

MUDITA STUDIOS RESEARCH — THE STUDIO SERIES

The Next Decade of Company Building

AI-Native Studios

AI makes software cheaper to write, so a studio's engineering bench is no longer its edge. What makes studios the biggest winners is everything they can share around the code.

PUBLISHED
October 2026

BY
Mudita Studios Research

LENGTH
21 min · 7 exhibits

EXECUTIVE SUMMARY

We assessed how AI is changing company building, from revenue per employee to coding productivity trials, and modelled what it does to a studio's cost per MVP. The analysis shows studios, founders and partners where AI gains are real and where they are capped.

$2.4M

median self-reported ARR per employee at the AI-native companies we tracked

Smaller teams carry more revenue

That compares with $437K for 29 US-listed SaaS companies in their latest 10-K filings. Every AI-native figure is self-reported, and compute partly replaces payroll.

24 mo

for the top 100 AI companies on Stripe to reach $5M in annualized revenue

Faster to revenue, though not unprecedented

The top SaaS cohort of 2018 took 37 months. Most AI-native leaders say they reached $100M ARR in about 20 months, roughly what Wiz and Deel did before generative AI.

36%

solo founders' share of new startups in 2025, up from 17% in 2017

Fewer founders, smaller teams

The average seed-stage company on Carta has 6.2 equity-holding employees, down from 10.3 in 2021. Carta's data cannot show that AI caused the shift.

+19%

change in task time for experienced developers using AI (METR, 2025)

The productivity evidence is mixed

Other randomized trials range up to 56% faster, on a greenfield task. The vendor estimate of an 80% time saving is not trial evidence.

6–24%

modelled fall in a studio's cost per MVP on the realistic trial range

Coding gains alone are capped

Even an 80% coding speedup cuts cost per MVP only about 37%, because discovery, design, legal and go-to-market work are untouched. The 50% engineering share of cost is our assumption.

01 1. Company building has a new production function

The clearest signal of the AI era is how much revenue a small team now carries. We built a benchmark from primary filings: revenue over total headcount for 29 US-listed SaaS companies, from Salesforce to SentinelOne, using each company's latest Form 10-K. The median is **$437K per employee**, with an interquartile range of $385K–$494K. That matches Benchmarkit's index of 134 public SaaS companies ($395K in 2025).

The AI-native companies sit on a different curve (Exhibit 1). Lovable reported **$400M ARR with 146 employees** in March 2026, about $2.7M per head. When it crossed $100M ARR in July 2025 it had 45 full-time staff. Gamma crossed $100M ARR profitably with about 50 people. Midjourney was reported at $200M of revenue with about 40 staff as early as 2023.

EXHIBIT 1

AI-native leaders report 2.6-11x the revenue per employee of public SaaS

Revenue or ARR per employee. AI-native figures are company-reported ARR over reported or estimated headcount

