FIRST STRATEGY · CASE STUDY

Venture Studio

A real engagement, anonymized to industry label. The full case study: the story, the deliverables we produced, and the plays that ran it.

The story

Day One Proposal

Day One Audit

Playbook and Delivery Proposal

Charter

The plays

A missed deal makes no sound

On a weekend in front of a whiteboard, the CEO of a medtech venture studio corrected our paraphrase of his ambition. We had said his studio needed to be twenty times better at sourcing deals than anyone else in the space. A hundred times, we tried again. "A thousand times better," he said, and he was not joking, because he had spent the previous day finding out what one times looks like: a full day of his own hours negotiating with a single university's tech transfer office, for one deal, from one source, in a market with hundreds of them.

His studio runs a model with no slack in it. It does not start companies from ideas; it acquires intellectual property, from universities, research hospitals, government labs, and the occasional inventor whose own university turned him down, and develops it into clinical assets: pre-revenue medical devices, built to be acquired. The thesis is a blade. Devices, never drugs, because drugs mean human trials and a hundred times the cost. A regulatory pathway with a predicate. A sellable asset in 18 to 30 months on one to two million dollars, in a category where the convention is five years. The team that runs this is three principals: the CEO, a regulatory lead who has taken six devices through FDA, and a chief science advisor who can tell an engineering problem from a magic trick.

The market they hunt is strange and rich. Institutions fund research by the billions, end up holding the patents, and mostly let them sit. The offices whose job is to market that technology are understaffed and unread. The researchers want publications and patents, not companies. The result is a standing inventory of engineered, often validated technology that almost nobody is systematically reading, and behind it, a research literature where the next inventory surfaces years earlier, roughly five journals per category across every category the studio cares about. Nobody reads that either. The CEO had been trying, with a custom GPT he built himself, in evenings, one opportunity at a time, and the searches remembered nothing from one session to the next.

That was the diagnosis, and it had a shape worth stating precisely, because everything afterward followed from it. The studio's edge was judgment: a thesis sharp enough to kill most of medtech on sight and the specialists to vet whatever survives. The studio's constraint was reading, and the loss it caused was invisible, because a missed deal makes no sound. The deals the studio never saw cost it nothing it could measure and everything it was for. The scarcest asset in the building, specialist judgment that runs about five thousand dollars per candidate evaluated, was being spent partly on a machine's work: finding things and reading them. Machines should read. People should judge. The whole engagement is that sentence, built.

The cheap inventory before the expensive insight

The sharpest thing the CEO said that weekend was also the most expensive instinct in the room. Really good technology, he had realized, often never reaches the tech transfer offices at all. It lives upstream, in

the labs, in the papers, before anyone thinks to commercialize it. He is right, and the studio had the receipt to prove how much being right costs: a quarter-million-dollar funded research project at a university, running at that moment, with five possible use cases and no certainty about which product it would become. Going upstream means paying for the future before it is investable.

So the instinct went up on the whiteboard and the costs went up next to it, and the CEO talked himself out of his own sequence, which is the best way for it to happen. The transfer inventory is sitting stock: engineered, often validated, listed by someone whose job is to sell it. The literature is earlier, riskier, and priced like research. The cheap test is to exhaust the inventory that is already for sale, and to go upstream only when the downstream pool proves thin. The decision came with an architectural rider that mattered more than it looked: the machine that reads a commercialization site must not care that it is reading a commercialization site. Point it at a journal next year and the funnel downstream is identical. The insight was not abandoned. It was sequenced.

The same weekend settled what a machine could be trusted to score. Of the studio's five investment gates, three are data disciplines wearing judgment's clothes: the patent landscape is searchable, the regulatory pathway is a predicate lookup, and acquisition appetite is M&A history. Those the machine would score, weighted and versioned. The other two, whether a device can actually be manufactured and tested inside the budget and the timeline, stay human, because that judgment is the studio's product. And one decision preceded all of it: the methodology is a trade secret. No patent would ever be filed on it, because a patent is a public recipe. The CEO put us under NDA in the first hour, cheerfully, mid-sentence.

He asked for the proposal the same day, in operator's language: get me something I can budget. It went out within days. Three phases, a month each, fixed fee per phase, each phase independently valuable, stop at any boundary. No board approval was needed. He could decide, and did.

Wrong early, on purpose

The build's first weeks were floor time, except the floor was a funnel. Agents went up against the first commercialization sites, and every run was watched: a machine reading listings, a human checking every call it made. The sites share a purpose and almost nothing else. Each publishes its own way, structures its data its own way, buries its technologies at its own depth, and each source failed differently before it ran clean. Everything the agents read went into the database, relevant or not, because relevance is a function of a thesis that moves, and re-scoring a stored record is free while re-scraping the world is not.

The thesis screen came up the same way: measured, not trusted. Device or drug, in or out of category, engineering problem or science project, the screen made its calls and a human made the same calls blind, and the gap between them was the number that mattered. Early, the gap was wide. That was the plan. A screen that is wrong in measured ways improves; a screen that is wrong in unmeasured ways just lies to you, so the feedback loop was wired in from the first gate, the CEO's corrections flowing back into a matrix that carried version numbers like code. The gates followed on top: patent landscape, regulatory

pathway, acquisition activity, each scored advisory, with feasibility flagged for the specialists rather than scored with false authority.

What came out the other end was the artifact the whole funnel exists to produce: a brief. One surfaced candidate, what it is, why it fits the thesis, what the gates found, readable in minutes. The trust moment was not a demo. It was the CEO, who used to spend an evening wrestling one opportunity out of one source, working through a morning's briefs in the time a single raw listing used to take, and realizing the reading was now happening continuously, without him, while he slept and while he ran his three live programs. The system was reading hundreds of opportunities a week. The obvious kills died quietly. What survived arrived scored and explained.

The discipline that kept the trust honest was the unglamorous one: the discard pile got sampled, every week, by a human. A wrongly surfaced candidate costs minutes; a wrongly killed one is invisible forever, the same silence the studio had been living in before, rebuilt at machine speed. Sampling the kills is how the funnel stayed humble, and more than one correction that taught the matrix something came off the pile that was supposed to be noise.

Scale is when adding costs nothing

The early sources were each a small build, an agent crafted to one site's shape, and the builds kept shrinking. The differences between sources, it turned out, were configuration, not code, and the engine generalized until the build disappeared into the architecture: a new deal source now onboards from its URL. Name the source, decide its scraping posture deliberately, hand the engine the address, watch its first runs the way every source gets watched, and it joins the rotation with its own volume baseline. Adding a source stopped being a project and became a decision.

That is the line the economics of coverage crossed. When the marginal source costs almost nothing, breadth stops being a budget item, and the rotation grew into a standing read of the commercialization landscape that no hand process approaches: hundreds of opportunities a week, with headroom to thousands, flowing through the screen and the gates without a person touching the reading. The journals, the expensive insight from the whiteboard, are still waiting upstream, unspent, because the cheap inventory has not run thin. The hypothesis test is still running and still paying.

Underneath the throughput, the quieter asset compounds. Everything the engine has ever read is logged: source, researchers, summary, categorization, the irrelevant alongside the relevant. When the thesis moves, and it moves, a new version of the matrix re-scores the entire database overnight, and last January's near-miss resurfaces by itself as this June's candidate. The CEO had named this asset on the whiteboard the first weekend, grinning: "I've got the best database of medical inventions on the planet. It'd be fun to say that to investors." It was a joke then. It is a database now, growing with every run, and it is the foundation of the studio's long game: the firm that has read everything is the firm researchers eventually come to, and deal flow that comes to you is the only kind that scales past a thousand times better.

