insightology.ai

A Guide to AI for Strategic Market Research Agencies

What is actually changing, what it means for your agency, and what to do about it

A WHITE PAPER BY INSIGHTOLOGY
OCTOBER 2026

Executive Summary

01

One thing is by now crystal clear: AI will have a profound impact on strategic market research agencies driving significant change at an accelerated pace. This white paper explores what these changes are and outlines some thoughts on how agencies can equip themselves to benefit from all the opportunities and protect themselves from the threats. Look at a typical strategic research project: anywhere from 50-70% of billable time goes into pure production (our working estimate across the agencies we work with) — building methodologies, writing proposals (that you don't always win), designing complex surveys and discussion guides, fieldwork management, data analysis and reporting. Broad AI tools like Claude and ChatGPT help, but they often fall short of professional standards, creating a frustrating "polishing tax". That is rapidly changing. Next-generation AI systems, built specifically for research, can now deliver elite-grade outputs in a fraction of the time.

The shift is here. Every agency leader we speak to already knows it. What follows is our view of how the next few years play out, and what to do about it.

The whole argument, briefly

If you are a strategic market research agency, the industry is restructuring in your favour. Seven shifts are underway at once, and six favour well-positioned agencies: production costs fall and reset every cost base; agencies get leaner and more senior, with new choices about the work they take; data becomes a commodity, so the value of true insight climbs; cheaper research means more of it gets bought; tactical and strategic work pull apart, squeezing the undifferentiated middle; and while the agency population consolidates, the survivors inherit the work. Section 2 takes each in turn.

The pressure lands on how you make money. Production of documents and reports is most of what your invoices bill for, and that is disappearing. The junior training route that ran on that 'grunt work' is closing with it. And once anyone can generate a plausible-looking deck, plausible stops being worth anything at all. Section 3 sets out the four uncomfortable truths.

The answer is to sell what we call 'warranted thinking'. Agencies sell their thinking but price like contractors — time plus third-party costs plus a markup — and that model breaks when AI halves the production days, because clients can see the hours aren't there. Charge instead for the thinking and the deliverables that carry it: the rigour of the design, the debrief, the strategic recommendations. Some of the saving should flow back to clients; all of it must not. What makes the new price stick is the ability to warrant it — the warrant is the method, the instruments, the real respondents, the data handling and the trail from brief to conclusion; the thinking is a set of skills only people have. That is a value and pricing decision, not a technology one, and it is the biggest one most agencies will make this decade.

Divide the labour deliberately. AI's powers are powers of execution — speed, coverage, multi-modal sight, consistency, patience. The eight human powers are what clients actually buy: framing, creativity, strategic thinking, contextual intelligence, insightfulness, judgement someone stands behind, relationships and collaboration. Agencies go wrong in both directions — people doing factory work, or generic AI making judgement calls — and Section 4 maps the split that works.

For strategic work, specialist beats generic. A generic assistant is a chat window; a specialist platform has the research method built into the software, with memory, enforced process and an audit trail. Law has already run this experiment: its top firms chose specialist, and held to it when the model makers arrived. Section 5 shows why.

Choose your AI path with your eyes open. Use a generic assistant well, build a skills layer on top of one, or use a specialist platform — built or rented. Each genuinely wins for a different kind of agency, including cases where the cheap option is the right one. Sections 6 and 7 give you the decision framework and the honest economics.

Then prepare deliberately. Audit where your hours go, pilot on live projects, reprice one service line, and rebuild how you train juniors — the 90-day plan in Section 8 turns all of the above into a quarter of manageable work.

Handled well, this transition leaves an agency with better margins, more projects and a stronger seat at the table on its clients' biggest decisions. Handled with a few ChatGPT licences and a hopeful memo, it leaves the agency in the middle of the market — and Section 2 explains why the middle is where the squeeze is.

What is actually changing in the industry 02

Most conversations about AI and market research get summarised in one line: there will be fewer agencies. That may be true, but it's unlikely to be the market change that we notice most, because seven things are changing at once, and six of them are good news for agencies that are well positioned.

01 Production costs will fall hard — and your cost base resets

On a typical strategic research project, anywhere from 50–70% of billable agency labour goes into pure production work (our working estimate across the agencies we work with) — by which we mean developing methodologies and proposals (that you don't always win), designing complex survey tools and discussion guides, fieldwork management, data analysis and reporting.

Specialist AI tools designed for market research can cut the labour cost of this production work by up to half. The specific changes will be different for each agency, but the driver will be the same for them all: AI will permanently reset the cost base of a research project, and your staffing plan, pricing and utilisation numbers all reset with it.