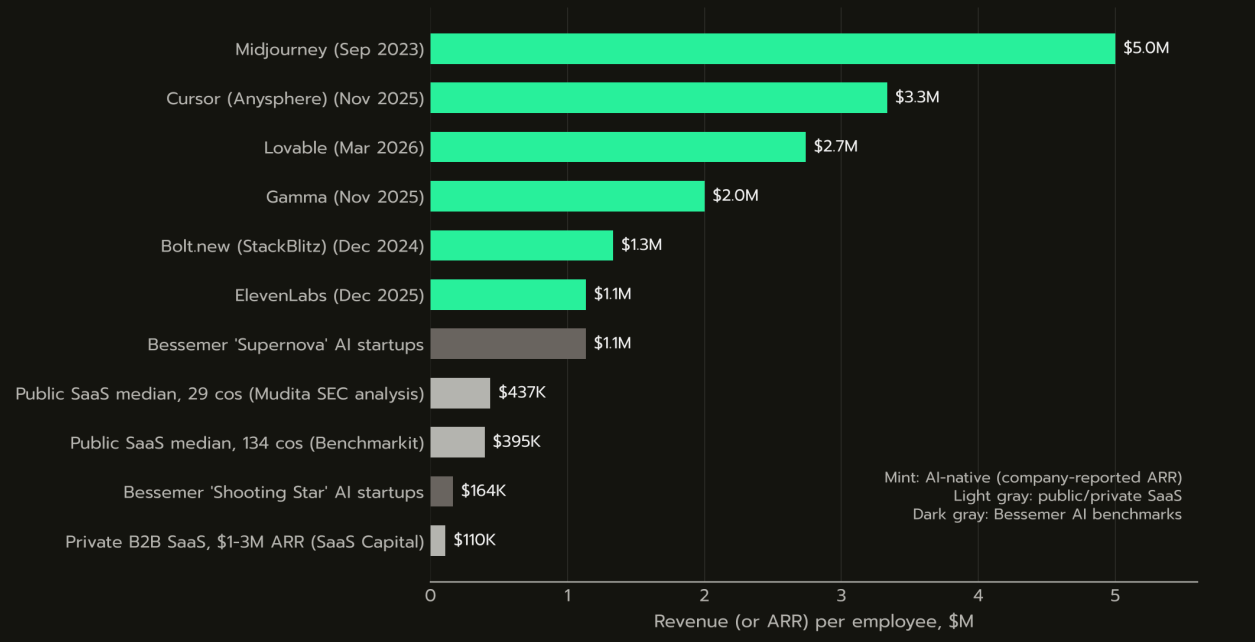

Source: Mudita Studios analysis of 29 SaaS 10-K filings (SEC EDGAR, FY2025-26); company announcements and press (Lovable, Cursor, Midjourney, ElevenLabs/Revelio Labs, Gamma, Bolt); Benchmarkit (2025); Bessemer State of AI (2025); SaaS Capital (2026).

MUDITA STUDIOS

Table 1: Every AI-native data point beats the public SaaS median ($437K), but all rest on company-reported ARR

Company	As of	Revenue/ARR	Employees	Per employee	vs public SaaS median	Reliability (revenue / headcount)	Sources
ElevenLabs	2025-12-31	$330.0M	291	$1.1M	2.6x	self-reported / independent-data	ElevenLabs; Revelio Labs
Bolt.new (StackBlitz)	2024-12	$20.0M	15	$1.3M	3.1x	self-reported / anecdote	Stripe; AI Engineer (YouTube talk)
Gamma	2025-11-10	$100.0M	50	$2.0M	4.6x	self-reported / self-reported	Gamma; SaaStr
Lovable	2025-07-23	$100.0M	45	$2.2M	5.1x	self-reported / self-reported	Lovable; TechCrunch
Lovable	2026-03-11	$400.0M	146	$2.7M	6.3x	self-reported / self-reported	Business Insider; TechCrunch
Cursor (Anysphere)	2025-11-18	$1000.0M	300	$3.3M	7.6x	self-reported / anecdote	SaaStr
Midjourney	2023-09-06	$200.0M	40	$5.0M	11.4x	self-reported / self-reported	The Decoder

Source: company announcements and press as listed; public SaaS median from Mudita Studios analysis of 29 Form 10-K filings (SEC EDGAR). ARR is a run-rate, not GAAP revenue; headcounts are point-in-time and sometimes estimated.

Read these numbers with two caveats.

- **They are self-reported.** ARR is a run-rate that companies choose when to announce. Headcounts are point-in-time and in some cases estimated: we used Revelio Labs for ElevenLabs and a rough "~300" for Cursor.
- **High revenue per employee is partly a cost substitution, not pure efficiency.** Bessemer's 2025 "Supernova" AI startups average $1.13M ARR per FTE but report gross margins of roughly 25%. Its more typical "Shooting Stars" run at $164K per FTE with 60% margins. Compute replaces payroll, so revenue per employee flatters the economics.

Speed to revenue tells a similar but more nuanced story (Exhibits 2 and 3). Stripe's platform data is the most independent source here, though it covers success-selected top-100 cohorts:

- The top 100 AI companies reached **$1M of annualized revenue in a median of 11.5 months**, about four months ahead of the fastest SaaS companies.
- They reached **$5M in 24 months**, against 37 months for the top SaaS cohort of 2018.
- Stripe's 2025 letter says the whole 2025 cohort of new businesses grew about 50% faster than 2024's. Twice as many reached $10M ARR within three months of launch.
- The share of Stripe Atlas startups that charged a first customer within 30 days rose from 8% in 2020 to 20%.

EXHIBIT 2

Most AI-native leaders hit $100M ARR in ~20 months, like the best pre-AI SaaS

Months to $100M ARR. Milestones are company-announced; start points differ (launch, first $1M ARR or founding)