Silence is a symptom

A scraping system's characteristic failure is not a crash. It is a quiet, healthy-looking run that read less than it should have, and the engine met it on schedule. One source restructured its site. The agent reading it did not break; it kept reading what it still recognized, which was a fraction of what was there, and run after run came back green: errors zero, volume quietly down. No error log catches that, because nothing erred. What caught it was history: every source carries its own volume baseline from the day it is onboarded, and the baseline flagged the drop for what it was, a silence where reading used to be. The agent was reconfigured, re-watched like a new source, and returned to the rotation. The incident was routine. The rule it confirmed was not: in a system built to make missed deals impossible, silence is a symptom, and aggregate health is not health. The engine's totals had looked fine the whole time, the way a portfolio looks fine while one position quietly dies.

That posture, watch the segments, sample the kills, version the judgment, is what the system runs on now that it is boring, and boring is the achievement. The three principals run the deal review of a hundred-person firm. The screen reads everything; the briefs surface the few worth a principal's minutes; the five-thousand-dollar specialist hours are spent only on candidates that earned them; and every pursue, pass, and correction keeps teaching a matrix that belongs to the studio, on the studio's machine, under the studio's NDA, scoring with a method no one outside can read. The CEO's job inverted along the way. He started as the funnel, the man whose evenings were the studio's sourcing capacity. He is now the judge of what the funnel surfaces, which is the job the studio actually needed him doing, and the reading never stops, never tires, and never forgets what it read.

Day One Proposal

The opening session

Prepared for the CEO, on a standing relationship.

What this is

A working session inside your deal flow before anything gets built. Real work, not slides.

One day, you and us, in front of the way the studio actually finds technology. How a piece of research becomes a candidate. Where candidates come from today: the commercialization sites, the journals, the labs, the inventors. What your thesis admits and what it kills. What a candidate costs you to evaluate once a human touches it. Where the hours actually go.

Day One usually means a day on a client's floor, watching the work as it happens. Your floor is a deal flow. We walk it the same way: one opportunity traced end to end, from the lab that produced it to the decision you make about it, with the person who does that work today. That person is you, which is the finding before the session starts.

We come ready to listen and to think on our feet. No prepared deck.

What you walk away with

A plan, within days. Not a deck. Not a recommendation memo hiding in a PDF. A written read for operators.

The session pulls the signal from the only person who has it: how the studio decides, what it pays for judgment, and where the search breaks. That work turns the signal into a sequenced plan you can run, and a build proposal you can budget.

The plan answers three questions:

- Where AI fits in your sourcing, and where it does not.
- The highest-leverage moves we see, sequenced so you can act on them in order. A roadmap, not a list.
- What it would take to run the sequence: with your own team, with another firm, or with us.

The plan is yours. Run it however makes sense.

What we need from you

- A day of the CEO's time, with the whiteboard and the market hypothesis as they stand.
- The thesis as you actually apply it: what kills a deal, what makes one, and the scoring instincts you carry in your head.
- The tools you already use, however rough: the manual searches, the prompts, the comparables work.

- Candor about the costs: what an evaluation takes in specialist hours, and what going upstream costs when you pay for research instead of finding it.

The terms

The session ran on the relationship, not a fee. The standard entry is a flat fee of [flat fee] for the day and the plan, travel and expenses at cost.

No retainer. No commitment beyond the session itself. If we are the right fit for what comes next, we will already have been talking about what that looks like. If we are not, the plan is still yours to run.

What happens next

After the plan, you decide. Run it with your own team, hand it to another firm, or build it with us. If the work points to a build we are right for, we will scope it in a phased proposal, priced so you can budget it, with each phase independently valuable.

Day One Audit

The one-line finding

Your bottleneck is not judgment. It is reading. The studio's edge, a thesis sharp enough to kill most of medtech on sight and the specialists to vet what survives, is starved by a search that runs on one person's hours. The search space is effectively infinite, the team is three principals, and the scarcest asset in the building, specialist judgment, is being spent on the cheapest task in the building, finding and reading research. Machines should read. People should judge. Build the funnel that way and the studio reviews deal flow like a firm thirty times its size.

How we looked, and how we measured

A working session inside the deal flow, with the person who runs it. There is no floor to walk in a sourcing operation, so the floor we walked was the funnel itself: where medtech IP is generated, where it surfaces, how the studio finds it today, what the thesis admits, and what every downstream step costs. One opportunity was traced end to end, from the lab that produced it to the studio's decision about it. Where a figure below is the studio's own number, named in the session, it is the engagement record. Where a figure is our estimate, this audit says so.

The supply side: where the technology is

Medtech IP is generated in four places, and the studio knew all four. What the session established is how differently they behave as sources.

IP holder	How it surfaces	State of access
Universities	Commercialization sites, run by offices whose job is to market the IP	Public, listed, marketed, and mostly unread by investors
Research hospitals	Same commercialization machinery, often above-average volume	Public, same shape
Government labs	The same model again: funded research, transfer offices	Public, same shape
Individual inventors	No machinery at all; surfaced by relationships	Invisible to any systematic search

Behind all of them sits the research literature itself: peer-reviewed journals, roughly five per research category across the five to ten categories the thesis covers, where work appears before it ever reaches a

commercialization office, if it ever does. The CEO's sharpest observation in the session, and the one that reframed the build: really good technology often never reaches the transfer offices at all. It sits upstream, in the labs.

The structural fact about this market is the perverse incentive behind it. The institutions fund the research, end up holding the IP, and mostly let it sit. The offices that market it are measured on deals closed, not value captured, and the researchers want publications and patents, not businesses. The result is a standing inventory of engineered, often validated technology that nobody is systematically reading. That inventory is the opportunity.

Stakeholder map

Role	What they own	Where their conviction is	Their definition of the problem
CEO	The studio's operations and its deal flow	The model: sourcing is the studio's edge if it can scale	I am the funnel, and my hours are the ceiling
Regulatory lead	The FDA pathway on every program	The pathway: regulatory is an open-book test if you pick devices with predicates	Do not bring me candidates without a predicate path
Chief science advisor	Scientific validity	The science: separate engineering problems from magic	Most of what surfaces is not real; filter it before it reaches me
IP counsel	The patent landscape	The landscape: a clear path to an issued utility patent or no deal	A bloody landscape kills a deal no matter how good the technology
The board	Capital and pace	The portfolio: more programs, sooner	When does the studio scale past three programs

Nobody disagreed about the problem. That is rare, and it is worth recording: every stakeholder independently described the same funnel and the same starvation. The disagreement, mild and productive, was about where to point the machine first, and the session settled it with a sequencing argument rather than a debate.

One opportunity, traced end to end

The path one piece of technology travels from a lab to a studio decision, as the work ran at intake.

1. **The research happens.** Funded by an institution, run by a researcher whose incentives end at publication and patent. Invisible to the studio.

2. **The work surfaces.** A paper in a peer-reviewed journal, a listing on a commercialization site, or both. Public. Nobody is watching.
3. **Discovery.** The CEO finds it: a manual search, a session with a custom GPT he built, a conference, a relationship. One opportunity at a time, read by one person, in hours that compete with running three live programs.
4. **The thesis screen.** Device, not drug. A Class 2 pathway. Reachable in 18 to 30 months on one to two million dollars. An engineering problem, not breakthrough science. The screen lives in the CEO's head and runs only when he is reading.
5. **The gates.** IP landscape, regulatory pathway, acquisition activity, manufacturability, testability. Some checks are quick; the real versions are specialist work.
6. **Deep diligence.** IP counsel on the patent landscape, the regulatory lead on the predicate path. About five thousand dollars of specialist time per candidate that reaches this step.
7. **The pursuit.** Negotiation with the institution. A full day of CEO time for a single university, in the engagement record, and that was a deal that went well.
8. **The decision and the learning.** A program starts, or the candidate dies. Either way, what the evaluation taught goes nowhere: no database, no memory, no way to re-find the near-misses when the thesis shifts.