02 Your agency will change shape — leaner, more senior, with new choices

Falling production costs don't automatically mean fewer agencies — that may be the outcome of this or other forces — but they do mean agencies with a different shape: leaner, with fewer junior roles and a more senior mix. But look, AI also gives agencies higher capacity, and a choice about how to use it: do more with the team they have, or be more selective about the work they take on, specialising far more than they could before.

It also creates the real opportunity to move up the thinking chain. Removing the grunt work frees up time for thinking and insight delivery — and the next shift explains why clients will need much more of exactly that.

03 The value of insight will go up

One more change runs in strategic agencies' favour, and it comes from the sheer abundance of data. The world will create around 230 zettabytes of it this year, a figure forecast to roughly triple between 2025 and 2029 (IDC/Statista) — and most of it sits untapped, unverified or, increasingly, synthetic. Gartner puts unstructured data at 80–90% of all new enterprise data, growing three times faster than the structured kind.

Abundance does to data what it does to everything else: it turns it into a commodity. And as the data itself becomes commonplace, the ability to find the true insight hidden inside it becomes dramatically more valuable — rarer, harder, and worth more precisely because everyone is now drowning in the raw material. The winners are no longer the organisations that hold the data; they're the ones who can make sense of it. Which puts insight and strategic thinking at the centre of where value gets created next. Knowledge is not power any more. Only insight is.

04 Falling costs + rising value of insight = more demand

This is the part that people tend to miss: falling production costs mean the price of a typical project will fall. When the price of good research falls, more organisations can buy it. The company that could never sign off an expensive strategic study now can. The client that ran four strategic projects a year now runs eight. And because everyone is sitting on more data than they can make sense of, the value of genuine thinking and insight goes up, not down. The evidence supports this: the insights industry reached $153bn in 2024 and keeps growing in cash terms even as unit costs fall (Esomar Global Market Research 2025).

05 Tactical and strategic work will polarise

Synthetic respondents, AI moderation and DIY platforms are turning the tactical end of research into a commodity: quick-turn testing, trackers, low-stakes studies. Research software is the fastest-growing part of a $153bn industry: $62bn, up 11.5%, against 4.8% for established market research (Esomar Global Market Research 2025).

But the same forces push the value of strategic work up. Once anyone can gather data and produce something that looks like professional research, being an established agency with a validated method, real human respondents and genuinely good thinking stops looking expensive and starts being the reason clients buy from you. Depth is where the defensible value now sits. The labour market tells the same story: roles requiring AI skills are growing eight times faster than the jobs market and carry a 62% wage premium (PwC, 2026). Both ends of the market are doing well; the squeeze lands on the undifferentiated middle.

06 Consolidation: fewer agencies, but stronger ones

Here is the headwind. The global agency population will probably shrink. But the agencies that go will mostly be the agencies that resist AI altogether, or agencies competing at the commodity end where deep insight and strategic thinking are less valuable and prices are collapsing fastest. AI-enabled agencies doing strategic work don't just survive this, they pick up the work the departing agencies leave behind.

Consolidation has a second edge for independents to watch: the big network groups will chase this wave with scale economics — highly advanced and specialised AI, offshored production, aggressive cost-cutting — and they can undercut on price for as long as their shareholders allow. Independents shouldn't try to win that race on size. The counter is the one this guide describes: specialist AI to level the cost base, plus the senior attention and strategic focus a network finds hard to industrialise.

07 Clients take the simple work in-house

Clients have been using DIY tools and in-house teams for simple research for over 20 years. AI speeds that up sharply. Concept screens, basic trackers and quick polls will keep moving client-side. What stays with agencies, and gets more important, is the complicated, high-stakes work involving multiple stakeholders, where judgement, method and independence are what the client is paying for. If your revenue is concentrated in work clients can now do themselves, know the proportion and plan without it.

What it adds up to

Two patterns run through all seven changes. First, the market is splitting: tactical work becomes a commodity while strategic work becomes more important and therefore more valuable. Second, total demand is growing, because cheaper research plus a greater need for insight and thinking means more research gets bought.

If you run a strategic agency, both patterns work in your favour. Your demand outlook is the best in the industry. The next section explains why that still doesn't make you safe.

What this means for strategic agencies

03

The market news is good. The problem is that AI hits strategic agencies somewhere less obvious than demand. It hits how you make money, how you train people, and how you prove your work is any good. Here are four things worth being honest about.

Your document production work is disappearing, and it is on your invoices

50-70% of the billable hours on a strategic project are production tasks: methodology design, writing proposals, preparing all the fieldwork materials, analysis, building the report. All of that is compressing, quickly. Agencies using specialist AI tools designed for researchers report time savings of up to 50% across proposals, fieldwork materials and analysis. One agency took survey design from two days down to one hour.

