FIRST STRATEGY · CASE STUDY

Communications and Public Relations

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

Two problems that should not have been in the same building

A national public relations agency was burning tokens haphazardly and could not get its people to use AI.

Both. At once. Some of the firm was spending real money on AI without much to show for it, and most of the firm was not touching it at all. Either problem alone has an obvious answer. Together they cancel out, because the usual fix for one makes the other worse. Tighten spending and you confirm every doubt the holdouts already have. Push adoption and you multiply the waste.

They brought us in to operationalize AI and to build communications agents their team would actually use.

What was at stake

The asks were specific, and leadership named them without prompting. Increase practitioner output. Build consistency across every client the firm serves. Build company memory that could be drawn on for any deliverable or any pitch.

Underneath was something less comfortable. Agents were being stood up in a no-code builder and pointed at critical systems, without service design and without anyone deciding what the machine was allowed to do. There was no AI governance in any real sense: nothing written down about what had to be true before AI output could be used, and nobody who owned that question.

And the practitioners were against it. Not indifferent. Against. In a craft business, where the product is judgment and voice and relationships, a workforce that believes the machine is coming for the thing they are good at will find ways to be right about that.

Ungoverned tools, built without planning, handed to people who distrust them. That is how an AI program dies quietly and stays dead, and the second attempt is always harder than the first.

The part where leadership was right

The conventional shape of a story like this one is that leadership believed something and the floor proved them wrong.

That is not what happened here, and the record should say so.

The CEO and the managing partners had the diagnosis correct. They knew AI was being used badly and not used at all. They knew the mundane work was eating their practitioners. What they did not have was a sustainable plan or program, and they did not know how to run service design methodology against their own operation to surface what the day-to-day work actually needed before anyone automated a piece of it.

The gap was method, not insight. That is a harder problem than a wrong belief, because there is nothing to correct. There is only work to do.

A week on the floor

We spent a week in review, and the substantial part of it was interviewing practitioners directly over several days.

There was no dramatic reveal. No dispatcher running the operation from a spreadsheet with the real software minimized. Leadership already knew that media relations carried a great deal of homework before anyone got to the work that clients pay for. What the week established was the size of it.

Trace one pitch.

A publicist starts with the news. Manually, site by site, looking for the day's angle. One to two hours a day, every publicist. Then reporter research: who covers this, who has written something adjacent, who might take it. Only after both does the publicist do the thing they were hired to do, which is craft the pitch.

The craft is the third step. The first two are homework, and they run first, every day, by hand, before the work can start.

That is the mundane work of running public relations. It is not a defect in the operation. It is the operation, and it was consuming the hours of people whose judgment was the product.

The shift: govern first, then build

Most AI programs start by building something impressive and figure out the rules afterward. Here the order had to reverse, because the workforce problem was the real constraint. Nothing would survive contact with a practitioner who believed we were there to replace them.

So the first thing built was not an agent. It was a boundary.

AI output was written into a separated directory that could not be used directly. Nothing produced there was available to a client, a deliverable, or a pitch until a person had looked at it, exercised discernment, and graduated it to approved. The machine could work. The machine could not ship.

This is not something we invented for this engagement. It is how we run AI in every client we work with, and it arrived here already proven, which is part of why it could go first.

That mechanism does three jobs at once. It answers the governance gap, because now there is a written answer to what has to be true before AI output is used. It makes the promise to the practitioners structural instead of verbal: we said we would respect the craft and the human experts doing the work, and the directory is what that sentence means in practice. And it settles what AI is for. A collaborator, not an autonomous doer.

The second decision was about sequence. The first tool through the door was writing tooling we already had, built before this engagement and proven elsewhere. Not the most impressive thing we could have shown them. The lowest-risk way to let a skeptical practitioner see what good looks like, before anything autonomous entered the building.

What got built, in order

Writing tooling, brought in rather than built. Voice-bound content work, already in use, so the first thing practitioners touched was something that had been holding up under real conditions.

A new business intake agent, which took the research burden off inbound leads before a human spent time on them.

Research, which is the answer to the first two steps of that pitch trace. The day's news, gathered agentically and delivered as a report, so a publicist arrives at reporter research and pitch craft rather than starting two hours behind it.

A knowledge graph, the company memory leadership had asked for in the first conversation, so that what the firm knows is available across deliverables and pitches instead of living in whoever happened to work the account.