Steps 1 and 2 are the market. Steps 4 through 8 are the studio's strength. Step 3 is where the operation starves, and step 8 is where it forgets.

The friction, quantified

Friction point	The cost, as measured or estimated
Reading capacity	One person, one opportunity at a time. Our estimate in the session: a few dozen opportunities meaningfully evaluated in a good month, against a space that generates orders of magnitude more
The miss rate this implies	Unknowable by definition, which is the point. The studio cannot count the deals it never saw. With hundreds of commercialization sites and roughly 50 relevant journals publishing continuously, coverage by hand rounds to zero
CEO hours per source	A full day for one university negotiation; evenings for GPT-driven searches with no memory between them
Diligence spend at risk	About five thousand dollars of specialist time per candidate. Every weak candidate that reaches step 6 burns it
The upstream alternative	A quarter million dollars to fund one university research project to get ahead of the transfer offices, with five possible use cases and no certainty which product it becomes
Staleness	No change detection. A source scanned last month and a source never scanned look identical today
No compounding	Each search starts from nothing. A near-miss evaluated in January cannot be re-found in June when the thesis moves

The cumulative effect: the studio's deal flow was bounded not by the market's supply of investable technology, nor by the studio's ability to judge it, but by one person's reading hours. Most of the loss enters at step 3 and it is invisible, because a missed deal makes no sound.

What the session disproved

The CEO came in believing the machine should chase the upstream research first, because that is where the undiscovered technology lives. The session reversed the order, and the CEO made the argument himself once the costs were on the whiteboard: the transfer offices are sitting inventory, already engineered, often already validated, with someone whose job is to sell it. The literature is upstream of that, earlier, riskier, and more expensive to act on, as the quarter-million-dollar research project demonstrates. The cheap test is to exhaust the inventory that is already for sale before paying to go upstream. Start with the commercialization sites. Prove they are insufficient before the journals earn their build.

The architecture decision that falls out: the scraper must not care which kind of source it reads. Point it at a commercialization site today and a journal next quarter; the funnel downstream is identical.

Where AI fits, and where it does not

- **Fits: the reading.** Scraping the commercialization sites, cataloging everything found, detecting what changed since last visit. This is the starved step, and it is entirely machine work.
- **Fits: the thesis screen.** Device versus drug, category fit, the first-pass read of whether a technology is an engineering problem or a science project. The screen the CEO runs in his head, run by machine on everything, with his corrections teaching it.
- **Fits: three of the five gates.** IP landscape (the patent databases are searchable), regulatory pathway (predicate devices are a lookup discipline), and acquisition activity (M&A history is data). Scored, weighted, versioned.
- **Fits: the briefing.** A surfaced candidate arrives as an intelligent summary with its thesis alignment stated, built for a quick read, so a principal spends minutes deciding whether to spend specialist dollars.
- **Does not fit: the thesis itself.** What the studio believes about the market is the company. The machine applies it; people author it and version it.
- **Does not fit: engineering and testing feasibility.** Whether a device can be manufactured and tested inside the budget and the timeline is judgment that lives with the regulatory lead and the science advisor. The machine can flag, never score with authority.
- **Does not fit: diligence, negotiation, and the decision.** The five-thousand-dollar steps stay human. The machine's job is to make sure they are spent on candidates that deserve them.

The opportunity, sized

The arithmetic is capacity, precision, and compounding, in that order.

Capacity. A machine that reads the sources continuously moves the starved step by one to two orders of magnitude without adding a person, from a few dozen evaluations a month of one person's evenings to machine reading that does not stop. It is the step that bounds everything else.

Precision. The gates exist to protect the five-thousand-dollar diligence spend. If gating holds the studio to its target of three to five opportunities a quarter reaching deep diligence, the diligence budget concentrates on candidates that earned it, and the costliest failure mode, specialist time burned on weak candidates, falls with it.

Compounding. Everything scraped is stored, relevant or not. The studio accumulates a proprietary database of medtech inventions that no competitor has a reason to build, re-scorable overnight when the thesis changes, without touching a single source again. The same store is what the inverted funnel is eventually built on: the studio that has read everything is the studio researchers come to.

What we did not size, on purpose: the value of a single deal found that would otherwise have been missed. One program that exits inside the studio's model repays the build by multiples, and the session's

honest answer is that the miss rate today is unknowable. The sizing above stands without it.

Risks and constraints we observed

- The methodology is a trade secret and must stay one. NDA before work begins, no patent filings on the method, access held tight. The system's value is partly that nobody knows how it scores.
- Source diversity is real work. Every commercialization site publishes differently; some sources tolerate scraping and some prohibit it. Each source is a small build, and the architecture must make the next source cheaper than the last.
- AI scoring will be wrong early. Without a feedback loop wired in from the first gate, wrong stays wrong. The studio's corrections are the training data, and the studio must commit the review time that produces them.
- The thesis moves. A scoring system that cannot be versioned and re-run against the stored database hard-codes last year's beliefs.
- Three principals run three live programs. The system must demand minutes of their week, not hours, or it joins the list of tools that made work instead of removing it.

The signal we leave with

The studio does not need a better way to judge technology. It needs to see more of it, and to stop paying its scarcest people to do a machine's reading. The first move is the funnel: scrape the inventory that is already for sale, store everything, screen against the thesis, gate what survives, and put a short intelligent brief in front of a human. Sequence the sources cheap-first, prove the inventory insufficient before funding anything upstream, and wire the feedback loop in from the first day so every human decision sharpens the machine. The plan, sized by impact, is the Playbook and Delivery Proposal.

Playbook and Delivery Proposal

The playbook and the delivery proposal are one document because they are one act. The playbook says where AI fits and sizes the moves in order. The delivery proposal scopes the build for the moves you choose to start with. The first earns the second. Nothing past the first move is committed until the first move proves the approach in your business.

Part One: The Playbook

A written read for operators, not a deck. It answers three questions: where AI fits in this studio's sourcing and where it does not, the highest-leverage moves in sequence, and what it takes to run them.

Where AI fits, and where it does not

It fits the gap between the search space and the team. The space that generates medtech IP is effectively infinite: hundreds of commercialization sites, a research literature publishing continuously, government labs, inventors. The team is three principals already running three programs. AI fits the reading, the screening, the scoring, and the briefing. It does not fit the thesis, the feasibility judgment, the diligence, or the decision, which belong to the people whose judgment is the studio's product. The evidence is in the Day One Audit. The short version: the bottleneck is reading, not judgment. Machines read. People judge.

How to read the roadmap

The first two moves we diagnosed in the opening session, and we can size them against the studio's own numbers. The third is the expansion pattern the architecture is designed for; it is proven by design and each source still earns its slot. The last two we saw the shape of and did not diagnose, and we say so rather than dress them up.

Each move is read across six dimensions: time, accuracy and quality, cost and recovered revenue, growth, employee experience, and risk. The first move earns the right to the next.

The roadmap at a glance

#	Move	Status	Leverage	Containment	Why it sits here
1	The reading machine: scrape the inventory, store everything	Diagnosed, sized	Highest	A few sources, watched by hand	Nothing downstream exists until the funnel has volume. Start here.
2	The qualification matrix: thesis screen, gates, feedback loop	Diagnosed, sized	Highest	Scores advisory until proven	Volume without filtering is noise. This is where the studio's judgment gets encoded.
3	Source expansion on the proven engine	Pattern proven, each source earns its slot	High	One source at a time	The engine pays for itself every time a new source costs less to add than the last.
4	Upstream: the journals and the labs	Candidate, half diagnosed	Medium	A bounded experiment	Only when the downstream inventory proves insufficient. The session's explicit hypothesis test.
5	The inverted funnel: submissions, scored deal flow, the fund	Candidate, not diagnosed	High, distant	None yet	Built on the database the first moves accumulate. Named so nobody forgets where this goes.