The agency's experience and judgement is the truly valuable bit. But hours billed for basic production time is the largest part of the invoice. That gap is about to become visible to every client you have — not because clients suddenly understand your cost structure, but because your competitors will tell them.

Your pricing model forces you to change what you sell

Strategic agencies sell (or you could say warrant) the reliability of the data they gather and their great thinking, but they price like contractors: they work out the number of days it will take, multiply by a day rate and then add the third-party costs with a small markup. The thinking time is priced into the number of days, but production time is greater than thinking time, so when AI halves the production days, an agency on that model hands the entire benefit to its client and is left with the same project margin on a smaller number — assuming its cost base falls in step, which for an agency with salaried staff it does not. The percentage holds on paper; the cash does not.

We believe the answer is to start selling 'warranted thinking' — charging clients for the thinking, and for the warrant behind the deliverables: rigorous, validated methodology, professional-grade instruments, real humans carefully selected and interviewed, data handling that survives an audit, and professional-grade AI used with care. In practice that means four fixes: share some of the saving deliberately, put the warrant on the invoice, sell the depth of your thinking, and move up the value chain. Some of the saving will flow back to clients, and should — sharing it is part of what makes the new model easy to buy. But even with production costs falling further than prices, margins on the same work can still climb. What makes the new price stick is the fourth truth below: the ability to warrant the process and the outcome. Strategic research projects underpin the most important decisions your clients will make, the make-or-break stuff. An agency that can show how its answer was reached, and why it can be trusted, is charging for warranted thinking; everyone else is still selling hours. Agencies that move early will fund years of growth out of it. Agencies that wait will have procurement make the move for them, in the opposite direction.

Your junior training route is closing

Juniors have always learned the craft by doing exactly the work AI now does: the drafting of key documents, coding, the first-pass analysis. Take that away and the traditional training ground goes with it. PwC's 2026 AI Jobs Barometer finds AI-exposed entry-level roles are seven times more likely to demand senior-level skills such as judgement.

If you are now selling warranted thinking, every agency has a five-year question to answer: where does the next generation of great thinkers come from? The skills that make thinking warrantable are the human ones in Section 4 — framing, creativity, insightfulness, contextual intelligence, judgement — and they are what training now has to build. The agencies that create a deliberate route, where juniors develop those skills by working inside validated workflows and critiquing machine output under senior review, not by grinding through production, will end up with a talent pipeline their competitors can't match.

Being plausible is no longer worth anything

This one is strange, and it favours agencies that do things properly. Your client's own team, with a generic AI assistant, can now produce a research deck that looks perfectly credible in an afternoon. So looking credible has stopped being worth paying for.

What clients will pay for is the thing the in-house deck can't offer: research findings and insights they can stand behind. Produced by methods that have been validated. Traceable to real human respondents. Auditable from brief to debrief. The job of good research craft has shifted from getting the work done quickly to being able to prove that the answer is right.

That is good news if you're rigorous, because it turns your standards into something you can sell. It also raises the bar. "Trust us, we've been doing this for 20 years" stops being an answer. "Here is the method, here is the trail, here is the person who signed it off" starts being one. It is also what makes the pricing shift above possible: the warrant is what lets you put a price on thinking.

Strategic agencies have the best demand outlook in the industry and the most broken pricing model. The winners will do four things: cut the production time, warrant their work so they can show why the answer is right, price the thinking rather than the days, and rebuild how they recruit and train. The rest of this guide is about how to do that.

What AI is good at, and what people are good at

04

People usually ask whether AI will replace researchers. A more useful question, we'd argue, is which jobs should go to the machine and which should stay with your people. Get that split right and both get better at their work.

What AI does better

Specialist (domain-specific) AI is clearly better than people at the production layer of research.

Speed. Survey design that took days takes hours. A first cut of the analysis on 40 depth interviews can land the same afternoon fieldwork closes.

Coverage. It reads everything. Every transcript, every open-end, every previous study you've run for that client. A human analyst samples.

Multi-modal data power. It reads every type of data — words, numbers, images, audio — as one joined-up set and checks every part against every other part. Think of a haystack with several needles: people find some of them; AI finds every single one, so nothing interesting is missed.

Consistency. A properly built workflow runs the same validated process the same way on project 1 and project 200. No Friday afternoon dip, no variation between one junior and another.

Patience. The tenth version of a questionnaire runs in minutes not hours, so quality stops being limited by budget.

Up to 50% faster

Specialist AI tools can cut the time it takes to complete complex tasks — like writing proposals or designing surveys — by up to 50%.

And because the output quality is higher than generic AI, you don't waste those savings making everything client-ready

What people do better

People are still clearly better at everything that makes research worth commissioning in the first place.

Framing. Working out the real question behind the question in the brief.