Training ran alongside all of it. Not a workshop on prompt tricks. How to use AI effectively and responsibly, delivered to the people who had every reason to resist it.

What changed

The gathering stopped being human work. An agent reads across the sources and produces a report, and the publicist reads that report and decides which angle is worth a pitch. Nobody hunts site by site any more. The judgment about what is newsworthy never moved.

That distinction matters more than the hours it saves. The report is AI output like any other, so it arrives in the same place everything else does, and a person still decides what it is worth. The machine took the gathering. It did not take the choosing.

Adoption increased dramatically. Stated plainly and without a figure, because the change of state is the drama and the numbers that circulate about this engagement are the part that does not reconcile.

The more interesting outcome is who is using it. This is the workforce that was against it. The people who believed the machine was coming for their craft are the ones running it, and they came around for reasons that were earned rather than announced: they were trained properly, and the boundary held.

What they run without us

We stopped building.

Self-improvement and new agent building were taught to the firm's core enablement group, which we also trained. They build for themselves now. The capacity to make the next agent sits inside the agency rather than with us, which is the outcome we tell every client we are aiming at and the one that is easiest to quietly not deliver.

What continues is a standing monthly conversation with that group. Questions, decisions about what to build next, and oversight of the governance that made the whole thing possible in the first place.

The engagement did not end. It turned into advice, which is what capability over dependency looks like when it actually works.

Why this one is different

Every other case study in this practice turns on a discovery: the floor disproves what the boardroom believed, and the engagement pivots.

This one has no pivot. Leadership was right, the floor confirmed it, and the work was the work. What made the difference was not insight. It was method, sequence, and a boundary drawn early enough that a hostile workforce could watch it hold.

That is a less dramatic story and a more common problem. Most firms are not wrong about their AI problem. They just do not have a way to do anything about it.

Engagement Audit

The one-line finding

The firm was overspending on AI and refusing to use it at the same time, and both halves had the same cause: no governance, and no method for deciding what should be automated.

How we looked

A week of review. The substantial part was interviewing practitioners directly across several days, alongside a read of what was already running.

We did not start from the tooling. We started from the work, because the question that mattered was not what the firm had bought but what its people actually did all day.

What was already in the building

What	State we found it in
AI spend	Real, ongoing, and concentrated in a minority of the firm. No line of sight from spend to output.
AI adoption	Absent across most of the practice.
Agents	Being stood up in a no-code builder and pointed at critical systems. No service design behind them.
Governance	None in any operative sense. Nothing written about what AI was allowed to do, what had to be true before its output could be used, or who owned that question.
Practitioner sentiment	Actively against. Not indifference; opposition.

The five rows are one condition. Ungoverned tools, built without design, in a firm whose people did not want them.

The stakeholder read

Who	What they own	How they framed the problem
CEO and managing partners	The commercial outcome and the firm's standards	Correctly. Output, consistency across clients, and company memory. Named without prompting.
Practitioners	The craft: judgment, voice, relationships	As a threat to the thing they are good at.
The people already using AI	Their own accounts	As a tool they were free to use however they liked, because nothing said otherwise.

Leadership's diagnosis was right. That is worth stating plainly, because it is unusual and because it changes what the engagement had to supply. They were not missing insight. They were missing a plan, a program, and a way to run service design against their own operation to find out what the day-to-day work needed before automating any of it.

The gap was method.

The work, traced

One pitch, as it ran:

Step	What happens	Who
1	News research, manually, site by site, hunting the day's angle	Publicist
2	Reporter research: who covers this, who has written something adjacent, who might take it	Publicist
3	Craft the pitch	Publicist

Step 3 is the work clients pay for. Steps 1 and 2 are homework, and they run first, every day, before the craft can start.

Where the hours go

Step 1 alone runs one to two hours a day, for every publicist.

That is the mundane work of running public relations. It is not a defect in the operation and nobody was doing it wrong. It is simply the shape of the job, and it was consuming the hours of the people whose judgment is the product.

The firm's stated ask was to increase practitioner output. This is where practitioner output was going.

What the week did not find

No hidden system. No workaround running the operation from a spreadsheet. No belief held by leadership that the floor overturned.

Recorded deliberately. Leadership already knew media relations carried heavy homework. What the week established was the scale of it, and the fact that the firm's AI problem and its adoption problem were the same problem.

