

02 / THE PLAYBOOKS

The Selectability *playbooks*

Eleven steps, fourteen builds, one worked example



The Selectability *playbooks*

The how-to companion to Seen Isn't Selected · Ben Stokes

A NOTE BEFORE YOU START

Seen Isn't Selected makes the argument and teaches the method. These playbooks hold the how-to, in the order you'd do it: get the strategy right, find out where you stand, plan, then build and run each of the fourteen builds. Every build follows the same anatomy, from its footprint on the map to the Loop that tells you whether it worked. One worked example runs through all of them: Kestrel Advisory, an invented operations consultancy for mid-sized manufacturers. Every figure about Kestrel, its clients and its people is invented. Every figure about the world is sourced where it's used.

VERSION	Playbooks v2.0 · October 2026
UPDATES	The playbooks are updated as the evidence changes. The latest edition is always free at rocmode.io/playbooks
RIGHTS	© 2026 Ben Stokes. All rights reserved.

Contents

PART ONE	Strategy	4
	Steps 1 to 4: what's true, why you, who you win, how you look and sound	
	STEPS 1–4 Strategy: the facts, Credentials, Brand Strategy, Brand Identity	5
PART TWO	Where you stand	13
	Step 5: map the moments, then score them	
	STEP 5 Map the moments	14
	STEP 5 Score the moments	31
PART THREE	Plans	46
	Steps 6 to 9: GTM, Website, Marketing and Sales strategies	
	STEPS 6–9 Plans: GTM, Website, Marketing and Sales	47
PART FOUR	Builds, and running them	59
	Fourteen builds in six categories, then steps 10 and 11	
	OWNED The website	60
	OWNED Content	99
	OWNED Social	120
	PAID Advertising	142
	EARNED Being named: search and AI answers	166
	EARNED PR and reviews	192
	PARTNERED Partners and channels	209
	PARTNERED Creators	224
	DIRECT Outbound and account-based	242
	DIRECT The pitch	263
	DIRECT Events	285
	RELATIONSHIP Email	299
	RELATIONSHIP Automation	315
	RELATIONSHIP Customers	333
	STEP 10 The Campaign Engine	350
	STEP 11 The Loop	367

PART ONE

Strategy

Steps 1 to 4: what's true, why you, who you win, how you look and sound



Part one: Strategy

One journey, eleven steps

Every build in these playbooks reads what was decided before it started: what's true, why a buyer should choose you, who chooses, and where you stand today. Get those right once and the builds agree with each other. Leave them out, and each build invents its own version: the website describes one firm, the pitch another, the ads a third.

So the playbooks follow one journey, in eleven steps. Each step answers one question and reads only the steps before it. RocMode's Brand Journey follows the same steps, with the same names, and so does the book.


One journey, eleven steps: the playbooks, the book and RocMode

Each step answers one question and reads only the steps before it.

Part one: Strategy



Part two: Diagnosis



Part three: Plans



Part four: Builds, and running them



The Brand Journey · eleven steps in four parts

Part	Steps	What it settles
One: Strategy	1. Get the facts · 2. Credentials · 3. Brand Strategy · 4. Brand Identity	What's true, why choose you, who you win, how you look and sound

Part	Steps	What it settles
Two: Diagnosis	5. Where you stand: map the moments, then score them	Every moment that matters, and where buyers are lost today, and to whom
Three: Plans	6. GTM Strategy · 7. Website Strategy · 8. Marketing Strategy · 9. Sales Strategy	Which moments to win, what the site holds, which builds and budget, how accounts are won and kept
Four: Builds, and running them	The fourteen build playbooks, in six categories · 10. Campaign Engine · 11. The Loop	The work itself, twelve months of campaigns, and what changes when the numbers move

This part covers steps one to four. It maps onto the Selectability Model: the facts are the evidence; Credentials hold the offer and the one truth; Brand Strategy holds the reason; Brand Identity is what makes a brand familiar, and familiar feels safe.

The worked example. Every playbook uses one firm: Kestrel Advisory, an invented operations consultancy. Kestrel, its buyers' questions and every figure attached to them are invented. A real firm fills these in from its own records.

1. Get the facts: what do we already have?

Everything after this step can only use what's gathered here, so gather it once, from everywhere, and keep the source of every fact. In RocMode, this is ingestion: the firm's website, documents, profiles and client records are read into one store, and every fact keeps a link to where it came from.

Kestrel's sources (invented):

Source	What it holds
The website and the old blog	How Kestrel describes itself today, which turns out to be "we help organisations transform"
Standard terms and the contract template	The fixed price, the six weeks, the way out
The client list and sign-off letters	Who Kestrel works for, and the figures clients have agreed to share
Proposals, tenders and briefs	What buyers asked for, in their words, and what they judged on

Source	What it holds
Sales call notes and lost-deal interviews	The questions buyers ask, and why deals were lost
Client debriefs	What happened after the review, and what clients needed then
People's profiles	Who will be on site, and their plant experience

Two rules make this step pay off later. **Keep the buyer's words**, not a summary of them: they become the moments in step five. And **gather the evidence for losses as carefully as wins**: the lost-deal interviews are where the switching moments hide.

2. Credentials: what's true?

Credentials come first because everything else repeats them. In RocMode, credentials are built by reading what the firm has already published and stored (its website, documents, profiles and client records) into one store of facts. Nothing that can't be traced to a source gets through. The credentials deck, twelve sections in three acts (context and belief; capability and proof; trust and future), is a view of that store, not the store itself.

Every claim ends in one of four states, and only the first two can be used:

- **Sourced**: traced to a document, a record or a page.
- **Signed**: approved by a named person who can vouch for it.
- **Bound**: a number that may only appear with its live figure, so it can't go stale.
- **Removed**: no source can support it.

Kestrel's offer. A fixed-price, six-week operations review, on site, for manufacturers with 100 to 500 staff in food and drink, packaging, metal and engineering, and plastics and rubber. An optional improvement programme follows it.

Kestrel's claims register (extract):

Claim	Source	Signed off by	State
Fixed-price, six-week operations review	Standard terms of engagement	Managing partner	Signed
On site, not remote	Engagement method; client references	Head of delivery	Signed

Claim	Source	Signed off by	State
For manufacturers with 100 to 500 staff, in four sectors	Client list	Managing partner	Signed
If the review doesn't find more than it costs, the client doesn't go ahead	Contract clause	Managing partner	Signed
Named plant cases, with before-and-after figures	Client sign-off letters	Head of delivery, and each client	Signed, case by case
"The review pays for itself"	Results across all reviews	—	Bound: shown only with the live figure from the results page
"Leading operations consultancy"	None	—	Removed

The register is short on purpose. Every build that follows can only say what's in it.

3. Brand Strategy: why choose us, and who do we win?

In RocMode, the Brand Strategy Blueprint has six sections, each with a verdict, the findings behind it and the reasoning: the big idea, where you play, who you win, how you sound, where you're going, and what to do first. It reads only approved credentials. This chapter adds one thing: **the reason to be chosen as its own line**, so every build quotes the same words.

Section	Kestrel
The reason to be chosen	"We find the savings on the floor in six weeks, at a fixed price, and the review pays for itself or you don't proceed." Every part traces to a claim above; "pays for itself" is bound to the results page
The big idea	Savings are found on the floor, not in new equipment
Where you play	Mid-sized manufacturers in four sectors. Against a large consulting firm, a specialist lean consultant, the in-house improvement team, buying new equipment, or doing nothing

Section	Kestrel
Who you win	Plants with flat output or rising costs first; investor-owned manufacturers second. The people inside them, the choosers, are listed below. Personas live here and nowhere else
What you say	The words to use and the words to avoid. Say "savings on the floor", not "operational excellence"; say "review", not "transformation"
Where you're going	The firm mid-sized plants think of first, with an annual benchmark the category quotes
What to do first	Close the gap in results data, so "pays for itself" can carry a figure; then the sector pages and the benchmark

Who you win: the choosers

Brand Strategy names the people as well as the markets, because everything after it needs one list of them. In RocMode, the blueprint's "who you win" section already holds the personas and their jobs to be done. They live here and nowhere else: the moment map, the plans and the builds all point back to this list. Step five finds them by role in the buying group, so none is left to chance; this is the list it confirms.