Creativity. Inventing a projective technique to unlock understanding, finding the angle of approach that doesn't exist yet — for this brief, this audience, this problem. Translating a new insight into a great new product idea.

Relationships and collaboration. Earning the client's trust and respondents' honesty, being in the room when the decision is made, and holding together clients, suppliers and analysts from different companies.

Insightfulness. The leap from what people say to what's actually going on — hearing the throwaway line, making a connection everyone else has missed, noticing what's conspicuously not said.

Strategic thinking. Connecting the finding to the commercial choice it exists to inform, and telling a client when the question they asked isn't the question they should have.

Contextual intelligence. Knowing the terrain the research lands in: the category's history, the client's politics, the culture being studied.

Judgement you can be held to. Choosing among five reasonable interpretations, putting your name on it, and defending it in front of a sceptical finance director.

Put together, these eight skills are the 'thinking' in warranted thinking. A workflow can prove how an answer was reached; only these skills produce an answer worth standing behind.

AI can draft the deck, but it can't own the recommendation — and the recommendation, with a person answerable for it, is what clients are actually paying for. Human in the lead, not human in the loop.

Two ways agencies get this wrong

The first mistake is putting people on production work out of habit or because the billing model rewards it. Senior researchers formatting charts and doing first drafts. That agency is now more expensive than it needs to be, and clients will notice.

The second mistake is quieter and more damaging: using a generic AI assistant for the judgement work. A general model will give you a confident, sensible-sounding interpretation whether or not it's correct, and a subtly wrong answer looks exactly like a right one. That agency has a reliability problem it may not find out about until a client does.

> The agency model that works is now clear. AI runs the production line inside workflows you've validated; your people do the interpretation, carry the judgement and own the client relationship. In practice the hard part is diary management — making sure the freed senior hours land on thinking rather than on more email and that the thinking gets billed

Specialist AI versus generic AI

05

Most agencies meet AI through a general assistant: ChatGPT, Claude, Copilot. For tactical, low-stakes work they may be all you need. Used for client-facing work, they levy the polishing tax: senior hours spent lifting almost-right output to standard. Nearly 40% of the time saved is lost to rework (Workday, 2026).

But three different things get talked about as if they were the same. A generic assistant: a chat window on a general model. A skills layer: your way of working written into prompts, templates and agents on the same model — prompts direct a model, they don't educate it. And a specialist AI platform, with research method built into the software itself, not supplied by whoever is typing. The difference is how they're built, not how polished they look.

Plausible magic

There's a tendency to think AI can work out anything. Of course, that's not true. The output from an AI system is only as good as the training data it can access.

But here's the rub: even when it hasn't been trained on a subject in any depth, it will still answer your question — and more often than not, the answer will sound very plausible. Plausible isn't what your clients need. They need facts, proof, certainty.

In industries like law, there's a huge canon of published material for the models to train on. In market research, most of the industry's knowledge and know-how sits unpublished — in client debriefs, in researchers' heads, in proprietary agency models. Generic AI tools can't access any of it.

That's why some industries don't just need AI. They need specialist AI.

Why the ceiling exists: the internet knows surprisingly little about research

In law, accountancy, medicine and software, vast bodies of professional know-how are public, so generic models arrive having read the library. There is no equivalent public library for market research. Three barriers keep research expertise off the public web that generic models train on. The craft is largely tacit: it lives in researchers' heads and hands, not on indexed pages. The written material that does exist — proprietary methodologies, client reports, panel databases, syndicated metrics — sits behind enterprise firewalls and NDAs, where web scrapers cannot reach it. And the raw material of fieldwork — transcripts, discarded hypotheses, regional nuance — never leaves private project repositories at all.

So a generic model can define a TURF analysis and draft a passable discussion guide, because some textbook knowledge is online. What it cannot supply is what was never there: the institutional memory built across hundreds of projects, the contextual nuance of why this audience answers this way in this market, and the methodological craft that turns human behaviour into commercial strategy. **Some people believe that prompt engineering and building agents is the answer**, but prompting and agent-building direct a model; they do not educate it. You cannot prompt twenty-five years of fieldwork into a machine, and an agent built on a model that does not understand method simply executes a flawed approach faster. What does close part of the gap is putting your own material in front of the model systematically — your methods, templates and past studies — which is a property of a specialist platform, not of a generic assistant used as it comes.

What's the difference between generic AI and specialist AI systems?