Move 1: The reading machine (start here)

Agents that read the commercialization sites the way the CEO does, continuously, and a database that keeps everything they find. Every opportunity logged with its source, its researchers, its summary, and its first-pass categorization, relevant or not, because the thesis will change and re-scraping is waste. Runs watched by hand at first, one source at a time, until the categorization earns trust.

- **Time:** the starved step unblocked. Reading moves from a few dozen opportunities a month of one person's evenings to hundreds a week by machine, with headroom to thousands. The CEO's sourcing hours fall toward zero while coverage multiplies.
- **Accuracy and quality:** a machine does not skim, does not tire, and reads the fiftieth listing of the day as carefully as the first. First-pass categorization is checked by a human while trust is earned, and every correction is kept.
- **Cost and recovered revenue:** the build runs on commodity infrastructure, [a modest monthly run cost] all in. Against it: the unknowable miss rate of hand reading, where one found deal repays the

system by multiples, and the quarter-million-dollar upstream alternative deferred until the cheap inventory is exhausted.

- **Growth:** the database is an appreciating asset from the first week. Every scrape adds to a proprietary record of medtech invention that no competitor has a reason to build.
- **Employee experience:** the CEO stops being the funnel. The evenings spent re-running searches that remember nothing become minutes reviewing what the machine flagged.
- **Risk:** low and contained. A few sources, runs watched by hand, nothing downstream depends on it yet. The scraping posture per source is a named decision, made deliberately, source by source.

Move 2: The qualification matrix

The studio's judgment, written down and run by machine. First the thesis screen: device not drug, the device class, the timeline, the budget, an engineering problem rather than a science project. Then the gates on what survives: IP landscape, regulatory pathway, acquisition activity, each scored, weighted, and versioned, with feasibility flagged for human eyes rather than scored with authority. Every opportunity that reaches a principal arrives as an intelligent brief: what it is, why it fits the thesis, what the gates found. Every principal decision flows back in as a correction the matrix learns from.

- **Time:** a principal's read of a candidate drops from a sitting with the source material to minutes with a brief. The five-thousand-dollar diligence step is reached only by candidates that earned it.
- **Accuracy and quality:** the screen runs identically on everything, instead of living in one head that is sometimes tired and sometimes traveling. Wrong early, by design, and wired to get less wrong every week, because the feedback loop is built in from the first gate, not bolted on.
- **Cost and recovered revenue:** the costliest failure mode, specialist time burned on weak candidates, falls directly. Diligence spend concentrates on the three to five a quarter the studio actually wants to see.
- **Growth:** a versioned matrix re-scores the entire stored database overnight when the thesis moves. Last January's near-miss resurfaces by itself when the market turns it into a fit.
- **Employee experience:** the regulatory lead and the science advisor see candidates worth their attention, with the obvious kills already killed. The work that reaches judgment deserves judgment.
- **Risk:** the methodology is the company, so the matrix is a trade secret: NDA-protected, never patented, access held tight. Scores stay advisory until the feedback loop proves them. The named risk is automation bias, a human rubber-stamping the machine; the brief format and the review rhythm are designed against it.

Move 3: Source expansion on the proven engine

The same engine, pointed at the next source, and the next. The architecture's promise is that each addition costs less than the last, converging on a source onboarded from little more than its address. Each source still earns its slot: a named owner decides it is worth reading, the scraping posture is decided deliberately, and the first runs are watched before the source joins the standing rotation.

- **Time:** coverage compounds without headcount. The path from a handful of watched sources to a standing rotation across the commercialization landscape.
- **Accuracy and quality:** every new source passes the same gate before it is trusted: categorization checked by hand on its first runs, error patterns logged, the matrix adjusted where the source's shape demands it.
- **Cost and recovered revenue:** the marginal source approaches zero build cost, which inverts the usual economics of coverage. Breadth stops being a budget line.
- **Growth:** this is the move that makes the database defensible. Coverage no competitor will replicate by hand, accumulating daily.
- **Employee experience:** adding a source stops being a project. It becomes a decision.
- **Risk:** low, gated per source. A source that misbehaves is dropped from the rotation without touching the rest.

The later moves: named, not yet diagnosed

- **Move 4: Upstream, the journals and the labs.** The session's explicit hypothesis test. The downstream inventory is the cheap test; if it proves insufficient, the engine points upstream at the research literature, where the technology is earlier, riskier, and unmarketed. The engine will not care; the funnel is identical. What changes is the thesis screen's job, because raw research needs a use-case read before it is a candidate. Not committed until the downstream pool runs thin.
- **Move 5: The inverted funnel.** The database and the scoring, once proven, become the studio's gravity: researchers and commercialization offices submitting to the studio, scored deal flow offered to other firms, the fund structure the model supports. Built on everything the first moves accumulate. Named here so the early decisions, store everything, version the scoring, protect the method, are made with this destination in view.

Each of these is a contained bet with its own measurable result. None is committed now. They earn their turn only after the moves ahead of them prove out.

What it takes to run the moves

The discipline matters more than the technology.

- **Protect the method before the first line of code.** NDA in place, trade secret posture explicit, no patent filings on the methodology, access limited to people under the NDA.
- **Store everything, discard nothing.** Relevance is a function of the current thesis, and the thesis moves. The database is the asset; the scores are just today's read of it.
- **Version the judgment.** The thesis screen and the gate weights are versioned like code. Every change is a record: what changed, why, and what it re-scores.
- **Keep a human on every gate while trust is earned.** Scores advise, people decide, and every decision feeds back. Autonomy is granted by evidence, never by schedule.
- **Spend specialist time only downstream of the machine.** The five-thousand-dollar hours go to candidates the funnel has already earned confidence in.

The plays that run each canon come from our reusable plays library. The ones selected for this engagement are instantiated in the Charter. ## Who runs it

This can run with your own team, with another firm, or with us. It needs a few clear accountabilities: someone who owns the thesis and its versions, which is the CEO; the specialists who judge what the machine surfaces, which the studio already has; and a builder who runs the engine, the agents, the database, and the matrix. The judgment seats are filled. The engineering seat is the one you would bring in.

The recommended first move and the 90-day frame

Start with the reading machine and the matrix together, sequenced inside one build: the funnel has no value empty, and the matrix has nothing to score without it. The first 90 days: the engine reading its first sources with every find stored, the thesis screen running with human checks on its calls, the first gates scoring with the feedback loop live, and the studio's principals reviewing intelligent briefs instead of raw listings. Prove the funnel on the cheap inventory first. The journals wait until the inventory says they are needed.

Part Two: The Delivery Proposal

The proposal to build the playbook's first moves: the reading machine, the qualification matrix, and the expansion architecture. Scoped only after the playbook showed what is worth building.

What we understand

A venture studio whose deal flow runs on one person's reading hours, in a market that generates more investable technology than any team could read by hand. A thesis sharp enough to encode. Specialists whose judgment is the studio's edge, currently spent partly on a machine's work. A methodology that is

the company, to be protected as a trade secret. And a target: three to five vetted opportunities a quarter, surfaced by a system that learns from every decision.

What we will build

A deal sourcing and qualification engine, private to the studio. Agents that read the commercialization sites and log everything they find. A database that keeps all of it, relevant or not. A qualification matrix that screens against the thesis and scores the gates, versioned, with a feedback loop from every human decision. A review surface that puts intelligent briefs in front of the principals: what it is, why it fits, what the gates found. Architecture built so the next source is cheaper than the last.