Chooser	Their part in the decision	What they need from Kestrel
Chief operating officer	Starts it and sponsors it	A diagnosis they recognise, and a case they can take upstairs
Owner or managing director	Has the final say. In owner-led firms, is the committee	Trust in the firm, its people and its judgement
Finance director	Approves the spend and owns the numbers	Price, terms, the downside, and numbers that survive the board
Plant manager	Lives with the result, and can make or break it on the floor	Proof the work happens on the floor, and what it costs the team in time
Continuous improvement lead	Often the hidden blocker: an outside firm can look like a verdict on their work	To be worked with, not around, and to keep the methods afterwards

Chooser	Their part in the decision	What they need from Kestrel
Procurement, with site safety	Runs the paperwork and site access	Supplier approval, insurance, inductions and payment terms
Private equity operating partner	In investor-owned firms, often starts it	Speed, portfolio experience and reporting the board can use
AI assistant	Builds the first shortlist, for any of the above	Quotable facts: who it's for, what it costs, what it returns, and who says so

Shift supervisors and workforce representatives matter too, but their questions reach Kestrel through the plant manager, so they sit in that row.

Of the alternatives, **buying new equipment** matters most. Capital spending competes for the same budget, and it's the option most plants reach for first.

The story spine. Every build tells the same story in its own way:

1. **The problem, in the buyer's words:** "Output is flat and costs are rising, and we can't see where the waste is."
2. **The reason:** savings found on the floor in six weeks, at a fixed price.
3. **The proof pattern:** every claim carries a named plant, a before-and-after number and the plant manager's own words.
4. **The way out:** if the review doesn't find more than it costs, the client doesn't go ahead.

One change from RocMode's current blueprint: its "how you sound" section becomes **what you say**, the words and the order of the message. How the brand sounds, its voice, belongs to Brand Identity, so it's decided once.

4. Brand Identity: how do we look and sound?

In RocMode, Brand Identity is built from the brand strategy: it reads the big idea and the positioning and turns them into how the brand looks and sounds. It has two jobs. **Be recognised before the name is read**, because buyers trust what feels familiar. And **sound like one firm everywhere**, so the website, the pitch and the ads read as the same people.

Element	Kestrel	Why
The name	Kestrel Advisory, in full, every time. Never plain "Kestrel"	Other firms share the word, and assistants mix them up

Element	Kestrel	Why
The mark	A kestrel hovering: the bird that holds still to see what others miss	The big idea: look closely on the floor before you spend
Colour	One owned colour, used the same way everywhere, with its values written down	Distinctive assets only work if they never change
Lines	"Savings on the floor." "Fixed price. Six weeks. Or you don't proceed."	Short enough to repeat, and every word is a signed claim
Pictures	Real plant floors and real people from client plants, with permission. No stock handshakes	The proof pattern, in pictures
Voice	Plain, practical and specific. Numbers with their sources. The plant manager's words, not the consultant's	A COO should be able to forward any page to the floor

One rule ties Identity back to step two: **a line can only say what Credentials can sign.** "Or you don't proceed" is a line because the contract clause is signed. "The leading operations consultancy" never becomes one.

What comes next

With the facts, the truth, the reason, the choosers and the identity decided, step five finds every moment that matters and measures where Kestrel stands in each one: [Where you stand, part one: map the moments](#). Then Part three turns the map into plans: which moments to win (GTM), what the site must hold ([the Website Playbook](#)), which builds and in what order (Marketing, with each build's footprint), and how accounts are won and kept (Sales).

Feeling a bit overwhelmed?

You should be, a little. Before a single moment is mapped: every source gathered with its provenance, a claims register with every signature, a brand strategy in six parts, a named list of choosers, and an identity that only says what's signed. That's for one firm with one main offer. Most firms skip straight to the builds, which is why their website, their pitch and their ads describe three different companies.

RocMode does the heavy part for you:

- **The facts, gathered once.** RocMode reads your website, documents, profiles and records into one store, and keeps the source of every fact.

- **Your credentials.** It builds your credentials from that store, and won't let a claim through without a source.
- **Your brand strategy and identity.** The big idea, where you play, who you win and what you say; then how you look and sound, built from that strategy. You review, change and sign off each one, and every step after this reads them.

[Join the RocMode beta](#)

PART TWO

Where you stand

Step 5: map the moments, then score them



Where You Stand, Part One: Map the Moments

What this step is for


Every build, every page and every score inherits the moment map. Get the moments wrong and everything downstream is precisely wrong. So this step comes before any plan: find every moment that matters, show where each one came from, and only then measure where you stand.

Most maps of buyers are brainstormed in a workshop, and they show it. They hold the buyers the team already talks to and the questions sales already hears. The first draft of Kestrel's map was made that way, and it looked complete. It wasn't. Derived properly, it gained two journey steps, eighteen cells and twenty-seven moments. Kestrel's buyers had always been there; the brainstorm had missed them.

The method: four things you already know generate the candidates. Evidence checks each one. Five questions check the whole.

Moments are derived from four inputs, then checked twice

Nothing goes on the map without a reason, and nothing stays without evidence.



How the moment map is derived · four inputs, two checks

The inputs come from the steps before this one. The industry and the offer come from getting the facts and from Credentials. The alternatives and the difference come from Brand Strategy: where you play, and the reason to be chosen. The output feeds the Score in part two of this step, and then the GTM Strategy, which decides which moments to win.

Input 1: the industry gives the structure

The industry decides four things before you know anything about your own buyers: who sits in the buying group, the steps they go through, the ways they buy, and what sets them off.

Who chooses: the buying group

Research on how organisations buy has named the roles for fifty years: the people who use the purchase, influence it, decide, buy, and can block it ([Webster and Wind, 1972](#)). Later writers added the person who starts it ([Bonoma, 1982](#)). Fill each role for your category, and add the machine: an AI assistant now influences, and often gatekeeps, the shortlist. Gartner puts business buying groups at five to sixteen people across as many as four functions ([Gartner, 2025](#)).

Role	Kestrel's chooser	Why they're on the map
Starts it	Chief operating officer; in investor-owned firms, the operating partner	Owens output and cost; feels the problem first
Uses it	Plant manager	Lives with the review on the floor
Influences it	Continuous improvement lead; the AI assistant	The lead's view carries weight inside; the assistant builds the first shortlist
Can block it	Continuous improvement lead; procurement, with site safety	The lead can treat an outside firm as a threat; procurement controls supplier approval and site access
Decides	Owner or managing director; finance director, for the spend	Final say, and the budget
Buys	Procurement	Runs the paperwork and the terms

The roles found every chooser in Kestrel's first draft, and gave each one a reason to be there. The improvement lead had been on the map by luck; now they're there by role, as the influencer who can block.

The steps: the jobs every buying group must do

Gartner's research describes six jobs every business buying group has to get done, in no fixed order: identify the problem, explore solutions, build requirements, select a supplier, validate the choice, and build agreement inside the business. Add what comes before them (the trigger) and after them (buying, delivery, and renewal or referral), and you have the journey steps.

Kestrel's step	The buying job behind it	In Kestrel's first draft?
1. Trigger	Before the journey: the situation that starts it	Yes
2. Framing	Problem identification	Yes
3. Options	Solution exploration	Yes
4. Requirements	Requirements building	No: added
5. Comparing	Supplier selection	Yes
6. Checking	Validation	No: added (it was hidden inside comparing)
7. Building the case	Consensus creation	Yes
8. Buying	After the jobs: contract and start	Yes
9. Delivery	The job itself	Yes
10. Extend and refer	Renewal, expansion and referral	Yes

Requirements is where procurement writes the brief and the finance director sets the criteria. Checking is where buyers ring references and test claims. A brainstorm skips both, because sales rarely sees them happen.