	Generic AI assistant	Specialist research platform
What it is	A generalist model trained on the public web: a mile wide and an inch deep, and far thinner on research than on law, medicine or code, because client work is never published. Output varies run to run, with nothing checking it.	The profession's rulebook delivered as software: thousands of methodologies and detailed research logic built in, plus your own private canon of know-how, connected to the workflow.
Being right	Improvises where evidence is missing, without flagging it. A subtly wrong answer looks identical to a right one.	Validated methods and an audit trail, brief to debrief. 'How do you know?' gets evidence, not reputation.
Memory	Very little accumulates: each session starts again from close to zero, and a prompt library is just words anyone can copy.	Remembers every detail of the project from brief to debrief. Keeps a proper project record, so project 200 benefits from the 199 before it, along with how your agency likes to work.
Process	Varies between runs and between people. Steps can be skipped, reordered or simply forgotten.	The process runs consistently from brief to debrief, with each stage remembering and feeding the next, plus sign-off and an audit trail.
Working together	One person and a chatbot.	Agency and client in the same workspace, with roles, approvals and shared visibility.
Data and security	Your responsibility. Consumer tiers fail client vendor assessments for respondent data.	SOC 2 and ISO 27001 held by the vendor; your data trains nothing. Slots into your clients' assessments.
Who it gets better for	Everyone at once — potentially including every competitor you have, and your clients' competitors too.	The platform improves for your agency only, no one else learns from your brilliant work.

What a specialist platform can give you

Here is what a specialist AI platform is built to deliver. Treat this as a checklist to test any supplier against, us included!

	Seven things to expect from a specialist platform
Quality	**Professional-grade output, brief to debrief.** A standard well above what generic AI assistants and DIY platforms produce, because the method is built into the software rather than improvised each time.
Efficiency	**Quality work in a fraction of the time.** Winning proposals, expert-level discussion guides and complex quant surveys, plus analysis that is both faster and deeper. With generic AI you save less: nearly 40% of the time saved is lost again to rework (Workday, 2026), lifting almost-right output to a client-ready standard.
Your DNA	**Tuned to how your agency works.** A secure store of your agency's own methodologies, templates and house style, so what comes out sounds like your agency rather than a generic tool.
Win rate	**More work won.** Better methodologies, lower costs and deeper findings, which is how you beat both competing agencies and the DIY platforms in a pitch.
Insight	**The findings other tools miss.** Surface-level tools, and stretched humans, tend to stop at the obvious answer. Depth is what turns a project into a long-term client.
Data Security	**Data security beyond a chat window.** Consumer chat tools put client and respondent data outside your control and rarely pass vendor security assessments. A specialist platform carries the SOC 2 and ISO certification burden for you, with contractual guarantees that your data trains nothing.
Margin	**Lower costs, higher margins.** Reducing production costs (proposal writing, fieldwork prep, analysis) means you can offer clients sharper prices and still make more profit.

The legal profession has already been through this

When generic AI arrived, law firms faced the same choice agencies face now: general assistants with legal prompts, or specialist legal AI. The top of the profession moved onto specialist platforms like Harvey, now used by 100,000+ lawyers. Not because a model couldn't draft a clause, but because law runs on three things that research also runs on: answers you can stand behind, confidentiality that survives an audit, and a process that follows professional method.

Then Anthropic launched Claude for Legal in 2026 — and the specialist platforms were inside it. Generic underneath, specialist on top: the model was never the product. There is no reason research will behave differently.

Whatever AI capability you're weighing up, ask one question: **if a client challenged a finding, could you show your workings?** If the answer is a chat history, you have a productivity tool. If the answer is a validated method, a trail back to real respondents and an audit log, you have something you can do strategic work with.

Build, skills, or specialist platform

06

Every agency will have to decide what AI to build or to buy, and this always ends up on one of three kinds of technology, by decision or by drift — with the third coming in two commercial forms:

Build your own specialist platform. Your workflows, methods and IP coded into software you own.

Generic assistant, or a skills layer on one. ChatGPT or Claude used well, or with your way of working written on top: prompts, templates, agents, training.

Rent a specialist platform. A research platform built and maintained by a third-party supplier.

Path one: build your own

Building means creating your own software — a bespoke platform with your agency's workflows, methods and IP coded into it, on top of rented foundation models. You own the product; engineers build and maintain it. The appeal is obvious; the reality less so. You'll need a permanent engineering team of two to four at minimum (in-house or outsourced) plus research leadership time to specify and validate every workflow. Expect twelve to eighteen months before it matches what you could rent today — plus ongoing security certification (SOC 2, ISO 27001) and a permanent QA function, because a platform producing subtly wrong outputs is worse than no platform at all. The part most agencies underestimate is the second bill. Foundation models change every few months and funded vendors ship continually, so an in-house build needs perpetual investment just to stay level. It is a change of business model: you become a technology company as well as a research agency. And ask early: what happens when the person who built it leaves?