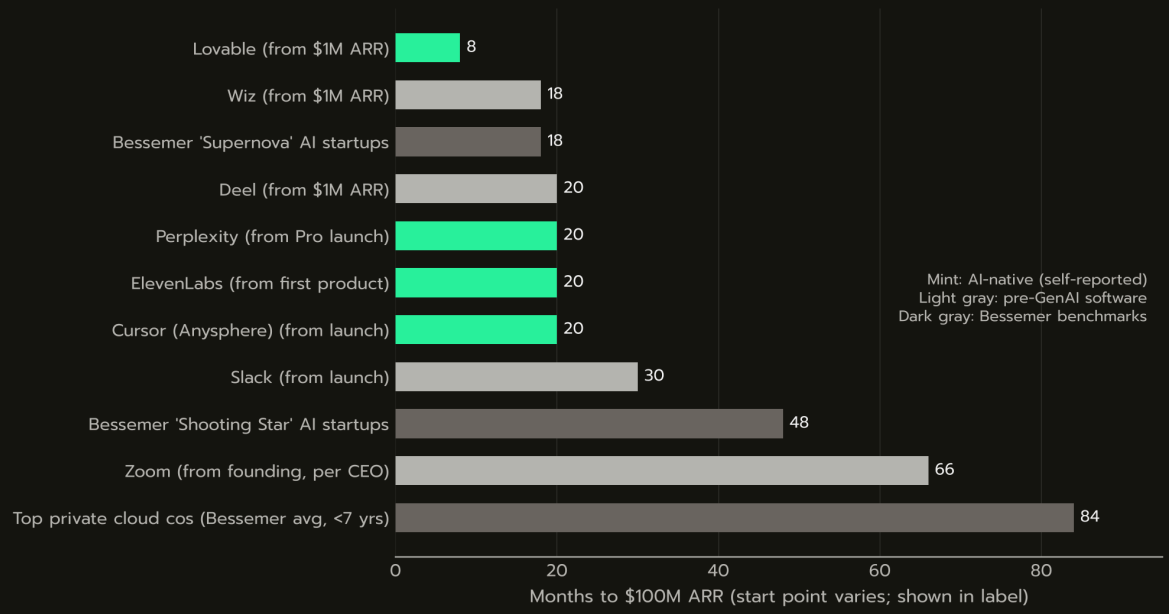

Source: Lovable (2025); SaaStr (Cursor, Deel); ElevenLabs and Perplexity CEO statements; TechCrunch (Wiz, 2022); Index Ventures (Slack, 2016); SaaStr interview (Zoom); Bessemer State of AI 2025. All company figures self-reported. Mudita Studios analysis.

MUDITA STUDIOS

EXHIBIT 3

On Stripe, top AI companies hit $5M revenue 13 months sooner than SaaS did

Median months to reach an annualized-revenue milestone; top-100 cohorts by revenue on Stripe

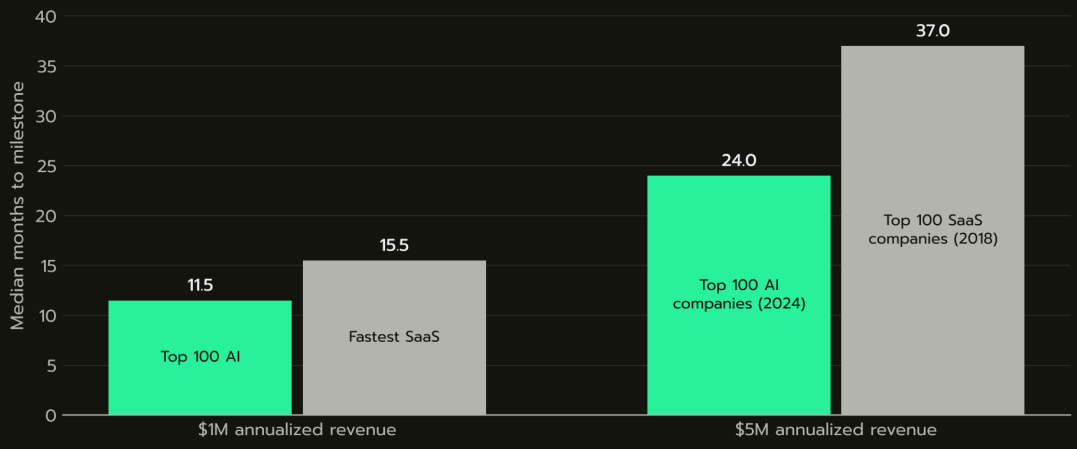

Source: Stripe, 'Inside the growth of the top AI companies on Stripe' (Jun 2025): $1M in 11.5 months, 'about 4 months ahead' of the fastest SaaS (15.5 derived); Stripe 2024 annual letter via TechCrunch (Feb 2025): $5M in 24 vs 37 months. Success-selected samples. Mudita Studios analysis.

MUDITA STUDIOS

The company-level race to $100M ARR is less dramatic than the headlines suggest. Cursor, ElevenLabs and Perplexity each say they got there in about 20 months. Wiz and Deel claimed 18 and 20 months before generative AI existed. AI has made the fast lane much more crowded, but the fastest pre-AI companies already used it. Lovable's eight months is the outlier.