How they buy: buy class and decision path

Robinson, Faris and Wind's buy classes ([1967](#)) say whether it's a first purchase (a new task), a changed repeat (a modified rebuy) or a straight repeat. Cross that with who decides, and Kestrel has five buying modes:

Mode	Buy class	Who's in the room
Planned committee	New task	COO, finance director, plant manager, procurement; the owner signs
Owner-led	New task	The owner, with the plant manager
Investor-led	New task	Operating partner, COO, finance director
Urgent	New task, compressed	COO and owner
Repeat	Modified or straight rebuy	The next plant, the next review, or the programme renewed; often procurement alone

The repeat mode is new. The first draft had renewal moments but no path for them.

What sets it off: category entry points

Jenni Romaniuk's category entry points are the situations that bring a category to mind. The Ehrenberg-Bass Institute finds them with seven W's: why, when, where, while, with whom, with what, and how the buyer is feeling ([Oxford University Press](#); [Ehrenberg-Bass](#)). For a business, ask each W from two sides: the business and the professional.

W	Kestrel's entry points	What it changes
Why	Flat output; rising costs; cash tied up in stock; scrap and rework; a new line below rate; after an acquisition; automation not paying back	The question. These seven become the needs, and split moments
When	After a bad quarter; budget season; year-end results	The timing
Where	One plant, or several	The scope
While	During a capital spending decision; during a customer audit	The timing
With whom	A new COO or plant manager in their first months; a new owner	The timing, and who's in the room
With what	A new line just installed; new automation	The question (already in "why")
How feeling	Pressure from a key customer; embarrassment after an audit	The tone

This gives a rule the first draft lacked. **An entry point that changes the question is a need, and splits moments. One that only changes the timing is an occasion.** Occasions don't go on the moment map. They go in the campaign calendar: budget season, a new leader's first hundred days, year-end results.

Input 2: what you sell gives the needs

Every offer is hired to get a job done, and every job runs through the same eight steps: define, locate, prepare, confirm, execute, monitor, modify and conclude ([Bettencourt and Ulwick, 2008](#)). Each step raises its own questions, for someone in the buying group. Map every offer's job to the end. Most maps stop at the purchase, which is step four.

Kestrel's review is hired for one job: **find and capture the savings in our plant.** The improvement programme is hired for the last three steps of it: make the gains last.

Job step	The question it raises	Whose	Where it lands	In the first draft?
Define	"Where's the money going: labour, scrap, stock or energy?"	Finance director	Framing	Yes

Job step	The question it raises	Whose	Where it lands	In the first draft?
Locate	"What will they need from our data and our team?"	Improvement lead	Requirements	No
Prepare	"What do we need from them on the floor, and what can't we spare?"	Plant manager	Requirements	No
Prepare	"How much of my team's time will it take?"	Plant manager	Building the case	Yes
Prepare	"What do we need ready before they arrive?"	Plant manager	Buying	No
Prepare	"What data do they need from us before week one?"	Improvement lead	Buying	No
Confirm	"Will the numbers stand up at the board?"	Finance director	Building the case	Yes
Confirm	"What does an operations review cost, and what does it return?"	AI assistant	Building the case	Yes
Execute	"Are we on track, and what are they finding?"	COO	Delivery	Yes
Execute	"What do my people need to do this week?"	Plant manager	Delivery	Yes
Monitor	"Are the savings real, and are they being counted properly?"	Finance director	Delivery	No
Monitor	"Are the savings showing in the accounts?"	Finance director	Extend and refer	Yes
Modify	"Results are slipping. What do we do now?"	Plant manager	Extend and refer	No (the first draft asked only "did it stick?")

Job step	The question it raises	Whose	Where it lands	In the first draft?
Conclude	"Will we own the methods when they leave?"	Improvement lead	Delivery	Yes
Conclude	"Can we run the next review ourselves?"	Improvement lead	Extend and refer	Yes
Conclude	"Did it pay back? Do we go into the programme?"	COO	Extend and refer	Yes

Six of the sixteen were missing. Nearly all sit after the purchase or just before the start, where marketing rarely looks and sales has already moved on. Those are exactly the moments that decide whether Kestrel is kept, extended and recommended.

Each step holds more needs than one question. Outcome-driven research typically finds 50 to 150 separate needs for a single job ([Strategyn](#)). Don't put each one on the map. A moment is one person at one step; the needs are held inside it, as the cell's pain, gain and objections.

Input 3: the alternatives give the switching moments

Every buyer is choosing between you and something else, and doing nothing is always on the list. For each alternative, four forces decide whether they move: the **push** of today's situation, the **pull** of the new option, the **anxiety** about it, and the **habit** of the present ([The Re-Wired Group](#)). A buyer switches when push and pull outweigh anxiety and habit. Each force is a question someone in the buying group will ask, or an objection they'll raise.

Five alternatives and four forces give Kestrel twenty switching needs:

Alternative	Push (away from it)	Pull (towards Kestrel)	Anxiety (about Kestrel)	Habit (of staying put)	Where it lands
Doing nothing	A key customer is threatening to move the contract	A fixed price and a way out make acting cheap	"What if it disrupts production?"	"We always firefight; next quarter will be better"	Owner · options; plant manager · building the case; COO · trigger
New equipment	The old line can't keep up	Find the capacity in the line you have, before you spend	"What if they find nothing, and we've delayed the purchase?"	The capital request is already in the budget	COO and finance director · options; owner · building the case

Alternative	Push (away from it)	Pull (towards Kestrel)	Anxiety (about Kestrel)	Habit (of staying put)	Where it lands
The in-house team	Lean didn't hold last time	Outside eyes for six weeks, and the team keeps the methods	"Will this make my team look bad?"	"We have people for this"	Improvement lead · framing, options, comparing, checking, building the case
A large firm	Day rates, long reports, and partners who vanish after the pitch	Fixed price, on the floor, and the people who pitch do the work	"Will the board trust a smaller name?"	"Nobody gets fired for hiring a big firm"	Owner · comparing and checking; finance director · comparing
A lean specialist	Methods, without the money case	Savings in the finance director's terms	"Are they lean consultants by another name?"	"We already know the lean playbook"	Improvement lead and finance director · framing; COO · comparing

The forces mostly don't add cells. They fill them: every switching need becomes a cell's pain, pull or objection, in the buyer's words. That's what turns a generic cell ("comparing options") into one a writer can brief from.

Two patterns show up in almost every category. **Habit is the force most plans ignore**, and it's usually the strongest: the capital request already in the budget beats any argument Kestrel makes. And **the anxiety about you is specific to each alternative**. A buyer leaving a large firm worries about the board; a buyer leaving the in-house team worries about the team. One generic reassurance answers neither.

Input 4: the difference gives the problem insight

A real difference implies that buyers misread their problem, or misjudge the risk of fixing it. Name what they misread, and you find the moments where your reason decides the choice. Then turn it round as a test: **if no moment needs a part of your difference, it isn't a difference**.

Part of Kestrel's difference	What buyers misread	The moment it decides
Savings are on the floor, not in new equipment (the big idea)	That flat output needs more capacity	COO · framing: "Is it people, process, equipment or planning?"; COO and finance director · options
Fixed price	That consulting risk is open-ended	Finance director · comparing: "Fixed price or day rates?"

Part of Kestrel's difference	What buyers misread	The moment it decides
The way out	That a review that finds little is money lost	Finance director · comparing and buying: "What if it finds nothing? How do we stop?"
Six weeks	That a review ties up the plant for months	Plant manager · building the case: "How much of my team's time will it take?"; COO · buying
On site	That reviews happen in meeting rooms	Plant manager · comparing: "On the floor, or in a meeting room?"
Findings in the plant manager's language	That consultants tell you what you already know, in words nobody uses	Plant manager · framing; plant manager · checking: "Can I see their people on a floor like mine?"

Every part decides at least one moment, so every part earns its place in the reason. One of those moments wasn't in the first draft: the plant manager's checking moment, where "findings in your language" is proved by meeting the people.