Path two: create skills

No software gets built here. Instead you take a general assistant off the shelf — ChatGPT, Claude, Copilot — and teach it your way of working through written instructions: a curated prompt library, templates, training for every grade, a governance policy, and a senior owner of quality. Done seriously, it's a programme, not a stack of licences. The strengths are real: a few months to build the skills manual and tools, immediately useful for tactical work like first drafts and desk research. Your people are probably using these tools already. The limits are structural. Output varies between runs and between people, so seniors end up checking everything that matters. Prompts direct a model; they don't educate it, and the underlying models know little about elite market research. Nothing accumulates, the library diffuses to competitors (it is only words), every model change forces a rebuild, and there's no audit trail. Counted honestly, a serious skills layer costs more to run than renting a platform. None of that matters for low-stakes work. All of it matters for the work strategic agencies are actually paid for.

Path three: rent a specialist platform

Renting a platform built for research, where the methodology lives in the software rather than in whoever happens to be typing. Workflows run brief to debrief, outputs follow validated methods, and the vendor carries the security certifications procurement teams ask about. Your side is smaller but real: a few weeks of onboarding, feeding in house methods so outputs sound like you, and the internal work of changing how projects run. The honest trade-offs of any rental apply: you depend on the vendor's roadmap, and the bill grows with usage.

Start with a better question

Most agencies start by asking what each option costs. That's understandable, and it starts in the wrong place: AI capability behaves less like a cost line than a revenue lever, because it changes what you can win, what you can charge and what you keep. The better question is what each path lets you win, and whether that beats what it truly costs to own. An expensive option that wins you work can easily beat a cheap one that wins you nothing — and a cheap option is only cheap if you have counted all of its costs.

The money, honestly

Midpoint estimates for a thirty-person agency in 2026. Adjust for your own scale.

Cost line	Generic	Skills layer	Build	Rent
Initial build cost	—	£60k	£370k+	—
Annual licence fees	£19k	£19k	£30k	£60k
Annual maintenance and hosting	—	£15k	£200k	Vendor
Annual upgrades and improvement	—	£45k	£185k	Vendor
Certification, year one	£33–70k	£33–70k	£33–70k	Vendor
Certification, annually thereafter	£19–33k	£19–33k	£19–33k	Vendor
Year one total	**£52–89k**	**£112–149k**	**£433–470k**	**£60k**
Annual cost thereafter	**£38–52k**	**£98–112k**	**£434–448k**	**£60k**

Source: Insightology internal cost data together with published estimates from a range of sources; 2026 midpoints. Certification: 2026 price guides for SOC 2 and ISO 27001 together. Skills layer assumes one full-time person building and maintaining it. Build is a floor: our own platform cost close to £900k in its first year. Senior checking and rework are excluded because they are time, not invoice; on the generic and skills paths they are usually the largest number on the page.

Only the build changes the order of magnitude; the other three sit in the low hundreds of thousands or less. The skills layer, done properly, costs more to run than renting: about £105k a year against £60k, once the person who maintains it, the rebuild cycle and certification are counted. And the cheapest column is not the cheapest option: a generic assistant looks unbeatable until you price the rework, every client-facing output lifted to standard by the people whose time costs most.

Remember you're chasing a moving target

A funded supplier doesn't stand still. Specialist AI platforms invest in the range of £750k to £1.5m a year improving what they sell. So a build budget buys you parity with where the market is today. Staying level means matching vendor-scale spend indefinitely. And parity isn't the win condition anyway: a build only pays off if it is better than what everyone else can rent, and stays better.

The best alternative test

Any platform, built or bought, only wins you more projects if it's better than what competitors can get. And competitors can rent the best specialist platform tomorrow. So measure a build against the best available platform, not against having nothing.

A rented platform at ~£60k a year delivering ~£295k of annual benefit nets ~£235k. To beat that, a build carrying ~£435k a year of running costs must generate ~£670k of margin a year just to draw level, before recovering the £370k+ build cost. Given the extra risk, a build should target five to ten times what renting returns. It earns its keep only on things a supplier can't replicate: proprietary data, unshareable methods, or a sellable product.

Five questions to work out your path

Answer these in order and stop when you get an answer.

01 **Is your technology what makes you different?** Would clients buy your platform without your best people attached to it? Is your tech the reason you consistently win pitches and make more margin than competitors? If yes, and you can fund £370k+ now plus ~£435k a year, build it in house. The bar is being better than the best platform anyone can rent, and staying there. If the answer is no, move to the next question.

02 **Which end of the market do you sell to?** If it's tactical work, quick-turn testing, trackers and low-stakes studies, a generic assistant or a skills layer is genuinely enough. If it's strategic work, keep going.