Table 2: The AI-native speed records are real on Stripe's data, but the company-level milestones are self-reported

Entity	Metric	Months	Definition	As of	Reliability	Source
Lovable	months to 100m arr	8.0	From first $1M ARR to $100M ARR	2025-07-23	self-reported	Lovable
Lovable	months to 17m arr	3.0	From launch to $17M ARR	2025	self-reported	Stripe
Cursor (Anysphere)	months to 100m arr	20.0	From product launch to $100M ARR (Jan 2025)	2025-01	self-reported	SaaStr
Cursor (Anysphere)	months 1m to 100m arr	13.0	From $1M ARR (Dec 2023) to $100M ARR (Jan 2025)	2025-01	self-reported	SaaStr
ElevenLabs	months to 100m arr	20.0	From first product launch (2023) to $100M ARR, per CEO	2026-01-13	self-reported	Business Insider via Yahoo Finance
Perplexity	months to 100m arr	20.0	From launch of Perplexity Pro to $100M ARR, per CEO	2025-03-27	self-reported	American Bazaar (CEO LinkedIn post; Bloomberg)
Bolt.new (StackBlitz)	months to 20m arr	2.0	From launch to $20M ARR	2024-12	self-reported	Stripe
Wiz	months to 100m arr	18.0	From $1M ARR to $100M ARR (company claim)	2022-08-10	self-reported	TechCrunch
Deel	months to 100m arr	20.0	From $1M ARR to $100M ARR	2025-06-29	self-reported	SaaStr
Slack	months to 100m arr	30.0	From launch (Feb 2014) to $100M ARR (mid-2016)	2016	self-reported	Index Ventures
Zoom	months to 100m arr	66.0	'Zero to hundred million in five and a half years' (CEO interview)	2024-02-03	anecdote	SaaStr
Top private cloud companies (Bessemer average)	months to 100m arr	84.0	'reach $100M ARR in < 7 years' on average	2025-08-12	industry-survey	Bessemer Venture Partners
Bessemer 'Shooting Star' AI startups	months to 100m arr	48.0	~4 years to $100M ARR ($103M in year 4)	2025-08-12	industry-survey	Bessemer Venture Partners
Bessemer 'Supernova' AI startups	months to 100m arr	18.0	~1.5 years to $100M ARR	2025-08-12	industry-survey	Bessemer Venture Partners
Top 100 AI companies on	months to 1m	11.5	Median months to $1M	2025-06-30	independent-data	Stripe

Entity	Metric	Months	Definition	As of	Reliability	Source
Stripe	annualized revenue		annualized revenue			
Fastest-growing SaaS companies on Stripe (subscription boom)	months to 1m annualized revenue	15.5	Derived: Stripe says AI cohort was 'about 4 months ahead' of this group	2025-06-30	independent-data	Stripe
Top 100 AI companies on Stripe (2024)	months to 5m annualized revenue	24.0	Months to $5M annualized revenue	2025-02-27	independent-data	TechCrunch
Top 100 SaaS companies on Stripe (2018)	months to 5m annualized revenue	37.0	Months to $5M annualized revenue	2025-02-27	independent-data	TechCrunch

Source: as listed per row; Mudita Studios analysis. Stripe figures are platform data for success-selected top-100 cohorts.

02 2. Fewer founders, smaller teams

The second change is who starts companies. Carta's cap-table data shows the share of new startups with a single founder **roughly doubling**, from 17% in 2017 to about 36% in 2025 (Exhibit 4). Carta's three reports use different cohorts and do not agree exactly: the 2024 figure is 30%, 31% or 35% depending on the report. All three show the same direction, though.

Teams are shrinking too. The average seed-stage company on Carta has **6.2 equity-holding employees, down from 10.3 at the 2021 peak**. YC says about a quarter of its Winter 2025 batch had codebases that were ~95% AI-generated. Both YC's figure and its growth claims have been publicly challenged.

Two facts temper the "one-person unicorn" story.

- **Solo founders struggle to raise capital**. They were 30% of startups founded in 2024 but received only **14.7% of cash raised in priced equity rounds**, and only 17% of the companies that closed a VC round that year.
- **The data cannot show that AI caused the shift**. Carta's figures cover companies on its platform and also reflect a tighter funding market.

The pattern matters for studios. Founders now arrive with more product leverage and less of everything else: co-founders, capital, distribution and operating support. That gap is what a studio exists to fill.

> **14.7%** of cash raised in priced equity rounds went to solo founders, who were 30% of startups founded in 2024

03 3. The productivity question: evidence on both sides

If AI makes engineers dramatically faster, studio economics change dramatically. The best-controlled evidence is far more modest than the anecdotes (Exhibit 5).