The difference also tells you which framing moments to own. Kestrel wins most when it reaches the COO before they've decided the problem is capacity. After that, the capital request is in the budget and habit takes over. That's why the framing moments, and the "before you buy another line" comparison, sit at the top of every build's priority list.

Check every moment against evidence

A derived moment is still a guess until something backs it. Give every cell a source, and one of three states:

- **Sourced:** at least one record in the buyer's own words. A sales call, a lost deal, a tender, a client debrief, a review, a forum thread, a question put to an assistant.
- **Inferred:** it follows from a framework or from the offer, but nobody has said it yet. Use it, flagged.
- **Hypothesis:** plausible, with no record. It can't drive spending until an interview confirms it.

Kestrel's evidence (invented for the example): notes from forty sales calls over the last year, twelve lost-deal interviews, seven tenders and briefs, nine client debriefs, the question logs from the first assistant scan, competitor reviews, practitioner forums and the investor press.

Evidence	Cells it backs	State
Sales call notes	17	Sourced
Lost-deal interviews	12	Sourced
Client debriefs	8	Sourced
Tenders and briefs	7	Sourced
Assistant question logs	6	Sourced
Forums (with sales calls)	1	Sourced
The job map	5	Inferred
Investor press	3	Inferred
The difference	1	Inferred
Nothing yet	4	Hypothesis

51 of the 64 live cells are sourced, 9 are inferred and 4 are hypotheses. By moments, that's 136 of 149 sourced.

The gaps aren't random. The inferred cells cluster in the investor's row and in the preparation steps of the job map. The four hypotheses are the investor's requirements and checking, the improvement lead's checking, and procurement's repeat buy. So the interview plan writes itself: three operating partners, three improvement leads and two procurement leads. Keep going until new conversations stop adding needs. In one well-known study that happened within twelve interviews, with the basics clear within six ([Guest, Bunce and Johnson, 2006](#)). Treat that as a guide, not a rule: their interviewees were similar to each other, and a buying group isn't.

The evidence also corrects the words. A cell is written in the buyer's language from its best source, never in the firm's. "Operational excellence" appears in Kestrel's old brochure and in none of the forty sales calls.

Check the map is whole

Evidence checks each moment. Five questions check the map as a whole, the same way the coverage test checks a sitemap.

Question	Kestrel's answer
1. Was every chooser considered at every step, with any absence stated?	Yes. 80 cells considered: 64 live, 16 absent, each with a reason (below)
2. Does every offer have its job mapped to the end, past the purchase?	Yes. All eight job steps for the review land in cells; monitor, modify and conclude cover the programme
3. Does every alternative, doing nothing included, have its switching moments?	Yes. Twenty switching needs, five alternatives, all landed in cells
4. Does every part of the difference decide at least one moment?	Yes. All six parts
5. Does every stage, Known to Kept, have moments?	Yes. Kept has eleven moments across delivery and renewal, up from ten, and they now follow the job past the purchase

Every buying mode can also be walked from start to finish: committee, owner-led, investor-led, urgent and repeat.

The sixteen absent cells, and why. A blank cell is a decision, not a gap, so each one says why it's blank:

Chooser	Absent at	Why
Owner or managing director	Framing, requirements, delivery	Leaves diagnosis and scope to the COO; not involved day to day
Plant manager	Options	The COO and finance director weigh the options; the plant manager's concerns arrive at comparing
Improvement lead	Trigger	Rarely the one who starts it; often the reason it wasn't started sooner
Procurement	Trigger, framing, options, building the case, delivery	Joins when there's a brief to write, and leaves after the contract
Investor	Framing, delivery	Delegates diagnosis to the COO, and delivery to the plant

Chooser	Absent at	Why
AI assistant	Trigger, buying, delivery, extend and refer	Not asked to start a project, sign a contract or run one; reputation checks moved to checking

If any of those reasons turns out to be wrong in an interview, the cell goes live.

Kestrel's map, re-derived

Here is the result: ten journey steps, eight choosers, 64 live cells. Each cell holds the chooser's main question at that step. A mark shows the question changes with the problem (P), the sector (S) or the urgent mode (U). Cells marked *inferred* or *hypothesis* are waiting on evidence; **new** marks a cell the first draft missed.

The decision makers

Step	Chief operating officer	Owner or managing director	Finance director	Plant manager
1. Trigger	"Output's flat and a key customer is complaining. What's going on?" (P)	"Margins are slipping. Is it the plant?" (P)	"Margins fell again. What's driving the cost?" (P)	"We're always firefighting, and nobody fixes the cause." (P)
2. Framing	"Is it people, process, equipment or planning?" (P)	—	"Where's the money going: labour, scrap, stock or energy?" (P)	"Will they just tell us what we already know?"
3. Options	"Fix it ourselves, buy equipment or bring someone in?" (P)	"Do we need outside help at all?"	"What does each option cost, and what does it return?" (P)	—

Step	Chief operating officer	Owner or managing director	Finance director	Plant manager
4. Requirements	"What should the review cover, and what does a good one look like?" (P) new	—	"What do we judge them on: fixed price, guarantee, references?" new	"What do we need from them on the floor, and what can't we spare?" <i>inferred</i> new
5. Comparing	"Who does this for plants our size, in our sector?" (S)	"Can I trust them with my business and my people?"	"Fixed price or day rates? What if it finds nothing?"	"Will they work on the floor with my people, or from a meeting room?"
6. Checking	"Can I talk to a plant like ours that's done this?" (S) new	"What do people I trust say about them?" new	"Who confirmed those savings, and how?" new	"Can I see their people on a floor like mine?" <i>inferred</i> new
7. Building the case	"How do I get the finance director and the board to back this?"	"Is this the best use of the money this year?"	"Will the numbers stand up at the board?" (P)	"How much of my team's time will it take?"
8. Buying	"How soon can they start, and what do they need from us?" (U)	"Who signs, and what are we committing to?" (U)	"What's in the contract, and how do we stop?"	"What do we need ready before they arrive?" <i>inferred</i> new
9. Delivery	"Are we on track, and what are they finding?"	—	"Are the savings real, and are they being counted properly?" <i>inferred</i> new	"What do my people need to do this week?"
10. Extend and refer	"Did it pay back? Do we go into the programme?"	"Was it worth it? Would I tell another owner?"	"Are the savings showing in the accounts?"	"Results are slipping. What do we do now?" new

The checkers and the machine

Step	Improvement lead	Procurement	Investor	AI assistant
1. Trigger	—	—	"The value-creation plan needs margin within 18 months."	—
2. Framing	"We've tried lean before. Why didn't it hold?"	—	—	"Why is my factory's output flat?" (P)
3. Options	"Can we do this ourselves?"	—	"Which firms deliver operations value quickly?" <i>inferred</i>	"How do I improve productivity in a factory? What are the options?" (P)
4. Requirements	"What will they need from our data and our team?" <i>inferred</i> new	"What goes in the brief, and how do we compare quotes fairly?" new	"What scope fits the 100-day plan?" <i>hypothesis</i> new	"What should I look for in an operations consultant?" new
5. Comparing	"Will they work with us, or replace us?"	"Are they on our approved supplier list, or can they be?"	"Have they done this in portfolio companies like ours?" <i>inferred</i>	"Best operations consultants for mid-sized manufacturers?" (S)
6. Checking	"Have they worked alongside an in-house team before?" <i>hypothesis</i> new	"Insurance, safety record, financial standing?" new	"Will other portfolio companies vouch for them?" <i>hypothesis</i> new	"Is Kestrel Advisory any good? Reviews and results?"

