

STEALTH NOTES // ISSUE 01 // JUNE 2026

Six hundred companies later.

By [Ali Sina](#) · Founder, STEALTH

READ TIME: ~8 minutes

COMPANION TO: The Studio Is the Founder (Manifesto) · The State of the Venture Studio Economy 2026

WORD COUNT: ~1,600

“ A workshop studio runs ten companies with a hundred people. The new shape runs a thousand companies with everyone. ”

Six hundred companies later, the studio I started looks nothing like the studio I started.

I want to tell you the story of how that happened — not because I think it's especially heroic, but because I think it accidentally describes a shape that's about to become very common, and almost nobody is naming it out loud.

This is a confession in the form of a thesis.

What we set out to build

In 2018 I read every venture studio post-mortem I could find. There were a lot. The good studios — Atomic, AlleyCorp, Pioneer Square Labs, Expa — had built reputations by doing one thing extremely well: taking a small number of carefully selected ideas, building them with hands-on intensity, and producing outsized outcomes.

The model worked. The math was good. The shape was a workshop. A few people in one room, building maybe ten companies at a time, with the kind of high-touch attention that doesn't scale and isn't supposed to.

That was the studio I set out to build.

For the first eighteen months we did exactly that. We had a deck. We had a cap table per company. We had calendars that looked like every other early-stage studio's calendar — back-to-back portfolio reviews, board prep, weekly standups with each founding team.

By company twelve, something was broken and I didn't know what.

What broke

I'll give you the moment exactly.

It's a Tuesday. I have eleven active portfolio companies. I'm about to spin up the twelfth. My calendar that week has thirty-two scheduled meetings with founders, four with potential LPs, six with operating-team members, and roughly two hours of actual unscheduled time.

I'm preparing for the most consequential meeting of the week — a board check-in with company six, which is about to either raise its Series A or implode — and I realize I haven't actually thought about company six's strategy in about three weeks.

I haven't been thinking. I've been showing up. There's a difference.

I sit with this for a few days and start asking the question every founder eventually has to ask: *am I the bottleneck?*

The honest answer was yes, and worse than I'd realized. Every company in the portfolio was throughput-limited by my attention, my partner's attention, and the original team's attention. We weren't building twelve companies. We were building one company, twelve times in parallel, badly.

I had three options.

I could slow down. Cap the portfolio at fifteen. Become a smaller, better-rested studio. Plenty of houses have done this and it works.

I could hire more partners. Add three more senior operators and double the cap. Plenty of houses have done this too. It also works, and it costs you the original culture by year three.

Or I could change the shape of the studio itself.

I picked the third one, and at the time I didn't realize how big a decision it was.

What we built without naming it

We didn't add partners. We added *distribution*.

The work that the founding team had been doing manually — the research, the validation, the early sourcing, the hiring drafts, the customer development, the first competitor maps, the cap table modeling, the legal templating, the supplier outreach — we started splitting into discrete tasks and putting them on a workspace.

Anyone capable of doing the task could pick it up. The institutional knowledge — the *how Stealth does it* — stayed in playbooks. The actual execution moved into the network.

We did not announce this. We did not give it a name. We just started doing it because we had to.

Then we kept building companies.

Twelve became fifty.

Fifty became two hundred.

By the time someone asked me how many companies were in the Stealth portfolio, I had to count.

The honest answer in June 2026 is **somewhere north of seven hundred and still climbing**.

That number should be impossible inside a traditional studio model. It isn't impossible. It just means we are no longer running a traditional studio model.

The thing nobody is saying

I want to be careful here.

There are roughly seven hundred and twenty-four venture studios in the world right now (the Global Startup Studio Network's number, not mine). Most of them are running the original workshop model — fifteen to twenty-five active companies, attention-bounded, intentionally boutique. That model is *correct* for what it's doing. I am not arguing against it.

What I am arguing is that there's a new shape forming inside the same category, and the people running the new shape have mostly been doing it without naming it.

A workshop studio runs ten companies with a hundred people.

The new shape runs a thousand companies with everyone.

The difference is not in the quality of the work — although a lot of the early skeptics assumed it was, and many of them have stopped assuming it. The difference is in *where the bottleneck lives*.

In a workshop studio, the bottleneck is the founders' attention. In the new shape, the bottleneck is the operating system — the playbooks, the workspace, the rules by which distributed labor reliably produces institutional-grade output. The bottleneck moved out of the room.

And once the bottleneck moves, you discover that the studio is doing something the original model could not.

It's compounding.

Every company built by a Stealth-shape studio makes the next one cheaper to build, faster to launch, and more likely to clear its Series A. By company three hundred, the marginal cost of starting company three hundred and one is closer to zero than to anything resembling a traditional startup formation budget.

This is a structural advantage that gets bigger every year, mechanically, whether the founders sleep eight hours or four.

The name we are putting on it

I have been circling around this for two years and it took our research team to push us to commit to the term.

We're calling it **The Distributed Studio**.

Not because the word is original — distribution is a familiar concept — but because the word does the work. A traditional studio is concentrated: a few people, one room. The distributed studio is, well, distributed: institution in the center, execution in the network, AI in the seams.

We just published the first annual data report on the category at stealth1000.com/research. 724 studios. 3.5× IRR premium over traditional seed. \$14.2B deployed in 2024. The report is free, sourced, and citable. We also published a companion manifesto — *The Studio Is the Founder* — that walks through the thesis in essay form.

This post is the conversational version of both.

What I would tell you if you were starting one

A few people have asked me, mostly privately, whether they should start a venture studio in 2026.

My honest answer is *yes, but not the kind you're picturing*.

The workshop model is full. The good ones are taking the good founders, the good capital, and the good operators. You will not out-execute Atomic at Atomic's model. They have been refining it for sixteen years.

But the distributed model is wide open. There are maybe ten studios in the world running it right now. Most of them aren't naming themselves that way. The category has barely been defined. The infrastructure layer underneath — the workspaces, the credential systems, the AI-augmented operating layer — is still being built in public.

If you are an operator looking at the studio category and you want to do something the next decade actually rewards: build the distributed version. Build it carefully. Document it. Hire from the network you're creating. Don't be in a hurry to look like Atomic. Atomic is doing something else.

We are six hundred companies in and the shape is still bending.

There is a lot of room.

What I'm not going to do

I'm not going to end this post by asking you to buy something.

I'm not going to make the closing argument that you should partner with Stealth, hire us for a corporate venture studio, or invest in Fund I. (Although if you want to do any of those, the links are at the bottom of every page on stealth1000.com.)

I'm going to ask you to do one specific thing.

If you work near venture or near the studio category — **read the report**. Print pages four and five and put them on your desk. The data is real. The growth curve is decisive. The category is leaving "alternative" and entering "infrastructure," and the next ten years of investment decisions are going to look different because of it.

If you're a founder — you do not need to be alone in a room. There is a model where the institution carries the load that nobody trained you to carry. We did not call ourselves founders by accident.

If you're an LP — there are maybe ten distributed studios in the world right now and they have not been compared head-to-head against the workshop houses yet. The fund-of-funds that does the head-to-head first will probably end up running the playbook for the rest of the asset class.

The studio is the founder.

Walk in.

– [Ali Sina](#)

Founder, STEALTH

ali@stealth1000.com · stealth1000.com

Next issue: *the IRR math. Why the 3.5× premium is real, where the selection-effect skeptics are right and wrong, and what LPs should actually model. Subscribe at stealth1000.com/research.*