Where AI fits, and where it does not

Fits. The gathering in step 1. Reading across the sources and reporting what is there is high-volume, low-judgment, and repeats daily. What does not fit, even inside step 1, is deciding which angle is worth a pitch. Step 2 splits the same way: assembling candidate reporters is retrieval; deciding who to approach is not.

Does not fit. Step 3. The pitch is the craft, and the firm's product is the judgment inside it. Nothing about this engagement is improved by automating the thing the client is buying.

The harder constraint. No AI output can be allowed to reach a client, a deliverable, or a pitch without a person having exercised judgment on it first. This is not a quality control preference. In a firm where the workforce is already opposed, an ungoverned tool that puts a machine's words in front of a client under a practitioner's name ends the program.

Risks

Risk	Why it matters here
A second failed attempt	The first one is already underway. A second is much harder to recover.
Automating the craft	Would confirm every practitioner's fear and lose the workforce permanently.
Ungoverned agents on critical systems	Already happening, and the exposure grows with adoption rather than shrinking.
Building faster than trust	Adoption in a hostile workforce is earned in sequence, not announced.

The signal we leave with

Govern before building. Then take the homework, not the craft.

Playbook

What this is

Where AI fits in a public relations practice, in the order the moves should run, sized by what each one is worth.

The order is not by size. It is by trust. In a firm whose practitioners are opposed to AI, the sequence is the strategy: each move has to earn the right to the next one.

Move zero: the boundary

Before any agent, a governance mechanism.

AI output is written to a separated directory. Nothing there can be used directly. A person looks at it, exercises discernment, and graduates it to approved before it reaches a client, a deliverable, or a pitch.

Dimension	Impact
Time saved	None directly. This move costs time.
Accuracy and quality	The control that makes every later move safe to attempt.
Cost and revenue	None directly.
Growth	None directly.
Employee experience	**The highest-value move on this list.** It is the promise to respect the craft, made structural instead of verbal.
Risk	Removes the largest one: ungoverned AI output reaching a client under a practitioner's name.

Sized last on four dimensions and first overall. Nothing else on this playbook survives without it.

Move one: writing tooling, brought in

Voice-bound content tooling that already exists and has been proven elsewhere. Not built here, not novel, and not the most impressive thing available.

That is the point. The first thing a skeptical practitioner touches should be something that already works, so the first impression of AI in this building is competence rather than a demo.

Dimension	Impact
Time saved	Real but not the reason to run it first.
Accuracy and quality	Consistency across clients, which is one of the three stated asks.
Cost and revenue	Modest.
Growth	Indirect.
Employee experience	High. This is the move that makes the case in a practitioner's own hands.
Risk	Low, which is why it goes first.

Move two: new business intake

An agent that carries the research burden on inbound leads before a human spends time on them.

Chosen second because it is adjacent to the craft rather than inside it. Nobody's professional identity is built on researching an inbound lead.

Dimension	Impact
Time saved	Direct, on every lead.
Accuracy and quality	Consistent intake rather than whatever the person had time for.
Cost and revenue	Faster qualification on new business.
Growth	Direct. This is the commercial front door.
Employee experience	Positive and uncontested.
Risk	Low. Internal-facing, nothing reaches a client.

Move three: research

The answer to the gathering in step 1 of the traced pitch. The day's news read across the sources and delivered as a report, so the publicist arrives at the choosing instead of starting one to two hours behind it. The angle is still chosen by a person.

This is the move that pays for the program, and it is deliberately third. It touches the practitioner's own daily work, which is exactly why it should not be first.

Dimension	Impact
Time saved	**The largest on this list.** Recovers the homework hours at the front of every pitch cycle.
Accuracy and quality	Wider coverage than a person hunting site by site under time pressure.
Cost and revenue	Practitioner hours redirected to the work clients pay for.
Growth	More pitches reached per publicist per day.
Employee experience	High, and this is the move that converts opposition. It removes the part of the job nobody defends.
Risk	Moderate. Output feeds client-facing work, so it runs through the boundary from move zero.

Move four: company memory

A knowledge graph, so what the firm knows is available across deliverables and pitches rather than living with whoever happened to work the account.

Third of the three stated asks, and last because it is worth most once the other moves are generating something worth remembering.