Step	Improvement lead	Procurement	Investor	AI assistant
7. Building the case	"Will this make my team look bad?"	—	"Will the board pack show the value?" <i>inferred</i>	"What does an operations review cost, and what does it return?"
8. Buying	"What data do they need from us before week one?" <i>inferred new</i>	"Site inductions, payment terms and data handling?" (U)	"Can they start inside the first 100 days?" (U)	—
9. Delivery	"Will we own the methods when they leave?"	—	—	—
10. Extend and refer	"Can we run the next review ourselves?"	"Can we re-engage them without a new tender?" <i>hypothesis new</i>	"Can we use them across the portfolio?"	—

What changed from the brainstormed draft

	Brainstormed	Derived
Journey steps	8	10: requirements and checking added, from the buying jobs
Live cells	46	64
Moments	122	149
Buying modes	4	5: repeat buying added, from the buy classes
Cells with a stated source	0	64: 51 sourced, 9 inferred, 4 hypotheses
Absent cells with a reason	0	16
Switching needs, by alternative	Scattered	20, all landed

	Brainstormed	Derived
Triggers vs occasions	Mixed	Separated: occasions go to the campaign calendar

Most of the new moments sit where nobody in sales or marketing usually looks: procurement's brief, the reference call, the week before the team arrives, and the months after the review. Those are the moments a brainstorm can't see, because nobody in the room was there.

Then score it

[Part two of this step](#) is the Selectability Score. It takes the map and asks, moment by moment, whether each of the seven stages holds: is Kestrel known, found, answered, believed, preferred, bought and kept, at this moment, for this chooser? The questions it asks of assistants and search come straight from the cells, in the buyers' words. A thin map gives a thin score. Kestrel's derived map gives the Score 149 moments to test instead of 122, including the reference call and the procurement brief, which the first draft never measured.

The Score's results then feed the GTM Strategy. Of the three tests for choosing which moments to win (worth winning, a true reason there, still open), "still open" can only be answered once the Score shows who is named today. Hypotheses are scored like any other moment but flagged, and no budget follows them until an interview confirms them.

One rule for the whole method: no moment without a reason, and no moment without a source. It's the same rule as the sitemap's, one level up.

Feeling a bit overwhelmed?

You should be, a little. That was one firm: six roles in the buying group, ten journey steps, five buying modes, sixteen job questions, twenty switching needs, six parts of a difference, sixty-four cells, each with a source, sixteen absences with reasons, and an interview plan for what's left. Every build, page and score after this depends on it.

RocMode does the heavy part for you:

- **Your buyers and their moments, mapped from what you already have.** RocMode reads your website, documents, profiles and records, maps your buyers across their journey, and fills each moment with its pain, gain, job to be done, objection and top question. You can change any of it, and your changes are kept.

- **Their questions, asked.** It turns each moment into the questions those buyers ask, and puts them to leading AI assistants, so the Score measures your moments rather than generic keywords.
- **Deriving every moment from your industry, offer, alternatives and difference,** with a source on each one and the completeness checks run before scoring, is what we're building next.

[Join the RocMode beta](#)

Where You Stand, Part Two: Score the Moments

What the Score is for

The map says where buyers choose. The Score says how often they choose you there, and when they don't, who they choose instead.

Part one of step five found every moment that matters: for Kestrel, 149 of them, across eight choosers and ten journey steps. Part two measures each one. It gives one number for how likely Kestrel is to be selected, and one list of what to fix first. Everything in the plans that follow reads from it: the GTM Strategy's test of which moments are still open, the website's priorities, each build's own slice of the Score, and the Loop that checks whether anything moved.

The method is the one set out in Chapter 30 of *Seen Isn't Selected*. This playbook runs it end to end on Kestrel, so you can see every number come from somewhere.

Three things it is not:

- **Not a count of sales.** It's an index of your odds at the moments that matter. Nothing yet ties it to share or revenue; that test is under way, and every version of the method is published.
- **Not a visibility number.** Being named is the gate, not the score. A brand that's named everywhere and recommended nowhere scores low.
- **Not a report card for the marketing team.** It belongs to whoever owns the choice: the person who decides which fix goes first, across marketing and sales.

What it reads from the chain

The Score invents nothing. Every input was decided earlier, and it reads each one by name.

Input	From	What the Score does with it
The moments	Step 5, part one: Map the moments	Scores every one of the 149, and only those. A moment that isn't on the map is left out, never counted as zero
The questions	The map: each cell's question, in the buyer's words	Asks them, in several phrasings, wherever that chooser asks

Input	From	What the Score does with it
The buying modes	The map: planned committee, owner-led, urgent	Sets the quality weights for each moment
The claims	Step 2: Credentials, signed claims only	Checks every answer that names Kestrel against them: right, wrong or missing
The reason	Step 3: Brand Strategy, the reason to be chosen	Checks whether Kestrel is recommended for its own reason, or someone else's
The rivals	Step 3: the alternatives, and who buyers name	Records who is named instead, moment by moment
The facts	Step 1: Get the facts, the source store	Checks what assistants say about Kestrel against what's true

Two rules follow. **If an input changes, the Score is run again**, because a new claim or a new moment changes what is being measured. And **a change of method resets the baseline**: when a question, a weight or the formula changes, the new number is a new series, never reported as progress.

The instrument: hear every moment where it happens

A moment is heard where the buyer actually asks. For Kestrel, that's three places, and each needs its own instrument.

Where the moment happens	Kestrel's moments	How it's heard	What one reading is
Asked of machines. The chooser types the question into an assistant or a search box	61	The 101 questions from those cells, on five assistants and one search engine, three runs each	1,818 answers a scan. The same scan feeds Being named ; it's run once, read twice

Where the moment happens	Kestrel's moments	How it's heard	What one reading is
Asked of people and places. A peer, a forum, an association, an adviser's list, a tender brief	38	Each moment's question in three phrasings on page one of search, plus a sample of the places that chooser asks: 40 threads, directories and adviser lists	114 searches and 40 places a quarter
Inside a deal. The buyer is already talking to Kestrel: comparing, checking, building the case, buying, working together, renewing	50	The deal record and the assets each person was given	The last twelve months: 31 fit calls, 19 proposals, 9 wins; 7 finished engagements, 5 extended

In every answer, five things are recorded: **whether Kestrel is named**, and where in the answer; **who is named instead**; **what is said** about each, checked against the signed claims; **the reason given** for the recommendation, if any; and **the sources cited**. A single answer means little, because assistants answer differently each time. Shares across many answers are what count.

Two more readings set the gate's first half, Known:

- **A recognition probe.** Each assistant is asked what Kestrel Advisory does, and for which problems. Is it recognised at all, and is it tied to the moments on the map?
- **Branded search share.** Searches for Kestrel's name against searches for the rivals buyers name instead.

One rule protects the whole instrument: **questions that name Kestrel never count toward Found**. Being named when someone asked for you by name proves nothing. Where a moment's questions all name Kestrel, as in an assistant's "Is Kestrel Advisory any good?", Found isn't measured, and the gate is Known alone.

Scoring a moment

Six checks

Each moment gets six checks, one per stage, each worked as a share from 0 to 100. What counts as evidence depends on where the moment happens.

Stage	The question	Asked of machines	Asked of people and places	Inside a deal
Known	Does the buyer, or the machine, already know Kestrel?	The recognition probe: named, and tied to this moment	Branded search share against the rivals	Open: they're talking to Kestrel
Found	Is Kestrel there when they ask without naming it?	Share of unbranded answers that name Kestrel	Share of searches and places that name Kestrel	Open
Answered	Does what Kestrel publishes answer this buyer's question?	Share of the moment's questions the site answers	The same	Did the asset this person was given answer their question?
Believed	Is what's said accurate, and backed by others?	Accuracy against signed claims, outside support, consistency, in answers that name Kestrel	The same, in the places that name Kestrel	Are the claims signed, and confirmed by clients?
Preferred	When Kestrel appears, is it the one recommended, for its own reason?	Strength of recommendation, and whether the reason given is Kestrel's	The same	Share of deals Kestrel kept at this step
Bought	Does the place they land make yes easy?	Does the landing page repeat the reason, show proof and offer a next step?	The same	Did this step run without a stall: terms, start date, paperwork?

Kept isn't a seventh check. It's the same six checks at the moments after the work starts: Kestrel's delivery and extend-and-refer moments.

The formula

Known and Found are gates. Nothing else is seen until they open. The other four are the quality of what the buyer finds.