The case for large gains:

- **GitHub Copilot lab RCT (Peng et al.)**. Developers given the tool finished a greenfield task, an HTTP server in JavaScript, **55.8% faster** (n=95).
- **Field trials at Microsoft, Accenture and a Fortune 100 firm (Cui, Demirer et al.)**. Pooled across 4,867 developers, completed tasks rose **26%**, with the biggest gains for less-experienced developers.
- **Google enterprise RCT**. Engineers took about **21% less time** on one enterprise task (n=96).

The case against:

- **METR's 2025 RCT**. Sixteen experienced open-source developers worked on their own mature repositories. With AI allowed, they took **19% longer** (95% CI +2% to +39%). Before the study they expected a 24% speedup, and afterwards they still believed AI had saved them 20%.
- **METR's late-2025 follow-up.** The point estimates moved toward speedup: 4% faster for new developers and 18% faster for returning ones. Both confidence intervals still include zero. METR says the results understate the gain, because 30–50% of developers withheld tasks they did not want to do without AI.
- **DORA's survey models.** In 2024 data, 25% more AI adoption was associated with **1.5% lower delivery throughput and 7.2% lower stability**. The 2025 report found throughput had turned positive but instability persisted.

The most bullish numbers come from vendors. Anthropic's staff self-report a 50% productivity gain (n=132). Its model-based estimate across 100,000 conversations puts task time savings at about 80%, and Anthropic notes this excludes the work of refining outputs.

Table 3: The best-controlled evidence on experienced developers is the least flattering

Study	Design	n	Setting	Time change	95% CI	Reliability	Notes
METR early-2025 RCT	RCT (task-level), preprint	16	Experienced OSS devs, own mature repos, Cursor + Claude 3.5/3.7	+19.0%	+2% to +39%	independent-data	Devs forecast -24% and believed -20% afterwards
METR late-2025 follow-up, new developers	RCT, METR flags selection bias	47	Newly recruited OSS devs, late-2025 tools	-4.0%	-15% to +9%	independent-data	METR says likely a lower bound on speedup
METR late-2025 follow-up, returning developers	RCT, METR flags selection bias	10	Original-study devs, late-2025 tools	-18.0%	-38% to +9%	independent-data	
Google enterprise RCT (Paradis et al.)	RCT (lab task, internal tools)	96	Google engineers, one enterprise-grade task	-21.0%	n/r	independent-data	'Confidence interval is large' (not reported here)
Field RCTs at Microsoft, Accenture, Fortune 100 (Cui, Demirer et al.)	Field RCTs pooled, working paper	4,867	Enterprise developers, Copilot autocomplete	-20.7%	-32% to -6%	independent-data	+26.08% tasks (SE 10.3pp); converted to time per task = 1/(1+effect)-1; CI from +/-1.96 SE
GitHub Copilot lab RCT (Peng et al.)	RCT (lab, greenfield task), preprint	95	Recruited devs writing an HTTP server in JavaScript	-55.8%	n/r	independent-data	Greenfield and well-scoped: closest analogue to an MVP
Anthropic internal survey	Self-reported survey (vendor)	132	Anthropic engineers/researchers	-33.3%	n/r	self-reported	+50% productivity self-reported; converted to time = 1/1.5-1
Anthropic conversation-level estimate	Model-estimated (vendor)	100,000	Claude.ai conversations, all task types	-80.0%	n/r	self-reported	Excludes refinement and rework; upper bound

Source: METR; Paradis et al. (Google); Cui, Demirer et al.; Peng et al.; Anthropic; Mudita Studios analysis. DORA (survey-based) is discussed in the text: +25% AI adoption was associated with -1.5% delivery throughput and -7.2% stability in 2024 data; the 2025 report found a positive throughput link but continued instability.

Our reading: **AI speeds up greenfield, well-scoped work the most, and complex work in mature codebases the least**. An MVP is close to the greenfield case, so studios should expect real gains at the build stage. Those gains shrink as a venture's codebase matures. Self-reported speedups are a poor guide, as METR's perception gap shows.

> **20%** time saving experienced developers believed AI gave them, after a trial that measured them 19% slower

---

04 4. What it does to studio economics: a model

We modeled a typical studio, starting from the most rigorous studio survey we could find: Inniches' 2024 survey of 123 studios.

Baseline from the survey:

- 36% of studios spend under $50K per MVP, 27% spend $51–100K and 37% spend more than $100K. The interpolated median is **$76K**.
- About half reach an MVP in 3–4 months (interpolated median **4.0 months**).
- 59% start just **1–3 ventures a year**.
- On average, a studio needs **200 ideas and 40 MVPs per exit**.

Model assumptions:

- Software engineering is **50% of MVP cost**, tested from 30% to 70%. This is an assumption: we found no survey that measures it.
- AI changes engineering time by the multipliers in Exhibit 5.
- AI tooling costs $300 per seat-month for three people, based on list prices.