How we will work

Three phases, one month each, a fixed fee of [fixed fee] per phase. Each phase is independently valuable and earns the next. The engagement can stop at any phase boundary with value already in hand. The phases run the WISER arc: prove the reading cheaply, build the one funnel that earns trust, then expand it source by source, with refinement running as a standing rhythm from the first feedback loop onward.

Phase 1: Foundation

Prove the reading, and stand up the spine.

- The agent framework: one agent per source to start, version tracking built in.
- The research database: everything found is stored with source, timing, researchers, and first-pass categorization.
- The first sources live, runs watched by hand, categorization checked by a human.
- The thesis screen stubbed and running, ready to refine.
- End of phase: the machine is reading, the database is filling, and the categorization error rate is measured, not guessed.

Phase 2: Vetting intelligence

Encode the judgment.

- The thesis screen refined against the studio's corrections from Phase 1.
- The gates built: IP landscape, regulatory pathway, acquisition activity, scored and weighted; feasibility flagged for human eyes.
- The feedback loop live: every principal decision recorded against the machine's score, and the matrix versioned with every change.
- The intelligent brief: each surfaced candidate arrives readable in minutes, thesis alignment stated.

- End of phase: the funnel surfaces scored candidates a principal can act on, and the system is measurably better than it was a month earlier.

Phase 3: Scale and interface

Make expansion cheap and the rhythm standing.

- The expansion architecture: source onboarding driven by configuration, converging on a new source from its address.
- The source rotation grown, each addition gated and watched before it is trusted.
- The review surface: the internal tool the principals work from, built for the quick read, no polish where polish buys nothing.
- Documentation: how to add sources, adjust the matrix, version the thesis, and operate the system.
- End of phase: a production engine the studio operates, reading at a scale no hand process approaches, with the quarterly review rhythm running on its output.

What we need from you

- The NDA executed before work begins, drafted for trade secret protection.
- The thesis as you apply it, and a working session to encode it: what kills a deal, what makes one, and the weights you carry in your head.
- Scored examples: a handful of opportunities your team has already evaluated, as the matrix's first training data.
- Priority sources: which commercialization sites matter most, and alternates if a source proves expensive to read.
- A weekly half hour to review the machine's calls while trust is earned. Your corrections are the training data.

Infrastructure

You provide the dedicated machine the agents run on, the AI and cloud accounts at [a modest monthly run cost], and the storage the system lives in. We provide the build: agent architecture, the database, the matrix, and the implementation. The system and its data are yours, on your infrastructure, under your NDA.

Who is working on this

A senior practitioner who leads the engagement, designs the system with the CEO, and builds it. Your principals fill the judgment seats: the thesis, the regulatory read, the science. Small team, close to the work.

Investment

Three phases at a fixed fee of [fixed fee] per phase, [total] for the build. Run costs are yours directly and stay modest. Each phase stands alone, and the engagement can stop at any boundary with the value already delivered.

Charter

What a Charter is

Not a project plan. Not a requirements document that executes once and collects dust. A Charter is the memory that survives the chaos. Its value is the decision log: when someone asks a year later why the system reads the commercialization sites instead of the journals, or why the discard pile gets audited as carefully as the survivors, the answer is here, with the alternatives that were weighed and the evidence that settled it. The Architect keeps it current, same-day.

Metadata

Field	Value
Project	A deal sourcing and qualification engine, private to the studio
Client	The venture studio (anonymized)
Charter Keeper	The Architect
Dates	Held in the private client record; relative markers used here
Current canon	Refine. The engine is live, reading at scale, learning from every review.
Version	Running state

Positions

The work was held together by clear accountabilities, not an org chart. On a team this small, people hold several Positions, and that is by design.

Position	Who held it	Tension owned
Sponsor	The CEO	Authority. Owned the why, decided without a committee, cleared the way.
Thesis owner	The CEO	Integrity. The matrix encodes his judgment; every version of it is his to approve.
Guide	First Strategy senior practitioner	Translation. Carried the method and kept the Charter honest.
Architect	First Strategy	Curiosity and stewardship. System design and Charter Keeper.
Builder	First Strategy	Execution. The agents, the database, the matrix, the review surface.
Sage	The regulatory lead and the chief science advisor	Context. What a predicate path looks like, and what separates engineering from magic.
Safety	Outside IP counsel	Protection. The trade secret posture and the IP landscape judgment the machine's gate approximates.

As the system proved reliable, the machine took over work inside documented constraints, with the humans shifting from doing to directing and reviewing. The Positions did not change. What changed is how much reading sat underneath them.

Objectives and constraints

The build specification: what the project set out to do and the lines it would not cross.

Scope

In scope: the sourcing and qualification funnel, from the commercialization sites to the brief in front of a principal. The agents, the database, the thesis screen, the gates, the feedback loop, the review surface, and the expansion architecture. Out of scope throughout: the diligence itself, the negotiations, the programs, and the thesis. The system applies the studio's judgment; it does not replace it, and it does not touch the live programs.

Objective and success criteria

Multiply the studio's reading capacity by orders of magnitude without adding headcount, protect the specialist diligence spend with gates that earn trust, and accumulate the proprietary database the studio's long game is built on.

Measure	Baseline	Target	Result
Reading capacity	One opportunity at a time, by hand, in the CEO's hours	Continuous machine reading across the source rotation	Hundreds of opportunities a week, with headroom to thousands
Cost of adding a source	A build per source	Each source cheaper than the last	A new source onboards from its URL
What reaches a principal	Raw listings and source material	An intelligent brief: what it is, why it fits, what the gates found	Briefs running; the read is minutes, not sittings
The record	No memory between searches	Everything logged, relevant or not	A growing proprietary database of the IP landscape, re-scorable as the thesis moves

Constraints

- The methodology is a trade secret. NDA before work, no patents on the method, access limited to people under the NDA.
- Store everything. No opportunity is discarded, because relevance is a function of a thesis that moves.
- Scores advise, people decide. No candidate reaches diligence, and none dies a final death, without a human in the loop while trust is earned.
- The principals' time is the budget that matters. The system asks for minutes a week, not hours.
- The studio's infrastructure, the studio's data. The engine runs on their dedicated machine, under their accounts.

Architecture and human-in-the-loop design

A funnel, top to bottom. Agents read the commercialization sites and log everything they find into the research database: source, researchers, summary, first-pass categorization, timestamps. The thesis screen runs on everything logged: device not drug, the device class, the timeline and budget shape, engineering problem not science project. The gates score what survives: IP landscape, regulatory pathway, acquisition activity, each weighted, with feasibility flagged for human eyes rather than scored with authority. What clears the gates surfaces to the principals as an intelligent brief with its thesis alignment stated. Every human decision on a brief, pursue, pass, or correct, flows back into the matrix.

The matrix is versioned like code. The thesis screen and the gate weights carry version numbers, every change is logged with its reason, and a new version can re-score the entire database overnight without touching a source. The judgment stays human at the top: the CEO owns the thesis, the specialists own feasibility, and the machine owns the reading.

During the build the humans watched everything: scrape runs reviewed one at a time, categorization calls checked by hand, every kill sampled. The grip loosened tier by tier as the error rates earned it, under the Hierarchy of Agency below, never before.

Current state at the start

Carried from the Day One Audit. A three-principal studio running three live programs, with a sharp thesis and the specialists to vet what survives it. Sourcing by hand: a custom GPT with no memory, a day of CEO time per university negotiation, a quarter-million-dollar funded research project as the price of going upstream. About five thousand dollars of specialist time per candidate evaluated. No database, no change detection, no compounding. The bottleneck was reading, not judgment.