Gate = square root of (Known x Found)
 Quality = weighted average of Answered, Believed, Preferred and Bought
 Moment = Gate x Quality / 100
 Score = weighted average of every moment, 0 to 100

A check that wasn't measured is shown as not measured and left out of the average. It is never guessed. If Kestrel isn't named at a moment, there are no answers to judge for Believed or Preferred, so those two drop out and quality is the average of what's left.

The weights

The quality weights follow the buying mode. Kestrel's map uses three of the five:

Mode	Kestrel's moments	Answered	Believed	Preferred	Bought
Committee (planned)	Every chooser except the owner	25%	30%	30%	15%
Owner-led business	The owner's 13 planned moments	20%	35%	30%	15%
Quick choice	The 4 urgent moments	15%	25%	15%	45%

The moment weights favour moments nearer the decision, and the choosers who decide. Steps run from 0.6 for a trigger to 1.0 from comparing to buying; the COO, owner and finance director count 1.0, the assistant 0.9, the plant manager 0.8, the improvement lead 0.7, procurement and the investor 0.6. These are judgement, version 1.0, published and open to challenge.

Three moments, worked

Moment	Gate	Quality	Moment score
COO · framing · flat output. Asked of machines: 2 questions, 36 answers; Kestrel named in 3, all from an old blog post still on page one of search	Known 14, Found 8: $\sqrt{14 \times 8} = 11$	Answered 5, Believed 35, Preferred 25, Bought 15, committee weights: 21	$11 \times 21 \div 100 = 2$
Owner · options: "Do we need outside help at all?" Asked of people: named in one association directory	Known 12, Found 10: 11	Answered 5, Believed 40, Preferred 35, Bought 20, owner-led weights: 29	$11 \times 29 \div 100 = 3$

Moment	Gate	Quality	Moment score
Finance director · building the case · flat output. Inside a deal	Open: 100	Answered 25, Believed 30, Preferred 47 (9 of 19 proposals won), Bought 40: 35	$100 \times 35 \div 100 = 35$

The multiplication is the point. The first two moments have half-decent quality, but almost nobody reaches it. The third is reached by every buyer who gets that far, and they're the only ones who do.


Kestrel's Score, worked in full

Kestrel scores 18 out of 100, with a range of 12 to 27. Rarely selected.

Kestrel scores 18 out of 100: almost all of it inside its own deals

Moment scores by chooser and step. 99 moments happen before a buyer contacts Kestrel, and almost all of them score 0.

The 50 moments inside a deal score between 27 and 81.



Kestrel's first full scan · 149 moments · 1,818 answers, 114 searches, 40 places, 31 deals · method v1.0

The grid splits cleanly in two. Before a buyer contacts Kestrel, it's almost absent: of the 99 moments that happen out in the market, 93 score zero. Inside a deal, it does well: 27 to 81, and best of all once the work starts. So Kestrel isn't a weak firm. It's a good firm that almost nobody finds. Of its 31 fit calls last year, 22 came by referral; the rest of the market never got as far as asking.

Where the moment happens	Moments	Average moment score	Points of Kestrel's 18
Asked of machines	61	0.2	0.1
Asked of people and places	38	0.1	0.0
Inside a deal	50	48.5	18.3

By stage

Averages of each check, where it was measured. The number in brackets is how many moments it was measured in.

Stage	Asked of machines	Asked of people and places	Inside a deal
Known	12 (61)	9 (38)	Open
Found	0.5 (60)	0.4 (38)	Open
Answered	10 (61)	6 (38)	43 (50)
Believed	31 (4)	40 (2)	48 (50)
Preferred	22 (4)	35 (2)	54 (50)
Bought	15 (61)	15 (38)	48 (50)

Believed and Preferred are measured in only six moments outside deals, because those are the only six where anyone names Kestrel. That isn't a gap in the method. It's the finding.

Four kinds of moment

Outside its deals, every moment falls into one of four kinds, set by two measures: does the machine or the place pick Kestrel, and does Kestrel's own site answer the buyer?

	Picked	Not picked
The site answers	Won: 0	Hidden: 0
The site doesn't	Borrowed: 5	Empty: 94

No Hidden moments is unusual, and it matters. Most firms have answers nobody is sent to. Kestrel's old site, 14 pages about the firm, answers almost none of its buyers' questions. The five Borrowed moments rest on things Kestrel didn't plan: two old blog

posts that still rank in search, an old directory listing, an association page and one adviser's list. And when an assistant is asked about Kestrel by name, it mixes it up with a mining-services firm of a similar name.

Why the range is wide


Twelve to 27 is a wide range, and it should be. Almost all of Kestrel's score rests on 19 proposals. Win two more or two fewer and the number moves by several points. The top of the range crosses into "sometimes selected". The bottom doesn't come close. Report the range every time, or a lucky quarter will be read as progress.

Where Kestrel loses buyers, and to whom

The Score is out of 100, so the gap to 100 is what's being lost. Kestrel loses 82 points. Split that loss by stage, by step and by chooser, and the biggest share is the biggest lever.

Of the 82 points Kestrel loses, 46 are lost at Found

Points lost by stage, out of 100. The two gates, Known and Found, account for 62.2; the four quality stages for 19.3.



Kestrel's first full scan · each moment's lost points split between its gate and its quality stages, weighted as in the Score

Three quarters of the loss is at the gates. Kestrel isn't losing buyers who compare it and choose someone else. It's losing buyers who never see it.

By step and by chooser

The most points lost	Points	What's happening there
Options	16	Buyers weigh fixing it themselves, new equipment or outside help, and Kestrel isn't among the outside help anyone names
Trigger	15	Seven problems, five choosers, and Kestrel is absent from 28 of the 29 trigger moments
Framing	14	Assistants explain flat output and rising costs every day. None mention Kestrel's insight, that the capacity is usually already in the line
The COO	26	The initiator, with 41 moments; 32 of them happen before anyone contacts Kestrel
The finance director	20	Lost twice: absent while framing the cost, then thin on proof once inside the deal
The AI assistant	14	Six steps on the map, a score of 0.3. It shortlists for the whole committee and has no reason to name Kestrel

To whom

When buyers ask without naming anyone, this is who gets named instead. Shares are of the 1,800 unbranded answers in the scan; an answer can name more than one firm.

Named in the answer	Share of answers	Where they win
No firm at all: general advice	44%	Trigger and framing. The buyer is told what to do, not who to call
Large firm A	29%	Options and comparing, every sector
Large firm B	22%	Options and comparing; investor questions
The lean specialist	16%	Framing and the improvement lead's questions
The engineering firm	8%	Options, where new equipment is the answer

Named in the answer	Share of answers	Where they win
Kestrel	0.4%	Eight answers: six from two old blog posts that still rank in search, two from an old directory listing in packaging

The first row is the opening. At 44% of answers nobody owns the moment yet. Large firms win comparing because they're named everywhere; nobody is winning framing, because nobody has published the explanation buyers are looking for. That's the moment Kestrel's reason was written for.

What the Score hands to the plans

One list

The loss becomes a list of fixes, ranked by how much each should lift the Score for the effort it takes. Every fix names its moments and its stages, and carries a predicted range, worked out by rerunning the Score with the fix in place.

#	The fix	Moments and stages	Predicted lift	Effort	Shows up
1	Be there before they call. Problem and sector pages, the facts fixed everywhere, the explanation buyers are looking for	97 moments outside deals, trigger to comparing; Known, Found, Answered, Bought	+4 (2 to 6)	Large	Over two or three monthly scans
2	The finance director's kit. The way out in writing; savings confirmed by clients' finance directors	The finance director's 12 moments inside deals; Answered, Believed	+1.5 (0.7 to 1.8)	Small	Next quarter's deals
3	A reference within a week. One client per sector, ready to take a call	16 comparing and checking moments; Believed, Preferred	+0.9 (0.5 to 1.2)	Small	Next quarter's deals
4	A proposal page per person. Each chooser's own question, answered on one page	12 building-the-case moments; Answered, Preferred	+0.9 (0.4 to 1.1)	Medium	Next quarter's deals
5	The procurement pack. Insurance, safety, terms; onto two approved-supplier lists	Procurement's 5 moments inside deals; Answered, Believed	+0.3 (0.2 to 0.4)	Small	Next tender

Together: **from 18 to about 26, with a range of 22 to 28.** That's still "sometimes selected" at best, after a year's work. Say so plainly to whoever owns the choice. A firm that's absent from 94 moments doesn't fix that in a quarter, and a list that promises otherwise is selling something.

