

What's Changing in What Your Research Delivers

Specialist Research Intelligence

You have seen this on your research bench

“Your clients’ procurement teams are no longer asking what your research costs to produce. They are asking what your interpretation is worth. By the time your project manager is in the room, the production line has already been benchmarked.”

01

Why is realisation drifting on research that our rate card says should hold?

02

What happens to our quality when our senior analysts are not in the room?

03

Why are clients resetting their turnaround expectations against AI?



You have seen this on your research bench

Every Head of Research Delivery or Chief Operating Officer at a mid-sized specialist intelligence firm in Europe has had the same week. Three client situations on the same calendar. First: a competitive intelligence subscription renewal. The client's procurement lead opens the conversation by saying their internal team is already generating a first pass with AI tools, and they want to understand what they are still paying for. Second: a due-diligence project on a potential acquisition target. Your junior analyst team delivered the information pack faster than usual, but your senior review caught three interpretation errors that would have changed the recommendation. Third: your 12-year intelligence retainer. The senior analyst who has held that relationship since the firm was a third of its current size schedules thirty minutes to talk about their plans for next year. The subject line says "timing."

You have run this delivery engine for over a decade. The model was clear: hire the best analysts you could find, build their sectors, protect your methodology, hold your turnaround commitments. The business was understandable.

Something shifted. Your junior analyst utilisation is under pressure because clients are questioning the production-research line items they used to buy without hesitation. Your senior intelligence professionals are the scarcest resource in the firm, and they are spending significant time assembling information that AI tools now handle in hours. Your clients' expectations around turnaround have reset against what they think AI should deliver. And you can name the three or four senior analysts whose departure inside the next two years would put your most important client relationships at risk.

These are not separate problems. They trace back to one unsolved thing, and the firms that solve it first are already pulling ahead in pricing conversations and client retention. The ones defending the old production model are watching margin drift every quarter.

This briefing covers what is happening in your sector, what it means inside your delivery engine, and the three questions your operations team should answer before the next board cycle.

The production model that held for thirty years is under direct substitution

For three decades the commercial logic of specialist research was stable. Gathering and synthesising high-quality information about a sector, a competitor, or a market was expensive. It required skilled analysts, proprietary sources, rigorous methodology, and time. Clients paid premium rates because the alternative was doing it themselves, and doing it themselves was more expensive than the fee. Both the value and the cost were real.

AI has changed one side of that equation. The information gathering, the synthesis, the structured analysis, the benchmarking, the report compilation: all of it that used to require analyst hours now runs at near-zero marginal cost. Clients with access to AI tools are already producing the first pass of analyses they used to outsource. AI-native research competitors have entered the market with a cost structure that no production-based research firm can match on price. And enterprise clients with internal AI capability are discovering that 60 to 70 percent of what they used to buy, they can now generate themselves.

Your clients' procurement teams are no longer asking what your research costs to produce. They are asking what your interpretation is worth. By the time your project manager is in the room, the production line has already been benchmarked.

The 30 percent that survives is genuinely irreplaceable: the insight, the interpretation, the "what does this mean and what do we do about it." It is the accumulated sector pattern recognition your senior analysts carry from 15 to 25 years of client work, not from a training dataset. The firms that migrate toward that 30 percent as their explicit commercial proposition are finding that advisory retainers carry margins of 35 to 55 percent against the 20 to 30 percent that production subscriptions deliver. The ones that defend the production model are repricing downward every quarter until the model stops working.

Your research bench, your quality, and your client relationships are feeling it now

Your research delivery engine. Project managers are working harder to hold turnaround commitments that clients now benchmark against AI alternatives. Scoping conversations are more contested than they were two years ago. The delivery methodology your firm refined over a decade was built around analyst hours as the primary input. The methodology has not been redesigned around a world where information assembly is no longer the bottleneck.

Your intelligence quality. The quality floor holds because your senior analysts are still in every review. The quality ceiling has not risen in years. Both are at risk. Your best intelligence professionals catch the interpretation error a junior would miss, read the signal the data is not showing, understand what a client's regulator actually cares about beyond what is written. That judgment was accumulated over twenty years of sector work and has never been written down. When one of them leaves, two client engagements run with a visible quality gap before the client notices.