Decision log

The decisions that shaped the build, each with the alternatives weighed and the evidence that settled it. This is the part of the Charter that answers "why did we do it this way."

When	Decision	Alternatives rejected	Rationale	Evidence
The session	Start downstream, at the commercialization sites	The journals first; fund more upstream research	The transfer inventory is engineered, often validated, and for sale; prove it insufficient before paying to go upstream	A day of CEO time per university deal; the quarter-million-dollar upstream precedent
The session	Store everything, discard nothing	Keep only what matches the current thesis	The thesis moves; re-scoring a stored record is free, re-scraping the world is not	The studio had already shifted its own market hypothesis once since founding
The session	Trade secret by NDA, never patented	Patent the methodology	A patent discloses the recipe; the method's value is that nobody knows how it scores	IP counsel's posture, named in the session
The session	The machine scores three gates; feasibility stays human	Score all five gates	IP landscape, regulatory path, and M&A activity are data; manufacturability and testability are judgment the specialists own	The specialists' own definition of their work
Foundation	One agent per source, runs watched by hand	One generic scraper pointed at everything	Every site publishes differently; trust is earned per source, and the error patterns differ	First watched runs: each source failed differently before it ran clean
Foundation	Measure the categorization error rate from the first run	Tune by feel	A screen that is wrong in unmeasured ways cannot earn autonomy	The watched-run checklist; every human check logged
Vetting	Scores advisory, feedback loop live from the first gate	Build the gates, tune them later	Wrong early is the plan; the studio's corrections are the training data, and they only exist if the loop is running	The screen's early calls against the principals' calls
Vetting	Version the matrix; every change re-scores the database	Static thresholds adjusted in place	An unversioned screen hard-codes last year's beliefs and cannot explain its own history	The first thesis adjustment, re-scored overnight against everything already stored
Scale	Generalize onboarding: a new source from its URL	Keep hand-building an agent per source	The architecture's promise was each source cheaper than the last; per-source builds break it	Successive source builds shrinking until configuration replaced code

When	Decision	Alternatives rejected	Rationale	Evidence
Scale	The review surface is a brief, not a dashboard	A full internal dashboard with views and filters	The principals' budget is minutes; the brief answers pursue or pass, and everything else is the database's job	The quick-read format holding in the principals' actual week
Refine	Per-source volume baselines and run monitors	Trust runs that return without errors	Scraping fails silently; a source that changes shape returns thin, not broken	A source restructure caught by its volume drop, not by an error
Refine	Audit the kills, not just the survivors	Review only what surfaces	A false kill is invisible forever; sampling the discard pile is the only way to see the screen's blind side	Kill-sample reviews feeding corrections the surfaced pile never would

The decision and experiment record

The supporting narrative behind the log. The project ran the full WISER method.

Witness

The session had already found the diagnosis: the bottleneck was reading, and most of the loss was invisible because a missed deal makes no sound. The build's first weeks added the texture the session could not: what the sources actually look like up close. The commercialization sites share a purpose and almost nothing else. Each publishes its own way, structures its listings its own way, and buries its technologies at its own depth. The watched runs were the floor time of this engagement: the machine reading, a human checking every call, and the error patterns accumulating into the first real picture of what each source demands.

Interrogate

Two questions had to be answered cheaply before the build earned its budget: can the sources be read reliably, and can a machine screen what they say. The first sources answered the first question source by source, each failing differently before it ran clean. The second question was answered the only honest way, by measuring: the screen's calls checked against a human's on every early run, the error rate tracked, the misses logged with their patterns. The screen was wrong in instructive ways early, and the instructive part was the point: every error pattern became a correction, every correction made the next run measurably better, and the trend line, not any single run, is what earned the screen its trust.

Solve

The funnel closed end to end when the first intelligent briefs surfaced: a candidate technology, what it is, why it fits the thesis, what the gates found, readable in minutes. The trust moment was not a demo. It was the CEO working through a morning's briefs in the time a single raw listing used to take, and the realization landing that the reading he had been doing in evenings and weekends was now happening continuously, without him. The funnel was doing the work of the search he could never staff: hundreds of opportunities a week flowing through the screen, the obvious kills dying quietly, the survivors arriving scored and explained.

Expand

Expansion was sources, and the architecture was the strategy. The early sources were hand-built agents, each a small project. The pattern across them converged: the differences between sources were configuration, not code, and the build generalized until adding a source meant giving the engine a URL and watching its first runs. Each new source still earns its slot, with its scraping posture decided deliberately and its first runs checked by hand before it joins the standing rotation. Coverage compounds, and the marginal source costs almost nothing. The downstream inventory has not yet run thin, so the journals, the explicit hypothesis test from the session, remain unspent: the cheap pool is still paying.

Refine

The standing lesson of the Refine canon arrived the way it always does, quietly. A source restructured its site, and the agent reading it did not break. It returned thin, run after run, errors zero, volume down. The per-source volume baselines caught what no error log would have: a healthy-looking run reading a fraction of what the source actually held. The fix was routine; the rule it produced was not. Every source now carries its own baseline, every run is judged against it, and silence is treated as a symptom. The same humility runs the screen: the kills get sampled, because a false kill is invisible forever, and the feedback loop runs on both piles, the surfaced and the discarded.

Hierarchy of Agency

Three tiers of human oversight, by the cost of being wrong. The tier governs how much human attention each kind of call gets.

Tier	Oversight	Applies to
1: Autonomous	Runs unwatched; monitored by baselines and logs	Scraping, logging, change detection, summarization. Earned after watched runs per source
2: Machine with sampling	The machine calls it; humans sample both piles on a cadence	The thesis screen's kills and passes; gate scores. The kill sample is non-negotiable
3: Human-led	The machine assists; a person decides	Surfacing to diligence, the diligence itself, thesis versions, gate weights, scraping posture per source

A call type moves to a lighter tier only on evidence: a sustained error rate the principals accept, measured against human checks, over enough runs to mean something. Nothing earns it by decree. If a tier drifts, it falls back to heavier review. The thesis and the matrix versions never leave Tier 3; they are the company.

Risk register

Risk	Mitigation	Status
The methodology leaks	NDA-gated access, trade secret posture, no patent filings, the recipe never in client-facing documents	Held
A source prohibits scraping	Scraping posture decided per source, deliberately, with alternates ready; the rotation survives losing any one source	Managed per source
Silent scraper failure	Per-source volume baselines; runs judged against history, not just error logs	Realized once; caught by the baseline; rule now standing
The screen kills good deals	Kill sampling on a cadence; corrections feed the matrix; the discard pile is permanent, so a recovered thesis can resurrect it	Active control
Automation bias at the brief	The brief states evidence, not just a score; principals decide; pursue-and-pass decisions logged against the machine's call	Active control
The thesis hard-codes	Versioned matrix, owner named, every change logged and re-scored against the database	Held
The principals' attention budget	Minutes-a-week design target; the brief format; the review rhythm sized to the three to five a quarter that reach diligence	Held

Evolution history

How the oversight posture changed over time, and why.

When	Change	Trigger
Foundation	Every run watched, every call checked	Trust not yet earned
Vetting	Scrape and log to Tier 1 on the proven sources	Watched runs clean; error patterns understood per source
Vetting	The screen to Tier 2 with both piles sampled	The screen's error rate trending down against human checks
Scale	Source onboarding generalized to URL-driven configuration	Successive source builds converging to configuration
Refine	Volume baselines per source; silence treated as a symptom	The thin-run incident
Refine	The feedback rhythm settled into the studio's quarterly review	The funnel's output matching the cadence the diligence spend runs on

Current status and what transfers

The engine is live and reading. The studio reviews hundreds of opportunities a week through it, with headroom to thousands, against a hand baseline of one at a time. Every opportunity it has ever read is logged, relevant or not, and the database grows with every run: a proprietary record of the medtech IP landscape that re-scores overnight when the thesis moves. New sources onboard from a URL. The principals work from intelligent briefs with thesis alignment stated, and their decisions keep teaching the matrix. The three-principal studio runs the deal review of a hundred-person firm, and the specialists' five-thousand-dollar hours are spent only on candidates the funnel earned.

