WORKING PAPER 01 · STEALTH RESEARCH

JUNE 2026 · DRAFT · OPEN FOR PEER FEEDBACK

The distributed studio operating architecture.

Four functional capabilities required to operate a venture studio at distributed scale. Working paper · open to peer review · changes documented before publication of v1.0.

⚑ WORKING PAPER · NOT FOR CITATION AS FINAL

Why this paper.

The manifesto declared the model. The Annual Report sized it. This paper does the engineering work underneath.

In June 2026 STEALTH Research published *The State of the Venture Studio Economy 2026, Vol I*, which named and sized the category of distributed studios. The companion manifesto, *The Studio Is the Founder*, described what the model is and why it works.

Neither document answered the next question — the one every operator who has read both has independently emailed to ask: **"What does the inside of a distributed studio actually look like?"**

This working paper is the attempt to answer it.

We do not believe the answer is one architecture. We believe there are perhaps three to five viable architectures, and the next decade of distributed studios will discover them. This paper documents the one architecture we know — STEALTH's — abstracted to the level of **capability** rather than specific tools, so the framework is portable across studios that will build differently.

We are releasing it as a working paper rather than a final publication for two reasons. First, distributed studios have only existed at meaningful scale for five years, and any architecture document released as "final" in 2026 will look quaint by 2028. Second, we want peer feedback from the eight to ten other studios already operating in this space, and the working-paper format invites that.

Comments are welcome at **support@stealth1000.com** or stealth1000.com/research/working-paper-01-comments. Substantive feedback will be incorporated into v1.0 (planned Q1 2027) with author attribution.

> *"A workshop studio runs on attention. A distributed studio runs on architecture." — author's note, 2026*

The four pillars.

A distributed studio that omits any one of these four capabilities will revert, structurally, to a workshop.

PILLAR 01 · THE WORKSPACE

A task-resolution system that produces institutional-grade output from distributed labor

The unit of execution moves from the meeting to the task.

The workspace is the most visible and the most misunderstood of the four pillars. It is not a project-management tool. It is a *task resolution system*: a structured environment where institutional knowledge is encoded as task templates, distributed operators pick up tasks, evaluation rubrics enforce quality, and finished tasks compound into reusable artifacts.

The defining test of a workspace is: can a new operator who has never worked at the studio before complete a non-trivial task to institutional-grade quality, supervised only by the rubric, within their first week? If yes, the workspace is real. If no, the workspace is a project board.

```
SIGNAL TEST: Time from operator onboarding to first credit-quality task delivered.
Target: < 7 days.
```

PILLAR 02 · THE CODIFIED OPERATING SYSTEM

Written playbooks that survive the original team's attention

The institutional knowledge moves from heads to docs.

The codified operating system is the studio's institutional memory written down. It includes: company-formation playbooks, supervisor evaluation rubrics, hiring funnels, vendor relationships, legal templates, pricing experiments, decision frameworks, and post-mortems. In a workshop studio, most of this lives in the founders' heads. In a distributed studio, it has to live in documents that distributed operators can reference, learn from, and update.

The defining test: if every founder of the studio disappeared for ninety days, would the studio continue to produce institutional-grade work? In a workshop studio, the answer is no, and that is fine. In a distributed studio, the answer must be yes — or the model collapses back to the attention ceiling.

```
SIGNAL TEST: Documented decision frameworks per portfolio company. Target: ≥ 1 per
company per quarter.
```

PILLAR 03 · THE CREDENTIAL LAYER

A way to recognize and rank distributed contributors

Talent observability moves from interviews to outputs.

Distributed studios cannot run on traditional hiring. The talent pool is too distributed, too anonymous at first contact, and too large for the founders' attention. Instead, distributed studios run on credentialing — a system that observes contributor outputs over time and produces a portable, verifiable signal of quality and trust.

STEALTH built this jointly with Forkaia® (the credentialing function on the student side). The Forkaia® Score (a 300–850 work-readiness index) is the protoype for the credential layer that we believe every distributed studio will need to build or borrow within five years.

The defining test: when the studio needs to assign a new high-stakes task, can the studio identify the right contributor without conducting an interview? In a workshop studio, every important task requires a conversation. In a distributed studio, the credential layer makes most assignments routine.

```
SIGNAL TEST: Percentage of high-stakes task assignments made without a 1:1 interview cycle. Target: ≥ 60%.
```

PILLAR 04 · THE WRITING CULTURE

The deepest cultural prerequisite — almost no traditional studio has it

Communication moves from verbal to documented by default.

This pillar is the one most workshop studios cannot adopt. A distributed studio that does not write things down — decisions, lessons learned, post-mortems, hiring rationale, pricing experiments — will revert to a workshop within twelve months, regardless of how good the other three pillars are. The writing culture is the substrate that allows the workspace, the OS, and the credential layer to function.

Writing culture is not the same as documentation requirements. Plenty of organizations have documentation policies. Writing culture is the social expectation that important decisions are made in writing, that disagreement is voiced in writing, that lessons are captured in writing, and that nobody is punished for writing something down that turns out to be wrong.

The defining test: when an important decision is made, is the rationale visible to a distributed operator who joins the studio six months from now? In a workshop studio, this is rare. In a distributed studio, it is the default.

```
SIGNAL TEST: Rate of post-mortem completion per portfolio company per quarter. Target: ≥ 1 per company per quarter.
```

Implementation order.