Dimension	Impact
Time saved	Compounding rather than immediate.
Accuracy and quality	Consistency across clients, structurally.
Cost and revenue	Indirect.
Growth	Strong on pitches, where prior work is the asset.
Employee experience	Removes the tax of reconstructing what the firm already knew.
Risk	Low.

Move five: hand over the building

Teach the firm's core enablement group to improve what exists and to build what comes next.

Dimension	Impact
Time saved	Not the point.
Accuracy and quality	The group closest to the work makes better decisions about what to automate.
Cost and revenue	Ends the dependency, and the spend that comes with it.
Growth	The firm's AI capacity stops being bounded by ours.
Employee experience	The strongest possible answer to the original opposition: the people who were against it now build it.
Risk	Requires the governance from move zero to be genuinely held, not just documented.

What it takes to run this

Discipline. The boundary only works if it is never quietly skipped under deadline. A public relations firm runs on deadlines.

Accountability. Someone owns the question of what AI is allowed to do. Before this engagement, nobody did.

A team. The enablement group is the accountable unit, which is why move five is on the playbook rather than assumed.

Where AI does not go

The pitch itself. The judgment about which reporter to approach. The client relationship.

Not because a machine could not attempt them, but because they are what the firm sells.

Charter

What a Charter is

One document at three points in time: the specification, the delivery record, and the current status. Its spine is the decision log.

Metadata

Field	Value
Client	A national public relations agency
Industry label	Communications and public relations
Entry	No Day One. A week of review, with multi-day practitioner interviews inside it. Terms agreed before the work began.
Sponsor	The CEO, with the managing partners
Stated asks	Increase practitioner output. Build consistency across all clients. Build company memory usable across every deliverable and pitch.
Delivery status	Delivered. Continuing as monthly advisory.

Positions

Position	Held by	What it carried
Sponsor	CEO	The commercial outcome and the mandate to change how the firm works.
Owner of the craft	Managing partners and practitioners	The standard the output is held to, and the reason automation stops where it stops.
Governance	Unowned at the start. Assigned during the engagement.	What AI is allowed to do, and what must be true before its output is used.
Build and enablement	First Strategy, transferring to the firm's core enablement group	Making the systems, then making the people who make the systems.

The unowned governance Position is the finding. Every other problem in the audit descends from it.

Objectives and constraints

Scope

In. Governance for AI output. Writing tooling. New business intake. Agentic news research. Company memory. Training for practitioners and for the enablement group.

Out. The pitch itself. The judgment about which reporter to approach. The client relationship. Autonomous anything.

Objective and success criteria

Move the practitioner's day so it begins at the craft rather than reaching it after the homework, without any AI output reaching a client unexamined, and without losing a workforce that started out opposed.

Success is not a spend figure or a tool count. It is whether the people who were against this are using it, and whether the firm can build the next thing without us.

Constraints

Constraint	Source
No AI output reaches a client, deliverable, or pitch without human discernment	The audit's hardest finding, and non-negotiable
The craft is not automated	Leadership's standard and the practitioners' condition
Trust is earned in sequence	The workforce was opposed at the start
Governance precedes capability	Agents were already running ungoverned on critical systems

---

Architecture and human-in-the-loop design

AI writes into a separated directory. Nothing in it is usable. A person reviews, exercises discernment, and graduates the output to approved. Only approved work moves on.

The machine can work. The machine cannot ship.

This is the whole architecture, and everything built afterward sits behind it. It is also not specific to this engagement: it is the firm's standing operating governance, applied in every client, which is why it could be the first thing in place here rather than something to design.

Current state at the start

Overspending on AI in one part of the firm, no adoption in the rest. Agents standing on critical systems with no design behind them. No governance. Practitioners opposed. Media relations homework consuming one to two hours before the craft work began.

Decision log

#	Decision	Alternatives rejected	Evidence
D1	Build the governance boundary before building any agent	Build first and write the rules after, which is the normal order	The workforce was already opposed and ungoverned agents were already running. A second failed attempt would have been much harder to recover than the first.
D2	Lead with writing tooling we already had rather than something new	Open with the most impressive capability	The first thing a skeptical practitioner touches should already work. Proven tooling makes the case in their own hands; a demo does not.
D3	Sequence by trust, not by value	Run the highest-value move first, which would have been research	Research touches the practitioner's own daily work. Going there first, into an opposed workforce, risks the program to save a few weeks.
D4	Automate steps 1 and 2 of the pitch, never step 3	Assist or generate the pitch itself	The pitch is the product. Automating it confirms the fear and removes what the client is buying.
D5	Teach the enablement group to build, rather than keep building	Continue as the firm's build shop	Capability over dependency. The group closest to the work makes better decisions about what to automate, and the firm's AI capacity stops being bounded by ours.