The order isn't the order of the list. Fixes 2 to 5 are small and show up within a quarter, so they go first, alongside the start of fix 1, which is slow and worth the most. That's exactly the year-one order in [the Plans](#).

The open test

The GTM Strategy sorts moments with three tests: worth winning, a true reason there, and **still open**. The Score answers the third. A moment is open when nobody owns it: no firm is named in most answers, or the firms named are named without a reason. For Kestrel, trigger and framing are wide open: 44% of answers name no one. Comparing isn't: the two large firms own it. So Kestrel wins first where the field is empty and its reason is strongest, and comes to comparing with references in hand.

Each build's slice

Every build reads its own part of the Score, and its own audit starts there:

- **Website:** Answered and Bought, in every moment. For Kestrel, 10 and 15 outside deals.
- **Being named:** Known and Found, at the 61 moments asked of machines. Found is 1.
- **Content and social:** Found and Answered at trigger and framing.
- **PR and reviews:** Believed, wherever Kestrel is named, and the references inside deals.
- **The pitch, events and outbound:** the quality of the 50 deal moments, and Found for the owners.
- **Email, automation and customers:** the delivery and extend moments, which are Kept.

Predicted, then measured

When a fix ships, the next scan records what moved, beside what was predicted. Predicted against measured is how the list earns trust, and how the weights get corrected. A fix that ships and moves nothing is a finding too.

Running it

The cadence

Reading	How often	Why
Asked of machines: the 101 questions, six engines, three runs	Monthly	Answers move as sources change. The same scan feeds Being named
Asked of people and places: searches and the 40 places	Quarterly	These move slowly; more often is noise
Inside a deal: the deal record and the assets	Quarterly	Kestrel closes too few deals for a monthly number to mean anything
The full Score, the range and the list	Quarterly	The number the owner of the choice reports
A rescan out of turn	When an input changes	A new claim, a new moment or a new rival changes what's being measured

Who owns it

One person owns the Score and the list, across marketing and sales, with the standing to say no to work that moves no moment. At Kestrel, that's the managing partner. They don't run the builds. They decide which fix goes first, and which team does it.

Reading AI's influence

Most buyers who act on an assistant's answer don't click through. They search Kestrel's name, or type the address. So each quarter, read three numbers together: the share of answers that name Kestrel at its moments, searches for its name, and direct visits. If the first rises and the other two follow a month or two later, that's the assistants at work. It isn't proof, so note what else changed that quarter.

The Loop

Step eleven, the Loop, starts here. Each quarter it sets the new Score beside the old one, the measured lift beside the predicted one, and asks what changed, and what the plans should change. A change of method, a new weight or a new question set starts a new series. It is never reported as progress.

Testing the weights

The weights are version 1.0: judgement, published and open to challenge. The test is whether brands that score higher hold more share, and whether a rising score comes before rising share, not after. Kestrel can contribute one line to that test: its Score each quarter beside its new clients and fees. If the two don't move together over a year or two, the weights are wrong, and they change, with a new version number.

Common mistakes

The mistake	What it does to Kestrel's number	Instead
Scoring only what machines can hear	0.2 out of 100, and no sign that Kestrel wins half its proposals	Hear every moment where it happens: machines, people and places, deals
Scoring only the deals	A 47% win rate, and a firm that looks healthy while 93 moments score zero	The same: the deal record is one instrument of three

The mistake	What it does to Kestrel's number	Instead
Counting questions that name you	"Is Kestrel Advisory any good?" names Kestrel every time, and would make Found look alive	Branded questions never count toward Found
Reading single answers	One lucky answer in packaging reads as a breakthrough	Shares across many runs, with a range on every number
Guessing what wasn't measured	Believed and Preferred are measured in only six moments outside deals. Filling the rest with an average would invent a firm that's trusted where nobody has heard of it	Show it as not measured, and leave it out
Counting activity	Publishing 20 pages would lift the number before a single buyer saw them	Only measured outcomes move the Score
Reporting a new method as progress	Adding questions or changing a weight can move the number by several points overnight	A new method starts a new series
Mistaking named for chosen	Large firm A is named in 29% of answers. Being named without a reason still loses the buyer who asks why	Preferred checks whether the reason given is yours
Giving it to marketing alone	Nearly all of Kestrel's 18 points come from the pitch and delivery, which marketing doesn't run	One owner of the choice, across marketing and sales

And the bias to watch in yourself: the **anchoring** of the first number. Kestrel's 18 will feel like a verdict on the firm. It isn't. It's a measure of how often buyers get the chance to choose a firm that, once they do, they usually like.

Feeling a bit overwhelmed?

You should be, a little. One firm: 149 moments heard in three places, 101 questions on six engines three times a month, 114 searches and 40 places a quarter, 31 deals read moment by moment, six checks per moment, three sets of weights, a range on every number, and a ranked list rerun after every fix. Most firms check a handful of questions in one assistant, call the result their AI visibility, and never connect it to the deals they win.

RocMode does the heavy part for you:

- **Your questions, from your moments.** It writes the buyer questions for every chooser at every step, from your personas, their pains, gains and jobs, and the rivals you name.
- **Asked across leading AI assistants.** Each question runs on five assistants, three times each. It records whether you're named, who's named instead, and which sources are cited.
- **Scored by stage, moment by moment.** Found, Answered, Believed, Preferred and Bought for every moment, with a range on the score and small samples held in check. Questions that name you never count toward Found.
- **The kinds of moment, and the biggest lever.** Each moment sorted into won, hidden, borrowed or empty, and the fix worth the most points, with the points it should add.

[Join the RocMode beta](#)

PART THREE

Plans

Steps 6 to 9: GTM, Website, Marketing and Sales strategies



Part Three: Plans

What the plans are for

The plans turn the map into decisions: which moments to win, what the site must hold, which builds to run, and how accounts are won and kept. They come after the Score for a reason. Until you know where buyers are lost today, and to whom, any plan is a guess about what matters.

Four plans, in order, each reading what came before it:

Step	Plan	Decides	Reads
6	GTM Strategy	Which moments to win, which segments, which buying paths, and the sales motion	The moment map and the Score
7	Website Strategy	Every page, why it exists, and in what order	The map, the moments chosen, the reason, the claims
8	Marketing Strategy	Which builds lead which moments, in what order, and with what share of the effort	All of the above
9	Sales Strategy	How accounts are won, delivered, kept and grown	The motion, the moments chosen, the signed claims and cases

Three rules hold across all four. **None redefines an audience:** they use the choosers and moments by name. **None adds a claim:** a new claim goes into Credentials first. **"Moment" means one thing:** a chooser at a step. A trade show or budget season is an occasion, and goes in the campaign calendar.

Kestrel is the worked example throughout: its strategy is in [Part one](#), and its 149 moments in [Where you stand](#).

6. GTM Strategy: which moments to win, and how to sell

In RocMode, the GTM Strategy holds the ideal customer, the segments, the win strategy, the sales motion, the messaging with its sales enablement, and how leads are covered and routed. Here it starts from the map and the Score, and makes four decisions.

Which moments to win first

Every moment worth winning gets covered; the GTM Strategy decides the order. Three tests sort them: **worth winning** (does it lead to work Kestrel wants?), **a true reason there** (does a signed claim decide it?) and **still open** (does the Score show nobody owns it yet?).