If you are building a distributed studio from a workshop, the order matters. We learned this expensively.

The four pillars are not equally easy to implement, and they are not independent. In our experience building STEALTH, the order in which a studio adopts these capabilities determines whether the model takes or rejects them.

Recommended sequence

1. Writing culture first. The single hardest pillar. Cannot be installed retroactively. Start by writing down decisions before you build anything else. If you cannot establish a writing culture in the first ninety days of attempting the distributed model, abandon the attempt — the other three pillars will not stick.

2. Codified Operating System second. Once writing is normal, codifying the existing playbooks becomes natural. You will discover that 60% of what you thought was institutional knowledge was actually a single person's preference. That is fine. Document anyway.

3. The Workspace third. Without writing and an OS, the workspace produces low-quality output. With them, it begins to compound. Resist the temptation to build the workspace first because it is the most visible pillar — it is also the one that fails most dramatically without the other two underneath.

4. The Credential Layer last. Cannot be built before you have enough distributed contributors to rank. Wait until the workspace has produced at least 100 evaluated outputs per active contributor before attempting to formalize the credential. Premature credentialing creates legibility without signal.

How long this takes

From workshop to distributed studio, in our experience: 18–36 months. Faster than that suggests something was already in place. Slower suggests one of the pillars is being skipped.

What we got wrong

Three honest mistakes from the STEALTH path, in case they are useful to other operators:

- **We built the workspace before the codified OS.** The first six months of workspace operation produced low-quality output until we backfilled the playbooks. Lesson: writing and OS first.
- **We over-rotated on credentialing too early.** Our first scoring system was a hand-rolled rubric that turned out to encode the founding team's preferences rather than universal quality signal. Took 14 months to recalibrate.
- **We assumed writing culture was a personality trait of the team.** It is not. It is an institutional norm that has to be enforced for the first 90 days and reinforced thereafter. Twice in our second year we had to publicly course-correct decisions that were made verbally.

Open questions.

Five things we genuinely do not know — and would like peer feedback on before v1.0.

1 · Is the credential layer transferable across studios?

STEALTH's Forkaia® Score is calibrated to STEALTH's portfolio composition. We do not know whether the same scoring methodology — or a different one — would produce comparable signal at a different distributed studio. Our hypothesis is that the methodology generalizes; the calibration does not. This is testable as the DSI matures and we have multiple studios reporting Score-derived metrics.

2 · How does the architecture interact with AI capability gains?

Every six months the marginal cost of structured intellectual work drops as AI models improve. The four pillars described here were designed under one set of assumptions about AI capability. We do not know how the architecture should adapt when that capability is meaningfully higher. Particularly: does the workspace remain task-based, or does it become agent-based? We are watching this carefully.

3 · Are there industries where the distributed model does not work?

Our portfolio is heavy on software, fintech, and digital services. The distributed model has not been tested at any scale in deep tech (semiconductors, biotech, hard sciences) or in heavily-regulated industries (banking, healthcare). We suspect the model adapts but slower; we have no data.

4 · What is the right governance structure for a distributed studio?

Workshop studios use the standard LLC partnership structure. Distributed studios have governance questions that traditional structures do not address — particularly around contributor profit participation, IP ownership of contributor-produced playbooks, and equity treatment for high-credential operators. We are working through these and would value comparative notes from other studios that have addressed them differently.

5 · How do distributed studios fail?

We do not yet have enough distributed-studio failures to do a clean post-mortem comparative study. The model has been live at scale for ~5 years. The first wave of distributed-studio failures will likely cluster around 2027–2028 as the early experiments mature. We intend to document those failure modes in a future working paper.

How to contribute

If you operate or advise a distributed studio and have an answer (or a counter-question) on any of the above, write us. Substantive contributions to v1.0 are credited with author attribution. **support@stealth1000.com**

About + license.

Bibliographic information, citation, and license terms for this working paper.

Status

This is a **working paper**, not a final publication. The framework presented here represents STEALTH's current best understanding of distributed-studio operating architecture as of June 2026. The framework will be revised based on peer feedback before publication of v1.0 (target: Q1 2027).

Authorship

This working paper was prepared by STEALTH Research, the institutional research function of STEALTH. Primary authorship is institutional rather than individual. Acknowledgments to the operating team, the workspace contributors who shaped these capabilities through daily practice, and the founder of Forkaia® for the parallel work on the credential layer.

Citation

> STEALTH Research. (2026). *The Distributed Studio Operating Architecture (Working Paper 01).* stealth1000.com/research

Cite as working paper. Final citation will update with the v1.0 publication.

Related work

- *The State of the Venture Studio Economy 2026, Volume I* — Annual Report (the data context)
- *The Studio Is the Founder* — Manifesto (the thesis context)
- *The Distributed Studio Index (DSI) — Methodology, Volume I* — Benchmark methodology paper (the measurement framework)
- *"Six Hundred Companies Later"* — STEALTH Notes Issue 01 (the narrative version)

License

Comments period

Open: **June 2026 → December 2026**. Submit via support@stealth1000.com or stealth1000.com/research/working-paper-01-comments.

STEALTH Research · Working Paper 01 · Distributed Studio Operating Architecture · June 2026

support@stealth1000.com · stealth1000.com/research