The decision and experiment record

Witness

A week of review, most of it interviewing practitioners directly. The finding was not a hidden system but a scale: the homework at the front of every pitch cycle, and an AI condition with no owner.

Plays: Field Observation, User Flow Mapping.

Interrogate

The assumption that mattered was the firm's own: that the AI problem and the adoption problem were two problems. The week showed they were one.

Solve

The boundary first, then writing tooling, then new business intake, then research, then company memory. Each move gated on the previous one landing with practitioners.

Plays: Human-in-the-Loop Design.

Expand

Adoption spread across the practice. The measure was never a seat count; it was whether the people who had been against it were using it. They were.

Refine

Oversight designed to loosen as reliability was demonstrated, through the graduation step rather than by decree.

Plays: Hierarchy of Agency Design, Graduation Decision Making.

Current status and the autonomy transfer

Everything scoped was delivered.

The core enablement group was trained, then taught to improve what exists and to build what comes next. They build for themselves now. The capacity to make the next agent sits inside the agency.

What continues is a standing monthly conversation with that group: questions, decisions about what to build next, and oversight of the governance.

The engagement did not end. It became advisory.

Outcomes

Outcome	Status
The gathering runs as an agent; publicists reach the choosing sooner and reach more pitches. The judgment about what is newsworthy stays human	Delivered. Directional by choice.
Adoption increased dramatically	Delivered. Stated without a figure; see `../Facts to Confirm.md`.
The opposed workforce is the workforce using it	Delivered. The engagement's real result.
The enablement group builds independently	Delivered.
Governance exists and is owned	Delivered.

Plays

Canon	Play
Witness	Field Observation
Witness	User Flow Mapping
Solve	Human-in-the-Loop Design
Refine	Hierarchy of Agency Design
Refine	Graduation Decision Making

No Interrogate play is indexed. This engagement ran no formal experiments, and the canon is not padded to look complete.

The plays

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

WITNESS PLAY

Field Observation

Witness play, instantiated for the communications and public relations engagement. Purpose: watch the work as it happens, not as documented.

How it ran

A week of review, the substantial part of it interviewing practitioners directly across several days.

We started from the work rather than the tooling. The question that mattered was not what the firm had bought but what its people did all day.

What we looked at	What we were listening for
The daily pitch cycle	Where the hours actually went before the craft work started
The AI already in use	Who was using it, on what, and under what rules
The AI not in use	Why not, in the practitioners' own words
Agents already built	What they were pointed at, and what design sat behind them

What it found

Not a hidden system. A scale.

Leadership already knew media relations carried heavy homework. The week established how much, and it established something they had not connected: the firm's overspending problem and its non-adoption problem were the same problem wearing two faces.

It also found the sentiment plainly. The practitioners were not indifferent to AI. They were against it, and they had reasons that were about craft rather than about technology.

Why it mattered

Every later decision came from here. Sequencing by trust instead of by value, leading with proven tooling rather than a new build, and putting the governance boundary before the first agent are all consequences of what a week of listening surfaced.

WITNESS PLAY

User Flow Mapping

Witness play, instantiated for the communications and public relations engagement. Purpose: trace one unit of work all the way through, every handoff included.

The unit

One pitch. The most representative thing a publicist does.

News research
manual, site by site, one to two hours a day

→

Reporter research
who covers this, who might take it

→

Craft the pitch
the work clients pay for

↻ Steps one and two run first, every day, by hand, before the craft can start.

Step	What happens	Judgment required	Repeats
1	News research, manually, site by site, hunting the day's angle	Split. Gathering is low; choosing the angle is not	Daily, one to two hours per publicist
2	Reporter research: coverage history, adjacency, likelihood	Mixed. Assembly is retrieval; selection is judgment	Per pitch
3	Craft the pitch	High. This is the product	Per pitch

What the trace decided

The automation boundary sits inside step 2, not at its edges.

