); it is credible only for a young company landing small in accounts with large expansion headroom — which is exactly this land-and-expand thesis.

SECTION SIX · THE SENSITIVITY BAND

One machine, three dial settings. *Both readings are honest.*

Run the dials at their base and the model lands \$24.7M. Soften net retention to about 145% and run velocity lower, and the same model lands near \$15M. That is not pessimism versus optimism — it is one machine at two settings, both inside the ranges the model already publishes.



Month-24 ARR at three dial settings of the same model. Conservative (\$15M), base case (\$24.7M), and the North Star stretch (\$50M). The North Star is marked as the stretch, never drawn as a gate the base case crosses.

Source: the model of record, Sensitivity sheet.

WHERE THE DIALS LAND THE ENDPOINT

ENDPOINT	LOGOS/MO	NET NRR
~\$15M	~25	145%
~\$25M base	~49	175%
~\$50M stretch	65+	195%

The \$15M reading is what the plan looks like if the aggressive assumptions do not all hold at once — the honest thing to know before committing to the \$25M reading.

Presenting both is not a weakness in the argument. **It is the argument's credibility.** A plan that can only be read one way is a plan that is hiding its own sensitivity.

WHY THE TWO LEVERS MOVE THE ENDPOINT DIFFERENTLY

LOGO VELOCITY — LINEAR

Add more logos and enterprise revenue rises roughly in proportion, for a fixed retention shape. Double the velocity and you roughly double the new-logo contribution — a straight-line relationship in the endpoint.

Velocity is what a channel machine delivers, and the payback discipline governs how hard it can run.

NET RETENTION — EXPONENTIAL

Retention compounds. Over twenty-four months a 20-point swing in annual net NRR moves the month-24 endpoint **more** than a proportional change in velocity does, because the installed base multiplies on itself every month.

That is why NRR is the hero lever — and the first slider on the simulator.

Both levers stay inside the ranges the model publishes; the sensitivity band is the same machine read at both ends of the credible set, never a separate optimistic and pessimistic plan.

SECTION SEVEN · THE COMPARABLES

The pattern is the argument, *not any single number.*

Capital has concentrated enormously in the builder aisle one layer below the finished-work layer, and no named, funded player occupies the layer above. Every figure below was followed to the source that owns or reports it and read there. Disclosed prints as disclosed; nothing is estimated to fill a gap.

THE COMPARABLES, EACH FIGURE WITH ITS VERIFICATION STATUS

PLAYER	FIGURE	STATUS	READS AS
Lovable	\$400M Series C at \$13.3B , 12 Aug 2026; crossed \$500M ARR (Jun 2026), projecting \$600M	DISCLOSED	A transactional, low-friction motion whose run-rate more than doubled since the prior round.
Anthropic	\$47B annualized run-rate (Series H disclosure) → ~\$65B end-July 2026	DISCLOSED + REPORTED	The character of the revenue: machines doing work, ~80% from business customers.
Midjourney	~\$500M ARR , \$0 venture capital , no free tier, profitable within six months	REPORTED EST. + FOUNDER	Kaldune's own no-free-tier, revenue-from-first-brief model echoes the shape directly.
Manus	\$75M led by Benchmark at ~\$500M , Apr 2025	REPORTED	An agent-orchestration player — the aisle, not the finished-work layer.
Vertesia	Seed ~\$4M (Elaia, Illuminate Financial), founded 2024	REPORTED	A landscape comp in the aisle, named for positioning honesty — never a valuation anchor.

HOW TO READ THE STATUS

- DISCLOSED** the company's own page.
- REPORTED** dated reputable coverage, not a company disclosure.
- REPORTED EST.** no owner page publishes it; a reported estimate.
- FOUNDER** founder-confirmed, publicly and repeatedly.

No single figure here is the argument, and the plan does not lean on any one of them. The argument is the pattern: money pouring from the hundreds of millions into the tens of billions in the builder aisle, and an empty layer directly above it.