Your analyst bench. Your junior and mid-tier analysts bill at rates justified by the depth of work the client expects. The information assembly layer that supported those rates is under direct substitution pressure. Realisation on production-heavy projects is drifting downward. Most of your junior analysts have not had a conversation yet about what the next version of their role looks like, and the best ones are starting to ask.

Your senior intelligence professionals. Your ten to twenty most experienced analysts carry the interpretive depth clients actually pay for. On a standard week, a meaningful portion of their time goes into assembling and checking information rather than doing the interpretation only they can provide. The scarcest resource in the firm is not fully deployed on the work that makes it scarce.

Your proprietary knowledge base. Years of client work have produced something most firms have not yet thought of as an asset: sector-specific pattern libraries, accumulated methodology, historical engagement archives, client-specific intelligence built project by project. Right now it sits unsystematised in project files and in the memory of senior analysts who may leave inside the next two years.

The firms pulling ahead have stopped treating information and insight as the same activity

Your sector is being compressed from below because AI commoditises information production. Your delivery engine is compressed from above because your senior intelligence professionals are spending time on information assembly instead of on the interpretation your clients cannot get anywhere else. Same force, two altitudes. But there is a third dynamic that most research COOs have not yet named, and it is the one that determines the next decade.

Your senior analysts have never been separated from the sector knowledge they carry. Every engagement makes the same few people more valuable. Not the firm. Every project file, every client briefing, every annotated research memo adds to individual expertise that walks out when the person does. Every margin pressure, every quality risk, every contested renewal is a symptom of that one unsolved thing. The firms that extract that accumulated knowledge as a side-effect of daily client work (from project artefacts, annotated briefs, client-specific intelligence) turn individual expertise into institutional infrastructure. The ones that do not are one retirement conversation away from their most defensible competitive asset walking out.

The firms pulling ahead have understood something that sounds counterintuitive: the substitution of production research is not the threat. It is the reveal. These firms were always selling sector judgment, not information assembly. AI has made that visible. And now that it is visible, the ones moving fast are charging explicitly for the interpretation that was always the real product, and building the infrastructure to hold it even when people leave.

You can tell which camp a firm is in from a single question. If your three most experienced analysts left tomorrow, how much of their sector-specific pattern recognition would remain accessible to the firm? Most intelligence firms cannot answer that. The ones that can have already made the shift.

You have 12 to 18 months. The commercial model migration must happen within that window. Each quarter of delay is a quarter where client expectations reset further against AI alternatives, and a quarter where an AI-native entrant gets closer to replicating a sector dataset it once took your firm fifteen years to accumulate.

What most firms are doing. What the firms pulling ahead are doing instead.

Every research COO has had the AI conversation. Most are making the same four decisions. A smaller group is making different ones. The difference is architecture, not tools.

THE DEFAULT PATH	THE WINNERS' PATH
Defend the production model on lower price Same deliverables, harder on cost. Margin drifts until it does not cover delivery.	Price the interpretation, not the information Migrate toward advisory retainers at 35–55% margin. Production becomes input, not product.
Treat senior departures as an HR succession problem Recruit a replacement, accept 18 months of quality gap. The client relationship goes to tender.	Extract senior judgment as a side-effect of daily work Capture from project artefacts they already produce. Knowledge stays when the person leaves.
Add AI tools on top of the existing delivery model Each analyst gets a productivity tool. Turnaround improves slightly. The model stays unchanged.	Rebuild delivery around interpretation from senior analysts Juniors handle AI-assisted production. Seniors own interpretation and client advisory exclusively.
Leave the proprietary data unsystematised Project files and analyst memories. An AI-native entrant closes the gap each quarter.	Systematise the accumulated sector knowledge now Turn years of client work into a queryable asset. The gap that took 15 years to build takes months to deploy.

Same research firm. Same twelve months. Opposite outcomes. The difference is what the delivery model is built around.

THREE QUESTIONS

Analyst Productivity. Intelligence Quality. Delivery Speed.