Step 1 splits. The gathering is high-volume, low-judgment and repeats daily, so it goes. Choosing which angle is worth a pitch stays with the publicist. Step 2 splits: assembling candidate reporters is retrieval, deciding who to approach is not. Step 3 is untouchable, because it is what the client is buying.

Mapping the flow is what turned "increase practitioner output" from a goal into an instruction. The output was going into steps 1 and 2. That is where it was recoverable.

SOLVE PLAY

Human-in-the-Loop Design

Solve play, instantiated for the communications and public relations engagement. Purpose: define who reviews AI output, how, and what gets logged, so a human directs the AI rather than the reverse.

The design

AI output is written to a separated directory that cannot be used directly.

AI produces
into a separated directory

→

Not usable
no client, deliverable, or pitch

→

Human discernment
review, judge, graduate

→

Approved
now it can be used

↺ The machine can work. The machine cannot ship.

Element	Design
Who reviews	A person, before any AI output is used
What they check	Whether it meets the standard the firm's name goes out under
The gate	Graduation to approved. Nothing ungraduated reaches a client, a deliverable, or a pitch
What it settles	AI is a collaborator, not an autonomous doer
Accountability	The human who graduates it owns it, exactly as they always did

---

Not bespoke

This is First Strategy's standing operating governance, applied in every client engagement rather than designed for this one. It arrived here already proven, which is part of why it could be the first thing built.

Why it came first

This was built before any agent, inverting the usual order.

The firm's practitioners were opposed to AI on craft grounds, and agents were already running ungoverned on critical systems. In that condition the boundary is not a control added after the fact. It is the precondition for being allowed to build anything at all.

It also did a second job. The firm was told we would respect the craft and the people doing the work. The directory is what that sentence means in practice, which is why the promise was believed.

REFINE PLAY

Graduation Decision Making

Refine play, instantiated for the communications and public relations engagement. Purpose: set the evidence thresholds for moving to less oversight.

What graduates, and what does not

Graduation in this engagement operates on two levels, and they are easy to confuse.

Output graduates. Every piece of AI work moves from the separated directory to approved through an act of human discernment. This is per-piece, it happens every time, and it does not get easier with reliability. It is the mechanism in `Human-in-the-Loop Design.md`.

Capability graduates. A class of work moves down the oversight tiers as it earns it. This is the one that changes over time.

Level	What graduation means	Does it loosen?
A piece of output	A person judged it good enough to use	No. Every piece, every time.
A class of work	Demonstrated reliability moves it to a lower tier	Yes, on evidence
Client-facing work	Nothing. It does not graduate.	Never

The thresholds

Capability moves down a tier when the people doing the work stop finding problems in it, and they are the ones who decide, not us and not the tooling.

That is a deliberately human threshold rather than a metric. In a firm where the workforce began opposed, an error rate that satisfies a dashboard is worth less than a publicist who has stopped checking because they no longer need to.

Why it was designed this way

The firm had agents running on critical systems with no thresholds at all, and practitioners who assumed AI would be forced on them regardless of whether it worked.

Making graduation explicit answered both. It gave the ungoverned agents a standard to meet, and it gave the practitioners something they had not had: a say in how far this goes, held by evidence they generate themselves.

Autonomy grew here because reliability was proven to the people doing the work. It was never granted by decree.

REFINE PLAY

Hierarchy of Agency

Refine play, instantiated for the communications and public relations engagement. Purpose: define tiers of human oversight by the risk the job carries.

The tiers

Oversight in this firm is set by one question: whose name goes on the output, and who is looking at it when it does.

Client-facing*pitches, deliverables. Human discernment on every piece, always* Internal-facing*research reports, intake summaries. Reviewed in use* Retrieval*gathering, assembling candidates. Checked by the work it feeds*

Tier	What sits here	Oversight
Client-facing	Anything going out under a practitioner's name	Discernment on every piece. This tier never loosens.
Internal-facing	The day's news report, new business intake summaries	Reviewed by the person using it, whose own work depends on it being right
Retrieval	Gathering news, assembling candidate reporters	Validated by the step it feeds rather than reviewed on its own

The tier that does not move

Client-facing work stays at full oversight permanently.

This is not a starting position to be relaxed once reliability is proven. In a craft business the output carries a person's professional judgment, and the moment a machine's unexamined words go out under their name the firm has sold something it does not have.

Autonomy is available in the lower tiers. It is not available in the one that matters most, and saying so early is part of why the practitioners came around.