An empty layer with no capital around it is a warning. **An empty layer sitting on top of the most heavily funded corner of the market is the definition of being early to the right place.**

Lovable: lovable.dev/blog/series-c and Menlo Ventures' own post. Anthropic: \$47B via CNBC (Series H); \$65B via Bloomberg/TechCrunch, end-July run-rate. Midjourney: Contrary Research; founder David Holz on the \$0-VC facts. Manus: Bloomberg/TechCrunch. Vertesia: vertesiahq.com. Re-confirmed 2026-08-26.

08

The things that are *easy to miss*.

Eight strategic consequences of the model — the things the numbers know that a first reading of the numbers misses.

SECTION EIGHT · EASY TO MISS · ONE THROUGH FOUR

Each item is a *decision* the plan already makes.

Each one turns a modeling fact into a strategic instruction. Read together, they are the operating logic of the plan.

8.1 THE ENGINE IS RETENTION, NOT LOGO COUNT

At 175 percent net retention the installed base out-produces the new logos. A pilot department that walks to twenty seats is worth more than several logos that sign and never grow. Every diligence conversation should lead with net retention — it is the difference between a linear business and a compounding one, and it is the plan's single biggest risk. Leading with it is the honest move rather than the confident one.

8.3 SOC 2 IS THE GATE THAT LETS VELOCITY ACCELERATE

The enterprise ramp is near zero for the first months partly by physics. Procurement often will not sign without SOC 2, and SOC 2 is an explicit use-of-funds line rather than a certification in hand. A seeded trial that cannot convert to a signed contract is not revenue. The grind months are grind partly because the door is still locked; funding SOC 2 is what converts trials into contracts, not overhead.

8.2 THE FLOAT IS THE COST OF WINNING

Unit economics are positive on every deliverable (64–85%+ GM), so the company does not burn capital on the work. It burns on timing: compute is prepaid while enterprise pays net-30 and net-60. The model carries a float line of ~45 days of enterprise revenue — and the float **grows as enterprise succeeds**. That is why the raise is sized against working capital, not survival. The float is a feature of the good outcome, not a warning.

8.4 THE TOURIST MOTION IS SELF-FUNDING LEAD GEN

Prosumer PAYG looks like a high-churn afterthought at 55% monthly churn. It is the opposite. Every prosumer is revenue from the first brief, so payback is immediate, measured from day one. A fraction convert up into Plans, and Plans feed the enterprise top of the funnel. Do not underwrite prosumer on its own retention — underwrite it on conversion into the ladder, where it is doing the most important job in the acquisition engine.

Each consequence is a real modeling fact, not a talking point — the section a founder asked for by name, and it earns its length.

SECTION EIGHT · EASY TO MISS · FIVE THROUGH EIGHT

The consequences the numbers *know*.

Continued: the seat-expansion end state, the payback discipline, the net-versus-gross honesty, and the concentration that eases as the base broadens.

8.5 THE \$5K ACCOUNT IS WHERE A LOGO ARRIVES, NOT LANDS

The ~\$5,000-a-month, twenty-seat account is an end state. A logo starts at ~8 seats and ~\$2,000/mo, a pilot department, and expands from there. If the plan is read as landing \$5K accounts, year one looks like failure. Read correctly — land small, expand relentlessly — the first year is *supposed* to look like small lands, and the value shows up in year two as they compound.

SEE IT ON ONE CURVE

One cohort's own expansion (Section 5) is the single chart that explains the entire snowball — and why small early numbers are the plan working, not struggling.

8.7 THE 175% IS NET, AND NET IS THE STRONGER NUMBER

The retention figure the plan rests on is net — expansion minus churn and contraction, per account — so it already absorbs the losses. Gross expansion runs materially higher. A sharp investor will ask which one it is, and the honest answer is the stronger one: net retention is the conservative framing of the snowball, and it still compounds the base ~3× over twenty-four months. The plan leads with why that expansion headroom is structurally available here rather than asserting the number and hoping it holds.