Table 4: On realistic trial evidence, AI moves a typical studio from 2 to ~2.1-2.6 ventures a year on the same budget

Scenario	Eng. time multiplier	Cost per MVP	Range (30-70% eng. share)	Change vs pre-AI	Months to MVP	Ventures/yr (same budget)	MVP spend per exit (40 MVPs)	MVPs per $1M
Pre-AI baseline	1.00	$76K	$76K-$76K	+0%	4.0	2.0	$3.0M	13
METR early-2025: 19% slower	1.19	$87K	$84K-$90K	+15%	4.4	1.7	$3.5M	11
METR late-2025: 4% faster	0.96	$78K	$77K-$79K	+3%	3.9	1.9	$3.1M	13
Enterprise field RCTs: ~21% less time	0.79	$71K	$68K-$75K	-6%	3.6	2.1	$2.9M	14
Greenfield lab RCT: 56% faster	0.44	$57K	$48K-$66K	-24%	2.9	2.6	$2.3M	17
Vendor estimate: 80% time saved	0.20	$48K	$35K-$60K	-37%	2.4	3.2	$1.9M	21

Source: Mudita Studios model. Baseline $76K and 4.0 months = interpolated medians of the Inniches 2024 survey (n=123); 2 ventures/yr = midpoint of the modal 1-3 bin; 40 MVPs per exit = Inniches funnel. Engineering share and tooling cost are assumptions.

The results (Exhibit 6, Table 4):

- **On the enterprise field-trial evidence**, cost per MVP falls about **6%**, to $71K. A studio with a fixed budget goes from 2.0 to 2.1 ventures a year.
- **On the greenfield trial**, the closest analogue to MVP work, cost falls about **24%**, to $57K. Throughput rises to 2.6 ventures a year, and MVP spend per exit falls from $3.0M to $2.3M.
- **On the METR 2025 result**, costs *rise* 15%.
- **Even the vendor's 80% figure** cuts cost per MVP by only about **37%**.

This is Amdahl's law applied to company building (Exhibit 7). Engineering speedups can only shrink the engineering share of the cost. Customer discovery, design, legal setup, recruiting, pricing tests, fundraising and go-to-market are untouched. Once coding is cheap, they become most of the cost.

Table 5: Once coding is cheap, the non-engineering share decides how far MVP cost falls

Scenario	Eng. 30% of cost	Eng. 50%	Eng. 70%	Eng. 50% + other work -20% (what-if)	Eng. 50% + other work -40% (what-if)
METR early-2025: 19% slower	+11%	+15%	+19%	+4%	-6%
METR late-2025: 4% faster	+3%	+3%	+2%	-8%	-18%
Enterprise field RCTs: ~21% less time	-2%	-6%	-10%	-17%	-27%
Greenfield lab RCT: 56% faster	-13%	-24%	-36%	-35%	-45%
Vendor estimate: 80% time saved	-20%	-37%	-54%	-48%	-58%

Source: Mudita Studios model. Cells = change in cost per MVP vs pre-AI baseline, including AI tooling. The 'other work' columns are what-if values for AI plus shared studio infrastructure cutting validation, design, legal and GTM effort by 20% / 40%; they are not measured data.

05 5. Our view: why the studio model gains the most

The model points to where the advantage lies. A solo founder with Cursor gets the coding gain. What they cannot get alone is a cheaper version of everything else. A studio can, because it repeats that work across a portfolio. We see three mechanisms. The evidence for each is suggestive, not proven.

1. **Shared infrastructure cuts the work AI coding leaves behind.** Table 5 shows that if shared studio infrastructure (templates, legal and finance stacks, design systems, distribution channels) cuts non-engineering effort by 40%, the greenfield scenario's saving rises from 24% to 45%. That 40% is a what-if, not a measurement, and it is the number studios should be trying to prove.
2. **Reusable agents amortize across ventures.** An agent that runs customer interviews, competitive research or regulatory checks costs roughly the same whether it serves one company or nine. Alloy Partners claims 1,133 agents running nine concepts for about $2,000 a month, though it reports no outcomes. The economic logic holds even if that claim does not: fixed costs spread across a portfolio favor the portfolio owner.
3. **A portfolio of experiments becomes cheaper.** The studio model is a numbers game: about 40 MVPs per exit. Inniches found no link between MVP cost, complexity or build time and eventual exits, and weak evidence that studios spending under $50K per MVP had higher success rates. If cheaper MVPs do not lower the odds of success, more of them is the rational strategy. AI lowers the cost of each shot.

The evidence that cuts against us.