What transfers is the discipline as much as the system. The thesis is theirs and versioned. The decision log is theirs. The kill sample, the volume baselines, and the review rhythm are operating habits now, not project artifacts. The destination named in the session still stands: the studio that has read everything is the studio researchers come to, and the database this engine accumulates is that gravity, building daily.

Outcomes

- Reading capacity moved from one opportunity at a time, by hand, to hundreds a week by machine, with headroom to thousands.

- Better qualified deals with the noise sorted by machine instead of human labor: the principals see scored, explained candidates, not raw listings.
- A proprietary and growing database of the medtech IP landscape, logging every opportunity seen, relevant or not, re-scorable as the thesis moves.
- New deal sources onboard from a URL, so coverage compounds without builds.
- The specialist diligence spend is protected by gates that earned their trust against measured error rates.
- A three-principal team running the deal review of a hundred-person firm, without the hundred people.
- The methodology protected as a trade secret throughout: the studio owns the system, the data, and the recipe.

Plays

The WISER plays this engagement ran, instantiated with the client's specifics. This is the index and what each produced. The high-value plays are held as standalone documents; the rest were applied inline in this Charter. | Canon | Play | What it produced | Source | |-------|------|------------------|--------| | Witness | User Flow Mapping | The deal flow traced end to end, one opportunity from lab to decision | Standalone play: Deal Flow Trace | | Witness | Documenting Current State | The hand-sourcing record: the GPT searches, the day-per-university cost, the missing memory | Inline in the Day One Audit | | Interrogate | Assumption Auditing | The register of instincts tested, upstream-first above all | Standalone play: Assumption Register | | Interrogate | Experiment Logging | The watched runs and screen-versus-human measurements that earned each trust step | Inline above | | Solve | Human-in-the-Loop Design | The advisory-score funnel, the brief, and the feedback loop on both piles | Standalone play | | Solve | Quality Objective Setting | The error-rate targets the screen had to hit before each tier move | Inline above | | Expand | Deployment Gating | What a new source must pass before joining the rotation | Standalone play: Source Onboarding Gate | | Refine | Drift Monitoring | The volume baselines and the kill sample | Standalone play | | Refine | Hierarchy of Agency Design | The three tiers and the evidence that moves a call type between them | Inline above |

The first funnel is built, and it will not be the last decision the database underwrites. The journals are still waiting upstream, the inverted funnel is still the destination, and the difference now is that the studio's reading never stops, never tires, and never forgets what it read.

The plays

The WISER plays this engagement ran, instantiated with the client's specifics, ordered by canon.

WITNESS PLAY

Deal Flow Trace

Witness play, instantiated for the venture studio engagement. Purpose: trace one unit of work end to end through every handoff. The unit here is one piece of medtech technology, followed from the lab that produced it to the studio's decision about it, as the work ran at intake.

The flow, at a glance

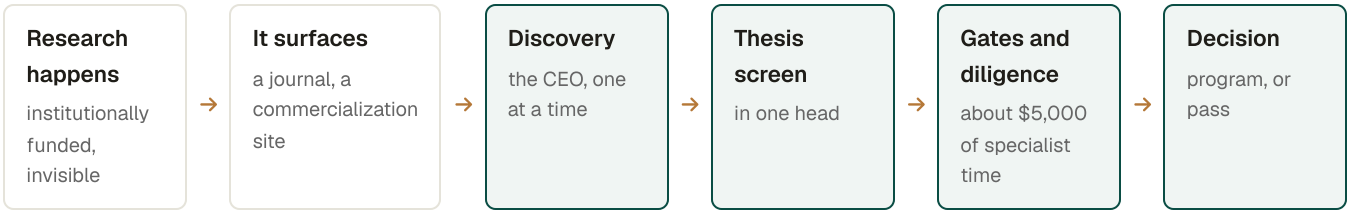

No loop: what the evaluation taught went nowhere. No database, no memory, no re-finding the near-misses.

The trace

Step	Who carries it	Where it breaks
Research happens in a lab	The institution and the researcher	Invisible to the studio. Researcher incentives end at publication and patent
The work surfaces	A peer-reviewed journal or a commercialization office	Public and mostly unread. Nobody is watching systematically
Discovery	The CEO, by hand: searches, a custom GPT, conferences, relationships	The starved step. One person, one opportunity at a time, against an effectively infinite space
Thesis screen	The CEO, from memory	Runs only when he is reading. Device not drug, the device class, 18 to 30 months, one to two million dollars
The gates	The CEO plus specialists	IP landscape, regulatory pathway, acquisition activity, feasibility. The real versions are specialist work
Deep diligence	IP counsel and the regulatory lead	About five thousand dollars of specialist time per candidate. Burned on every weak candidate that gets this far
The pursuit	The CEO	A full day for one university negotiation, in the engagement record
The decision and the learning	The principals	A program starts or the candidate dies, and either way the system forgets what it learned

The root-cause read

Steps one and two are the market and they are healthy: the institutions keep funding research and the research keeps surfacing. Steps four through eight are the studio's strength: a sharp thesis and the specialists to apply it. The flow starves at discovery, where one person's reading hours are the ceiling on everything downstream, and it forgets at the end, where evaluations leave no record. Most of the loss is invisible, because a missed deal makes no sound. The build brief fell straight out of the trace: put a machine on the reading, keep the people on the judgment, and never forget anything again.

INTERROGATE PLAY

Assumption Register

Interrogate play, instantiated for the venture studio engagement. Purpose: surface the assumptions the build was about to spend money on, including the CEO's own strongest instinct, and test each one cheaply before it got expensive.

The verdicts, at a glance

REFRAMED

The gold is upstream

True, and not first. The downstream inventory is the cheap test; the journals wait until it runs thin.

CONFIRMED

The cheap inventory can feed the studio

The commercialization sites alone moved reading to hundreds a week. The pool has not run thin.

REFRAMED

A machine can run the gates

Three of five. IP, regulatory, and M&A are data. Feasibility stayed human, flagged not scored.

KILLED

Every source needs its own build

The differences converged to configuration. A new source now onboards from its URL.

REFRAMED

A trusted screen can be left alone

A false kill is invisible forever. The discard pile gets sampled on a standing cadence, trust or no trust.

The register

#	Assumption	Source	Cheap test	Verdict
1	The best technology is upstream, in the research, so the system should start there	The CEO's read of the market	Cost the two paths on the whiteboard: sitting inventory someone is paid to sell, against literature that costs a quarter million dollars to act on	Reframed. The instinct is right about where undiscovered technology lives and wrong about sequence. Exhaust the cheap inventory first; prove it insufficient before the journals earn their build
2	The commercialization inventory is deep enough to feed the studio's targets	The session's sequencing bet	Point the engine at the sites and count what survives the screen	Confirmed so far. Reading runs at hundreds a week with headroom, and the downstream pool has not run thin. The journals remain unspent
3	A machine can score the investment gates	The build's premise	Build the gates against the specialists' own definitions and measure the calls	Reframed. IP landscape, regulatory pathway, and acquisition activity are data disciplines the machine scores. Manufacturability and testability are judgment; the machine flags, the specialists decide
4	Each source needs its own hand-built agent	The first builds	Build several and watch what actually differs	Killed. The differences were configuration, not code. Onboarding generalized until a new source costs a URL and a watched first run
5	Once the screen's error rate is earned, its kills can be trusted	The natural drift of a working system	Sample the discard pile on a cadence and look for false kills	Reframed. The sampling stays forever. A wrongly surfaced candidate costs minutes; a wrongly killed one is invisible and unpriced

Why assumption 1 was tested at all

It was the CEO's sharpest insight, which is exactly why it had to be tested. He was right that the best technology often never reaches the commercialization offices, and the studio had the quarter-million-dollar research project to prove how expensive acting on that insight is. The test did not disprove the insight. It found the order it gets acted on in: the inventory that is already engineered, validated, and for sale is the cheap test, and the engine was built so that pointing it upstream later changes a target list, not an architecture.