8.6 PAYBACK DISCIPLINE KEEPS THE VELOCITY HONEST

The demand-generation spend is not sprayed. The doctrine is the founder's own: measure cost-per-paying-customer and payback per channel, scale only what clears a short bar, cut what does not. Because every user is revenue from the first brief, payback is measured from day one. The base case asks a channel machine to deliver ~564 enterprise logos and ~2,000 PAYG users a month by month 24 from a standing start — the aggressive assumption made concrete, and the payback discipline is what keeps it honest. Distribution is the one open risk the founder names.

8.8 CONCENTRATION EASES AS THE BASE BROADENS

Early enterprise revenue is concentrated in a few large accounts, and those same accounts are on net-30/60 terms — so concentration and slow-pay compound each other in the grind year. But the risk is not static. Concentration falls as the base broadens, with more logos, more plans, and more prosumer conversion. The plan shows concentration declining over the ramp rather than treating it as a fixed hazard, because the shape of the risk over time is the honest thing to show.

Best-in-class enterprise SaaS net retention sits ~120–140%; 175% is credible only for a young company landing small in accounts with large expansion headroom, which is exactly the land-and-expand thesis — and the number a diligence team will push hardest on.

Kaldune

THE REVENUE PLAN · IN CLOSING

THE PLAN IN ONE LINE

The tourists fill the funnel, the plans hold the middle, and the enterprise base *compounds into the moat.*

The first million is the hardest dollar; the snowball does the rest; and net revenue retention is the number the whole thing turns on. \$25M by month 24 is the target, shown with its aggression on the page. \$50M is the North Star, marked. \$15M is the same machine read conservatively. All of it is one model, and every figure can be moved and watched to recompute — the difference between a forecast to be believed and a machine to be tested.

\$24.7M

BASE CASE, MONTH 24

~\$15M

CONSERVATIVE
READING

~\$50M

NORTH STAR, THE
STRETCH

KALDUNE, A PRODUCT OF STUDIO58, INC.

BASE CASE CONFIRMED BY THE FOUNDER · \$50M MARKED AS THE STRETCH, NEVER A FORECAST

K

COLOPHON · ON THIS PLAN

Every number has *lineage*.

One model, read three ways, with the aggression on the page and the receipts under every comparable. This is what a plan that means to be both credible and aggressive looks like.

The endpoint is a driven output, not a hard-coded wish. The 24-month ramp is what a small set of top-line dials produces when set at the credible-aggressive end of their ranges and held there for two years. Change a driver and the ramp, the mix, and the milestone months recompute.

Every Kaldune-internal figure in this plan is measured from the repository on **2026-08-20**. Every public comparable was followed to the source that owns or reports it, read on that page, and re-confirmed on **2026-08-26**. Disclosed prints as disclosed, reported as reported, and nothing is estimated to fill a gap. Where a false-precise number would be a lie, the range is the answer.

The base case is the founder's confirmed target, \$24.7M at month 24. \$50M is his North Star, shown as a stretch overlay at a stated cost and never as a gate. \$15M is the same machine read conservatively. The discipline throughout is the valuation defense's own: state the aggressive figure as aggressive, show the machine that produces it, and let the reader price it.

SOURCES OF RECORD

The model — the revenue workbook, Assumptions and Sensitivity sheets. Every figure reads from the assumptions via live formulas.

The argument — the plan's committed prose, reconciled to the model.

The positioning thesis and moat facts — the canonical brief, engine figures measured 2026-08-20.

The comparables — each on the owner's or reporting source's own page, re-confirmed 2026-08-26.

THE DATING IS PART OF THE ARGUMENT

This system moves fast enough that any printed number is already conservative. The staleness is itself a moat fact.

Set in Fraunces and Archivo. ARR means gross annual run-rate revenue, $MRR \times 12$. Kaldune, a product of Studio58, Inc.