Three questions every research COO is tracking. None of them used to be the same question. They are now.

01 · ANALYST PRODUCTIVITY

Why is realisation drifting on research that our rate card says should hold?

The rate card has not moved. Analysts are busy. But realisation on production-heavy projects is down, and clients are explicitly comparing your cost against what their internal teams now generate as a first pass. Junior analyst hours are under direct substitution pressure. Mid-tier are next. The analysts who migrate toward interpretation work are the ones whose billing the firm can defend. The ones still anchored in information assembly are becoming the question every renewal surfaces.

02 · INTELLIGENCE QUALITY

What happens to our quality when our senior analysts are not in the room?

Four to eight of your most experienced intelligence professionals are within two to three years of a departure decision. Each carries sector-specific pattern recognition that took two decades to build and has never been written down. Your quality floor holds because they review every major deliverable. When one leaves, two engagements run with a gap before the client notices. Three before the client mentions it. The succession conversation is on your calendar. The knowledge-extraction work is not.

03 · DELIVERY SPEED

Why are clients resetting their turnaround expectations against AI?

A standard competitive intelligence project runs eight to twelve weeks on a mid-complexity brief. Roughly half that elapsed time is information assembly, data aggregation, and synthesis that now compresses to days not weeks with AI-assisted tools. Firms that have rebuilt delivery around this are closing comparable projects in four to six weeks, at the same quality standard. The same rate card. Faster delivery. 40 percent more client projects per senior analyst. The clients whose expectations have already reset are the ones asking why yours have not.

What the firms winning this play are actually doing

Four concrete moves any research COO can start this quarter. Each touches one part of the delivery engine. Together, they become one architecture.

1 **Delivery: decompose every project type to the task level, rebuild pricing around the tasks that command a premium**

List the task bundles inside each engagement type. Mark each: AI-assisted with analyst review, analyst-led with AI support, senior-only interpretation, retired. What remains is a different engagement at a different cost base. The COO who walks into the next pricing committee with the task-level map owns the conversation. The ones defending engagement-level rate cards are defending a 2019 model.

2 **Quality: extract senior judgment now, as a side-effect of daily client work, not as a documentation project**

Knowledge-capture programmes have failed in professional services for thirty years. The judgment is not articulable in a memo. Extract from artefacts your senior analysts already produce: annotated briefs, client advisory notes, redlined drafts, post-engagement reviews. Senior consent is straightforward: the role shifts from delivery to interpretation. The sector knowledge stops being biographical. When the analyst leaves, the intelligence stays.

3 **Bench: redesign the junior analyst pathway around interpretation from day one, not information assembly**

The information assembly work that trained juniors into mid-tier analysts is moving to AI tools. No production work, no training ground, no senior bench in 2030. New analysts start on quality review, interpretation audit, and paired work with senior intelligence professionals. The five-year competence curve compresses to eighteen months. The junior who joins today builds interpretive judgment from month one, not year three.

4 **Knowledge base: systematise the proprietary sector intelligence your firm has accumulated**

Years of client work have produced something no AI-native entrant can replicate quickly: sector-specific pattern recognition built through real engagements. Right now it sits in project files and analyst memory. Systematised and queryable, it becomes the foundation for intelligence products no competitor can offer and the quality infrastructure your senior review runs on.

Where research-delivery advisory meets working knowledge infrastructure

Four places to go for help. None can do what we do, and each for a specific reason.

The major consulting firms' strategy and operations practices will sell a diagnostic at €500K to €1.5M delivered by a partner and analysts billing hours. They produce strategy decks. Their own commercial model is structured around the billable-hour logic your industry is trying to leave.

The knowledge-management and enterprise search platforms treat your senior analysts' judgment as a filing problem. Every research firm has watched that purchase fail once already. The knowledge that makes the firm valuable lives in how senior analysts interpret documents, not in the documents themselves.

The AI research and intelligence platform vendors are the layer already commoditising your production revenue. Adding more speeds up the substitution of the work you need to move away from, not toward.