- **AI commoditizes the studio's historic edge.** The engineering bench that let studios build faster than lone founders is now available to every founder for a few hundred dollars a month.
- **Studios have not proven they beat venture capital.** In Inniches' independent PitchBook analysis, studio-backed pre-seed deals exited **24%** of the time, against **38%** for VC-backed deals.
- **No AI-native studio has outcomes yet.** Of the 13 studios we catalogued (Table 6), one has its core facts independently verified and eight rest entirely on their own claims. None publishes audited venture outcomes.

The studio's case for the next decade therefore cannot rest on "we build faster". It has to rest on selection, distribution, shared infrastructure and capital, the parts of company building that AI coding does not touch.

Table 6: Of 13 studios positioning as AI-native, none publishes audited venture outcomes; 8 rest entirely on self-reported claims

Studio	Launched	Model	Claims	Capital	Outcomes disclosed	Status
AI Fund	2018-01	AI-thesis venture studio (co-founds with entrepreneurs and corporate partners)	3 months idea-to-funded (site claim)	$175M Fund I (2018); ~$190M Fund II (May 2025); 'backed by over $370M'	No portfolio-level outcome data published	mixed
Pioneer Square Labs (PSL)	2015 (AI pivot 2023)	Studio + separate venture fund	33 startups since 2015 (at May 2023)	$20M studio Fund III (2023); $100M PSL Ventures (2021)	GeekWire: majority raised follow-ons, a few acquired, no unicorn after 8 years	verified
Forum Ventures AI Venture Studio	2023 (pilot); fund 2025	Studio arm of an accelerator/VC	17 companies launched; plans 30 more; '63% raised follow-on within 12 months'	$10M pilot (2023); $20M fund per Forum ($11.5M successor fund per BetaKit)	Follow-on rate self-reported; no exits	mixed
Alloy Partners - One Health Studio	2026-04	Multi-partner corporate/government studio (Elanco, State of Indiana)	9 concepts; ~$2,000/month agent cost vs '100-300 people'; Series A plan 'in days'; ~90% of agent output 'solid'	Not disclosed	No revenue, funding or survival outcomes reported	self-reported
Inevitable AI Group (IAIG)	2026-01	AI-native studio launching multiple software businesses in parallel	5 ventures since January 2026; 'dozens more' expected by year-end	$6M pre-seed led by Aleph (Aug 2026)	No venture outcomes yet	mixed
Start.vc (Start Holdings)	2026-02	Full-stack studio with shared engineering/design infrastructure	First venture expected early 2026	Not disclosed	None yet	self-reported
AWAKE (Making Science)	2025-12	Corporate studio of a listed digital consultancy	'2x acceleration in tech feature	Not disclosed	Single-company	self-reported

Studio	Launched	Model	Claims	Capital	Outcomes disclosed	Status
			development'; portfolio company ad-machina 'multiplied its value by 10'		claim, unaudited	
Meta3Ventures	2026 (essay Jun 2026)	AI-native studio plus fund	Discloses that '30+ companies, 8 exits' figures are the founder's prior career, not the studio's	Not disclosed	None yet	self-reported
Builders	2021	AI venture studio with batch model and own operating system (VSI)	Up to EUR500K per venture	Fund I (studio) and Fund II (seed); sizes not disclosed on site	Not disclosed	self-reported
Shape Labs	2026 (article May 2026)	Boutique AI venture studio / product builder	Claims 2 people ship a v1 in 6-8 weeks vs 4-6 people, ~6 months and ~$500K in a traditional studio	Not disclosed	Two internal products named; no financials	self-reported
Misfit Labs	2026 (essay Mar 2026)	Builder-led studio	None quantified	Not disclosed	Not disclosed	self-reported
A studio operator (2026 industry report)	2026 (report Jun 2026)	'Distributed' AI-augmented studio	Claims 700+ portfolio companies	Not disclosed	Report figures loosely sourced; internally inconsistent	self-reported
Thrive Holdings (adjacent model)	2025	AI-native holding company acquiring and automating services firms (not a classic studio)	Claims tax prep time cut from 180 to 15 hours (site)	$1B+ committed at launch; $2B raised at $12B valuation (Aug 2026)	No audited operating results	mixed

Source: studio websites, press releases and independent press coverage; Mudita Studios analysis. 'verified' = core facts confirmed by independent press; 'mixed' = funding independently reported but performance claims self-reported; 'self-reported' = studio's own site, essay or press release only.

Measure your own inputs. Any studio can measure its own engineering share of MVP spend and its cost per MVP before and after adopting AI tooling. Those are the inputs this model is most sensitive to, and they matter more than any industry benchmark.

What would change our mind:

- Audited outcome data showing AI-native studio ventures do no better than solo AI-native founders at reaching revenue or funding.
- Evidence that agent-based validation and go-to-market save less than 20% of non-engineering effort.