SOLVE PLAY

Human-in-the-Loop Design

Solve play, instantiated for the venture studio engagement. Purpose: define who reviews AI output, how, and what gets logged, so the system can be wrong early without being dangerous, and so every human decision makes the machine better.

The funnel, at a glance

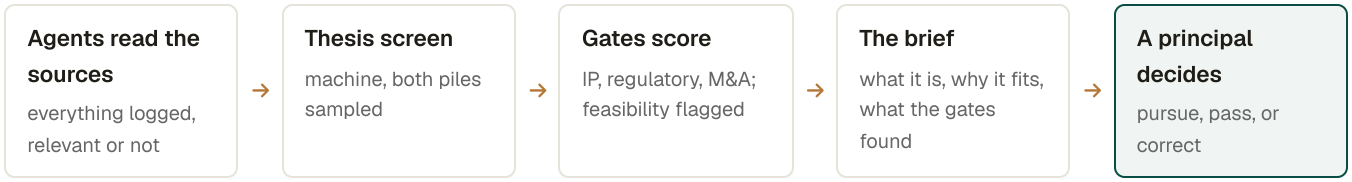

↺ Every decision flows back: corrections version the matrix, and the matrix re-scores the database.

The design

Element	The rule	Why
What the machine produces	Logged opportunities, screen calls, gate scores, and the brief	The reading and the arithmetic are machine work
Who reviews it	The CEO on thesis calls; the regulatory lead and science advisor on what reaches diligence	The judgment seats were already filled; the system routes to them instead of around them
The score's authority	Advisory until proven. A score recommends; a person decides	Wrong early is the plan. Authority is granted by measured error rates, never by schedule
The kill sample	A standing sample of the discard pile, reviewed by a human	A false kill is invisible forever. The surfaced pile audits itself; the discard pile has to be forced into view
What gets logged	Every screen call against every human check; every pursue and pass against the machine's score; every matrix change with its reason	The log is the training data and the audit trail. A correction that is not logged teaches nothing
The brief's format	What it is, why it fits the thesis, what the gates found, readable in minutes	The principals' budget is minutes a week. Evidence shown, not just a number, to resist rubber-stamping
The feedback rhythm	A weekly half hour while trust was earned, settling into the studio's quarterly review	The corrections are the point. The rhythm is sized so it survives contact with three live programs

What it protected

The five-thousand-dollar step. Every candidate that reaches deep diligence spends specialist time, and the funnel's whole economic argument is that machine judgment is cheap and human judgment is not. The loop design makes the cheap judgment improve continuously while keeping the expensive judgment in charge of everything that matters: what the thesis says, what reaches diligence, and what becomes a program. The machine got better every week precisely because the humans never stopped deciding.

EXPAND PLAY

Source Onboarding Gate

Expand play, instantiated for the venture studio engagement. Purpose: set the gate each new deal source must pass before going live in the standing rotation. Expansion in this engagement was sources, not locations, and the gate is what let onboarding get cheap without getting careless.

The gate, at a glance

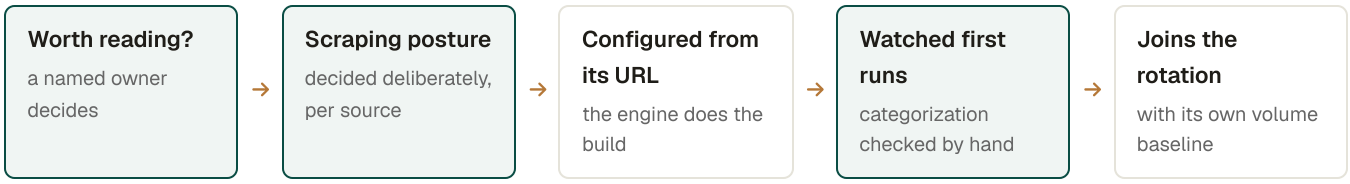

A source that misbehaves is dropped from the rotation without touching the rest.

The gate

Check	The rule	What it prevents
The case for the source	A named owner says why this source is worth reading, against the thesis	Coverage for its own sake. Breadth is cheap now, which is exactly when it needs a reason
Scraping posture	Terms reviewed, posture chosen deliberately, alternates named if a source is closed	A rotation built on access that can vanish, and decisions made by default instead of on purpose
Configuration, not code	The source onboards from its URL through the standard path	Bespoke builds creeping back in. If a source demands real code, that is a decision, not a drift
Watched first runs	First runs reviewed by hand: what it found, how it categorized, where it failed	Trusting a source's shape before it has shown its shape. Every source failed differently before it ran clean
The baseline	The source enters the rotation with its own expected volume and cadence recorded	Silent failure later. The baseline set at onboarding is what drift monitoring judges against

The economics it protects

The architecture's promise was that each source costs less than the last, converging on a URL and a watched run. The gate is what keeps that promise honest: cheap onboarding without the gate becomes indiscriminate onboarding, and the rotation fills with sources nobody can name a reason for. The gate

keeps the marginal source cheap and deliberate at the same time, which is why coverage compounds without the noise compounding with it.

REFINE PLAY

Drift Monitoring

Refine play, instantiated for the venture studio engagement. Purpose: watch for drift, including drift hidden in aggregates, in a system whose characteristic failure is silence. A scraping and scoring engine rarely breaks loudly. It returns thin, or it scores stale, and the dashboards stay green.

The incident that wrote the rule

- **In rotation**
  A proven source, runs clean, errors zero.
- **The site restructures**
  The agent does not break. It reads what it still recognizes, which is a fraction of what is there.
- **Thin runs, green logs**
  Run after run returns healthy: errors zero, volume quietly down.
- **The baseline catches it**
  The source's volume history flags the drop. Reconfigured, re-watched, returned to rotation.
- **The standing rule**
  Every source carries its own baseline. Silence is a symptom.

What is watched

Signal	Against what	Why
Volume per source, per run	The source's own baseline, set at onboarding	A restructured site fails silently. Errors measure breakage; only history measures thinness
The screen's calls	A standing human sample of both piles, surfaced and discarded	The discard pile is where drift hides. A false kill is invisible unless it is forced into view
Pursue and pass decisions	The machine's score on the same candidates	A widening gap between what the matrix scores and what the principals choose is the matrix aging
The thesis itself	The versioned matrix, every change logged and re-scored against the database	The market moves. An unversioned screen hard-codes last year's beliefs and cannot say when it stopped being right

The response pattern

Contain, diagnose, correct, and keep the lesson. The thin source was pulled from trust, not from rotation: its reads continued but its output was flagged while the agent was reconfigured and re-watched, the same gate a new source passes. The correction went into the engine; the lesson went into the standing rules. Nothing was credited to luck. The baseline that caught it became mandatory at onboarding, which is how an incident in one source bought protection for every source after it.

The posture

Aggregate health is not health. The engine's totals looked fine while one source ran thin, the same way a portfolio looks fine while one position quietly dies. Drift is watched per source, per pile, and per version, because that is where this system drifts: in the segments, in the silences, and in the slow divergence between what the machine believes and what the people decide.