03 **Do your clients audit you?** If not yet, they soon will — AI is pushing vendor assessments and respondent-data scrutiny up every client's agenda. That is exactly where generic AI assistants fall over, and where a build makes compliance your permanent job. Either way, you need capability that carries its own compliance. Keep going.

04 **Does your value depend on what your agency has learned over time?** If project 200 ought to benefit from the 199 before it, all those methods, templates and findings, you need a system that gets better with every project rather than a chat window. A skills layer won't, without a shared knowledge store you maintain. Builds can. Keep going.

05 **Do the maths.** If a specialist platform costs less per year than the margin on a handful of strategic projects it helps you win or deliver better, use the platform. If you're confident an owned platform clears its full ongoing cost, build it, but go in knowing a self-build should return five to ten times what renting would, to justify the extra risk.

Where each path wins, including where we lose

Building wins for the small number of agencies whose technology genuinely is the product: unique IP, proprietary data, methods they can package and sell, and a balance sheet that supports a permanent engineering team. Few agencies clear that bar, and it means becoming a technology company too.

A generic assistant or skills layer wins for tactical agencies, and for the tactical part of any agency. Fast, cheap and very good at commodity production. If that's your book of business, don't pay for more than you need, including from us.

A specialist platform wins for strategic agencies whose real product is warranted thinking. You rent the arms race instead of running in it, the compliance burden sits with someone else, your clients get a process they can audit, and you pay in proportion to the work you win.

The short version. If your tech is the product, build. If your work is tactical, a generic assistant or a skills layer is enough. If your judgement is the product and your work is strategic, rent a specialist platform and spend the difference on the people who do the interpreting.

The financial benefits of AI — and how to calculate yours

07

AI pays an agency back through two streams: it helps you win and keep more revenue, and it cuts the cost of doing the work. Both are real, but they are not equal — cost savings are bounded by your production bill and can only be taken once, while revenue growth is bounded only by the market and builds year after year. The paths from Section 6 deliver the two streams to very different degrees, and this section puts numbers on it in one worked comparison.

Benefit one: more revenue

Better proposals, and more of them. AI carries the production work of a proposal — the method, the costings, the document, the polish — which frees the senior time that actually wins pitches: thinking about the client's problem. Faster proposals also mean more proposals, and the time-sensitive brief you used to decline becomes winnable. Win rate and pipeline both move, and they compound.

Better thinking on every project you deliver. Fewer production hours means more thinking time per project, and the quality of insight rises with it. Clients notice, and the payoff is the cheapest revenue there is: repeat rates and expanded briefs. Add capacity without headcount, sharper prices, thinking priced at its premium and new revenue up the value chain, plus one that costs nothing: a governed audit trail is a quality signal procurement now looks for.

Where the paths differ. A specialist platform and a well-built in-house system perform at professional grade, so you get the full revenue effect. A skills layer on a generic assistant does not — output varies, subtle errors slip through, and the polish still needs senior rework — so it helps you win and retain **less revenue than the other two options.** That gap is why the comparison below credits a skills layer with half the revenue growth.

Benefit two: lower costs

The production bill shrinks. The 50–70% of project labour that is production work compresses hard, and one particularly satisfying line of waste shrinks with it: the cost of proposals you don't win. When a credible proposal takes hours rather than days, losing a pitch stops burning a week of senior time.

Where the paths differ. Task-level savings run up to half. Across a whole book of business we model renting and building at a 30% production saving and a skills layer at 15%, and the skills saving erodes further in practice, because lower-quality output means senior people spend real hours checking and correcting it.

The catch, and the better use of it. Cost savings only turn into cash if you remove the heads whose hours were saved — and many agencies, rightly, won't want to. The alternative is better anyway: redeploy the freed time into Benefit one. More proposals out, more client contact, more thinking on every project. The saving doesn't show up as a smaller payroll; it shows up as a bigger top line.

A worked cost–benefit comparison: a thirty-person agency

Here is the whole section in one table, for a thirty-person strategic agency: ~£4.0m gross billings, £2.8m net fee income after pass-through, and roughly £780k of payroll doing production work. Assumptions, so you can swap in your own: production costs fall 30% on rent and build and 15% on skills — below the task-level savings agencies report, since not every hour saved converts to cost. The saving is split three ways: a third to clients as sharper prices, a third banked through natural attrition, a third kept as project margin. Fee income grows 10% on rent and build and 5% on skills, at a 50% margin on incremental work. Costs: platform ~£60k a year; skills ~£130k in year one then ~£105k a year; build £370k up front, ~£450k across year one and ~£440k a year once live from year two. Benefits are half-weighted in year one; a build yields nothing until live. Three-year view, constant prices; sterling at $1.35/£.