Your internal IT function and digital transformation teams do not sit close enough to how your senior intelligence professionals reason to codify their judgment. The work that needs capturing is the work IT has never watched happen.

The market offers four things you have already considered. GRAIL is the fifth option.

WHAT YOU GET

Advisory that speaks research-delivery, not strategy consulting

Working infrastructure your analysts use from week one

Custom intelligence agents built on your methodology and sector patterns

A knowledge base your senior analysts actively feed.
No vendor dependency

WHY IT MATTERS

Founded by practitioners who ran professional-services delivery

Business model first, technology second

Three weeks from first session to working infrastructure

Your analysts keep everything. You own the IP.
Every next step earned

Three sessions with your research leadership. Working infrastructure from week one.

We start with your research leadership team because the migration cannot outrun the team's personal understanding of what is possible. Three workshops. Each delivers standalone value. Each earns the right to the next.

WORKSHOP 01

How to think

Walk through what is happening in your sector and inside your delivery engine. Build the first piece of infrastructure in the room: a client-specific intelligence agent trained on one of your real sector engagements. Your research leadership team leaves with a shared view of what to migrate toward and the first piece of working infrastructure already in their hands.

WORKSHOP 02

Your intelligence infrastructure

A half-day build. Your Head of Research, Head of Quality, Senior Intelligence Leads, and Head of Analyst Development each leave with five to ten pieces of working infrastructure tuned to their function. Research production agents. Quality-review workflows. Client-specific knowledge capture. Analyst pairing frameworks. Infrastructure your teams are running the same week, not slides about infrastructure.

WORKSHOP 03

Ready to run

Refined infrastructure based on real use. The rollout plan to your wider analyst team. The knowledge-extraction plan for your senior intelligence professionals. The commercial model migration path for your CEO. And the CFO conversation with numbers they can check: revenue per analyst, advisory margin by client segment, knowledge-continuity exposure by senior.

Three weeks. No IT project. No board approval needed to start. The infrastructure your research delivery engine is already running on when we finish.

For your next operations review

The questions where your research leadership cannot agree on an answer are the ones worth an hour on the agenda.

Of your last ten client renewals where pricing was harder, how many cited internal AI capability or AI-native alternatives as the comparison? If you do not know the number, what does that tell you?

If your three most experienced intelligence professionals left in the next twelve months, how much of their sector-specific pattern recognition would remain accessible to the firm? Can you name the client relationships at risk within eighteen months of each departure?

On your last five major projects, what share of elapsed time went into information assembly that AI tools could now handle in days? What would those projects look like if that time went entirely into senior interpretation?

What does junior analyst onboarding look like today versus two years ago? Is the first year still built around information production? If so, what are you training for the bench you will need in 2029?

Can you point to one thing in your firm that is measurably smarter this quarter than last? Not faster, not cheaper. Actually more intelligent from accumulated experience. If not, that is the conversation worth having first.

The hours your senior intelligence professionals spend assembling information are the most expensive hours in the firm. They are also the only ones that disappear when you build the infrastructure to hold their judgment. The question worth asking is not what AI will do to your cost base. The question worth asking is what AI will do to the capacity of the only people your clients cannot replace.

IN THIS SERIES

There's a version of this briefing for every seat at the table.

GRAIL has produced industry briefings calibrated to the specific reality of each C-suite role.

CEO

Revenue. Margin. Valuation.

Three numbers your board reads. AI is rewriting all three in your industry.

CFO

Capital. Forecast. Story.

What you allocate this year is the story your board reads next year.

COO

» YOU ARE READING

Cheaper. Better. Faster.

The triple threat that used to be a tradeoff is now table stakes.

CHRO

Capacity. Capability. Promise.

The biggest talent transition of a career. Compete on capacity, not headcount.

CPTO

Moat. Speed. Allocation.

The platform moat that survives 2028 is being chosen this year.

CCO

Margin. Pace. Position.

The whole GTM playbook is being rewritten in real time, across every commercial team at once.

Visit grail.works for the full series.

GRAIL

An AI-native management consultancy helping European mid-market leadership teams figure out what fundamentally changes about their business as AI enters their industry.

www.grail.works