> 45% possible MVP cost saving in the greenfield case if shared studio infrastructure cut other work by 40%, an untested what-if

IMPLICATIONS

What this means for studio operators and their partners

01 Expect real but modest savings from AI coding, largest on greenfield MVP work and smallest in mature codebases.

02 Measure your own engineering share of MVP spend and cost per MVP before and after adopting AI tooling.

03 Compete on selection, distribution, shared infrastructure and capital, the parts of company building AI coding does not touch.

04 Use lower MVP costs to run more experiments, since studios need about 40 MVPs per exit.

APPENDIX

Method, data and sources

Methodology & limitations

- **Datasets**. We assembled four datasets for this piece: AI-era company-building metrics (55 data points), AI coding productivity studies, AI-native studios and a public SaaS benchmark. Every data point carries a source, a date and a reliability rating (independent data, industry survey, self-reported or anecdote). All charts and tables are produced by a single reproducible analysis. The full dataset and sources are available on request.
- **Public SaaS benchmark.** Revenue is the XBRL-tagged annual revenue from each company's latest 10-K. Headcount is the total figure in the 10-K human-capital section, checked by hand against each filing. The basket is US-listed, so it skews toward large companies. The company-level figures are available on request.
- **AI-native figures are self-reported**. ARR is a run-rate and is not GAAP revenue. Headcount dates do not always match ARR dates; Cursor's headcount is approximate and ElevenLabs' is a third-party estimate. Midjourney's figure is reported revenue, not ARR. Start points for "time to $100M" differ by company, and the labels in Exhibit 2 show each one.

- **Stripe and Bessemer samples are selected for success.** They describe the top of the distribution, not the median startup. The "fastest SaaS" $1M figure (15.5 months) is derived from Stripe's "about four months ahead".
- **Carta reports disagree.** The 2024 solo-founder share is 30%, 31% or 35% depending on the report. We plot each report separately and do not interpolate between years.
- **Productivity conversions.** Effects reported as output gains were converted to time per task as 1/(1+gain) - 1. Cui et al.'s confidence interval is our approximation from the reported standard error. The trials use different tools, tasks and populations, so they are not directly comparable, and none studied a studio.
- **Studio model.** The baseline comes from a self-selected survey (11% response rate), and bin medians are interpolated. Three inputs are our assumptions, not measured data: the engineering share of MVP cost, the tooling cost and the non-engineering "what-if" multipliers. The model covers MVP spend only, not a studio's full cost per venture, overhead or follow-on capital. "Ventures per year" assumes a fixed budget and no other constraints. In practice, founder recruitment and partner attention also limit throughput.
- **Studio dataset.** Studios were found through web search; the list is a sample, not a census. "Verified" means core facts appear in independent press. It does not mean performance claims are audited.

Sources

A full source list with URLs, dates, measurement notes and reliability ratings is available on request. Principal sources:

- **Studio survey:** Inniches, *Big Venture Studio Research 2024* (Dec 2024).
- **Platform data:**
- Stripe, *Inside the growth of the top AI companies on Stripe* (Jun 2025), *Indexing the AI economy* (2025), 2024 annual letter via TechCrunch (Feb 2025) and 2025 annual update (Feb 2026).
- Carta, *Solo Founders Report 2025* (Dec 2025), *Founder Ownership Report* (Jan 2025) and *2026 edition*; Carta *State of Seed 2025* via a16z speedrun.
- **Public company filings:** SEC EDGAR Form 10-K filings for 29 SaaS companies (FY2025–26).
- **Benchmarks:** Benchmarkit (2025); SaaS Capital (Jul 2026); Bessemer Venture Partners, *State of AI 2025* (Aug 2025).
- **Productivity research:**
- Randomized trials: METR (Jul 2025; Feb 2026); Cui, Demirer et al., SSRN 4945566 (2025); Paradis et al., arXiv 2410.12944 (2024); Peng et al., arXiv 2302.06590 (2023).
- Surveys: DORA (Mar 2025; Sep 2025).
- Vendor estimates: Anthropic (Nov 2025; Dec 2025).
- **Company disclosures and press (self-reported):** Lovable; TechCrunch; Business Insider; SaaStr; The Decoder; ElevenLabs; Revelio Labs; Gamma; Index Ventures; Wiz via TechCrunch; Perplexity CEO via American Bazaar; YC via CNBC; EA Forum critique.
- **Studios:** AI Fund; TechCrunch (2018); GeekWire (PSL, 2023); Forum Ventures; BetaKit; Alloy Partners; The AI Insider (IAIG); Business Wire (Start.vc); Making Science; Meta3Ventures; Builders; Shape Labs; Misfit Labs; a studio operator's 2026 industry report; New York Times and OpenAI (Thrive Holdings).