Path	Cost over 3 years	Annual benefit once running	Net gain over 3 years	Break-even point
Rent a specialist platform	~£180k	~£295k (£140k margin on new fee income + £78k attrition + £78k project margin)	**+£565k**	Within six months
Create skills	~£340k	~£150k (£70k margin on new fee income + £39k attrition + £39k project margin)	**+£30k**	During year three
Build in-house	~£1.33m (from reserves, before any benefit)	~£295k — but only from year two, once the platform is live	**–£740k**	Not within three years

Same benefits engine, three very different outcomes. Revenue is the largest component of the benefit on every path, on a modest 10% growth assumption. A build does not pay back for an agency this size: around £1.33m of cost against ~£590k of benefit over the period, using the floor of Section 6's build range, not a realistic one. The skills layer barely washes its face: ~£340k of cost against ~£370k of benefit, positive only in year three. Renting pays back within months. That is Section 6's best alternative test, drawn as a picture.

The short version. Whichever AI route you take, you should see both revenue and cost-saving benefits — otherwise there's no point starting this transition. Build your financial evaluation using your own numbers. Our only watch-outs: work in fee income rather than turnover; account for all the costs of the AI route you choose; and when you get to the revenue calculations for rent versus build, be realistic — it's unlikely that your in-house build will massively outperform the rented version. Not impossible. Just unlikely!

A 90-day plan

08

All of the above turns into three decisions. Renting gets you through them in a quarter; a skills layer takes four to nine months and a build twelve to twenty-four.

Days 1 to 30: work out where you stand and pick a path

- **Count your production hours.** Take your last ten projects and split every billable hour into production and thinking, on real timesheets rather than impressions. Most agencies land between 50 and 70% production. That number is both your opportunity and your exposure.
- **Find the AI already in the building.** Someone is using something on client data right now. Find out who, what and under what rules.
- **Check your audit position.** Pull the last three vendor security assessments a client sent you and ask honestly whether your current AI use would pass them.
- **Run the five questions from Section 6** with your leadership team, build question first.

Decision at day 30: which path, written down.

Days 31 to 60: pilot it on real work

- **Run your chosen approach on three to five live projects,** not test cases, baselined, measuring hours per deliverable, proposal turnaround, and senior time spent making output client-ready.
- **Reprice one service line to outputs.** Pick something you understand well, proposals or a standard study type, and quote a fixed price for the result rather than a number of days. One line rather than the whole rate card. This is how you learn the margin shift without risking much.
- **Write one page of governance.** What AI can touch, what it must never touch, and who signs off.

Decision at day 60: does the pilot data justify rolling it out?

Days 61 to 90: roll out and rebuild

- **Extend across the project lifecycle,** brief to debrief, for the teams and study types where the pilot worked.
- **Reprice the next service lines,** using the cost data the pilot gave you.
- **Rebuild your junior training route.** Move trainee development onto validated workflows so juniors learn method by doing supervised strategic work, rather than by grinding through production that no longer exists.
- **Six quarterly numbers:** hours per deliverable, proposals, win rate, thinking time as % of fees, project margin, platform cost vs revenue, audit pass rate.

Decision at day 90: your operating model for the next year, with numbers attached.

None of this requires betting the agency. It takes one quarter of proper management attention, which is a good deal less than the market is about to ask of the agencies that do nothing.

About Insightology

Insightology.ai is the first AI platform built specifically for market research agencies, rather than generic AI with research prompts bolted on. It is method agnostic — qual, quant and mixed-method in one platform — rather than a point tool for a single technique, and it is built for the complex, multi-faceted programmes strategic agencies actually sell: multi-market, multi-stage work managed end to end in one place. It covers the whole strategic project: decoding the brief, methodology design, proposal writing, survey and questionnaire design, discussion guides, fieldwork materials, analysis and reporting. Everything runs as a validated workflow with an audit trail, built to reproduce the quality of expert researchers consistently.

Go to www.insightology.ai to find out more and book a demo if it looks relevant to you

It is built to do the seven things listed in Section 5: professional-grade output from brief to debrief, the same work in a fraction of the time, tuning to your agency's own methodologies and house style, a better win rate in pitches, deeper insight than surface-level tools reach, data security your clients' auditors will accept, and lower production costs that let you sharpen prices and still make more margin.

Agencies using it report time savings of up to 50% across proposals, fieldwork materials and analysis. Pricing is per won project with unlimited users. Proposals are free and unlimited, and credits only apply once a project is won and reaches fieldwork, so your cost follows your revenue rather than your headcount.

The argument in this guide applies to agencies doing strategic work, for clients who audit them, where warranted thinking is the real product. If that describes you, the quickest way to test it is the way Section 8 recommends: a proper pilot on three to five live projects, baselined, with the numbers measured and a written definition of what would make you say no.