Tier	Kestrel's moments	Why
Win first	Flat output and rising costs, at the trigger, framing and options steps, for the COO, the finance director and the assistant; every comparing and checking moment by sector; the finance director's price and way-out questions; "before you buy another line"	Most of Kestrel's work starts here; the big idea and the way out decide them; the Score shows assistants name no one consistently
Win next	The other five problems; the investor's path; the improvement lead's moments; procurement's brief	Worth winning, with a reason there, but smaller or slower
Hold	The four hypotheses: the investor's scope and checking, the improvement lead's checking, procurement's repeat buy	Scored and watched, but no spend until interviews confirm them

Which segments

Mid-sized manufacturers with flat output or rising costs, in the four sectors, first. Investor-owned manufacturers second: fewer of them, but each brings a portfolio, which is why the repeat path matters.

Which buying paths

All five modes stay on the map, but the motion is built for two first: the planned committee, where most of Kestrel's work comes from, and the urgent path, where speed decides. The owner-led and investor-led paths reuse the committee's assets with a different order. The repeat path is run by Sales.

The motion and the win strategy

The sales motion is sales-led and consultative, and it's set by the firm, never guessed: a fit call, a one-page scope, a proposal the committee can forward, a reference call, the business case, the contract, week zero, the review, then the programme or the next plant.

The win strategy names how Kestrel wins against each alternative, using the switching forces from the map: **against new equipment**, find the capacity in the line first; **against a large firm**, the people who pitch do the work, at a fixed price; **against the in-house team**, they lead and keep the methods; **against doing nothing**, the way out makes acting cheap.

7. Website Strategy: what the site must hold

The website is the only build lit in every live cell: whatever a buyer hears elsewhere, they check on the site. So its plan comes first, and it's worked through in full in [the Website Playbook](#). In RocMode, the website strategy has nine sections, from the site's thesis to its page inventory, and checks every page for the moments it covers.

For Kestrel, the result is **47 pages, plus one per case study**, each traced to the moments it answers. The site leads 31 of the 149 moments and supports the other 118. Twelve pages are P1 and launch first; the rest follow in two waves.

8. Marketing Strategy: which builds, in what order

In RocMode, the marketing strategy has eight sections, from the growth thesis through channels, content pillars, campaign concepts, measures and budget to the roadmap. Campaigns are its delivery plan, not a separate strategy. Its central decision is which build leads each moment.

Six categories, fourteen builds

Every way a buyer meets a brand sits in one of six categories, sorted by one question: **who is speaking, and how did they reach the buyer?** It extends Gini Dietrich's PESO model (paid, earned, shared, owned), which was built for communications. Shared becomes partnered, and two are added: direct, because for most firms a person in a room decides more moments than any channel; and relationship, because the system that remembers a customer decides whether they stay.

Category	Who is speaking, and how	Builds	Nods
Owned	Your voice, on ground you own	Website; Content; Social	Packaging and physical
Paid	Your voice, in space you buy	Advertising	Sponsorship; retail media; broadcast and outdoor; direct mail

Category	Who is speaking, and how	Builds	Nods
Earned	Someone else's voice, freely given, including AI answers	Being named: search and AI answers; PR and reviews	Analyst relations; awards; word of mouth
Partnered	Someone else's voice or audience, by agreement	Partners and channels; Creators	Affiliates; marketplaces; associations; employee advocacy
Direct	A person, live	Outbound and account-based; The pitch; Events	Account management; customer service
Relationship	Your system, to someone you already know	Email; Automation; Customers: onboarding, loyalty, community and referral	CRM; SMS and chat

A channel earns its own build only if it has its own shape, its own craft, and moments it leads better than anything else. The rest get a nod: a short section in the nearest build's playbook. Formats such as video, podcasts and webinars live inside builds, not beside them. Offer, price, product and distribution aren't builds at all: they're foundations, set in the strategy.

Fourteen builds, fourteen footprints

A build's **footprint** is the moment map with that build's cells lit. It's worked out, not chosen. Two rules keep it true:

1. **A cell lights up only where the build is the best answer to the moment, or where buyers expect to find one.** That a build could touch the moment isn't enough.
2. **Each cell has one lead build, and any number of supporting builds.** The lead does the work at that moment. The supporters echo the same message.


Fourteen builds, one map: each build lights different moments

Rows, top to bottom: COO, owner, finance director, plant manager, improvement lead, procurement, investor, AI assistant. Columns, left to right: journey steps 1 to 10.

Leads the moment Supports another build Live cell, not this build

Blank: that chooser isn't present at that step. Each panel names its category.





Kestrel's fourteen build footprints · 64 cells, 149 moments

Build	Category	Leads	Supports	Where it leads
Website	Owned	10 cells, 31 moments	54 cells	The COO's and finance director's options, scope and comparing; the finance director's checking; procurement's brief, comparing and checking
Content	Owned	6 cells, 24 moments	11 cells	The finance director's trigger; framing for the COO, finance director and improvement lead; the improvement lead's options and next review
Social	Owned	1 cell, 7 moments	4 cells	The plant manager's trigger
Advertising	Paid	1 cell, 7 moments	2 cells	The COO's trigger, for each of the seven problems
Being named: search and AI	Earned	5 cells, 20 moments	8 cells	Every assistant question before the pitch: framing, options, scope, comparing, cost and return

Build	Category	Leads	Supports	Where it leads
PR and reviews	Earned	4 cells, 7 moments	15 cells	The reference and reputation checks for the COO, improvement lead and investor, and what assistants say when a buyer checks
Partners and channels	Partnered	3 cells, 3 moments	4 cells	The investor's trigger, options and portfolio follow-on, reached through the advisers and lenders investors already use
Creators	Partnered	1 cell, 1 moment	5 cells	The plant manager's framing
Outbound and account-based	Direct	1 cell, 7 moments	3 cells	The owner's trigger: named firms, one problem at a time
The pitch	Direct	19 cells, 28 moments	9 cells	Building the case and buying for almost everyone; the plant manager's, improvement lead's and investor's scope; comparing for the owner, plant manager, improvement lead and investor; the plant manager's site visit; procurement's renewal
Events	Direct	2 cells, 2 moments	3 cells	The owner's options and checking: round tables where owners hear from owners
Email	Relationship	3 cells, 3 moments	14 cells	Delivery for the COO and finance director; the finance director's follow-on

Build	Category	Leads	Supports	Where it leads
Automation	Relationship	4 cells, 5 moments	9 cells	Procurement's paperwork; the week-zero preparation for the plant manager and improvement lead; the plant manager's week-by-week needs
Customers	Relationship	4 cells, 4 moments	4 cells	The improvement lead's handover; the follow-on for the COO, owner and plant manager: the results review, the client community and the referral

The leads add up to 64 cells and 149 moments: every moment has exactly one lead.

What Kestrel's footprints show: **the website is the hub**, lit everywhere but leading only a fifth of the moments; **the pitch leads the most cells**, because a committee builds its case and buys in meetings, not on a screen; **PR and reviews support far more than they lead**, because a reference or a review backs almost every claim made elsewhere; **the four newer builds each lead a few moments nothing else wins as well**: events for owners, partners for investors, outbound for named firms, and customers after the work is done; and **advertising, social and creators lead very little**. That's right for a consultancy sold to committees. A consumer brand's footprints would look almost reversed. Footprints are worked out for each firm, never copied from a template.

The check across all fourteen builds

- **No moment without a lead.** Every one of the 149 moments has exactly one lead build.
- **No build without a lead.** Every build leads at least one moment. A build that only supports is a cost without a job. For Kestrel, creators come closest, with one.
- **One message per crowded cell.** The COO's trigger is the most crowded: advertising leads, and social, PR, content, the website, events and outbound support. All seven say the problem in the buyer's words and give the same reason.
- **Hand-offs, not dead ends.** The ad sends the COO to a problem page; the page offers the business case kit; the kit leads to the pitch; the pitch offers a reference call; the reference call is arranged by PR and reviews; the results review, run by Customers, asks for the next plant and the referral.

In what order

The order follows the GTM tiers and what each build needs from the others. For Kestrel's first year:

Quarter	Builds	Why then
One	Website wave 1; being named (the facts fixed everywhere, the sector sentence); the pitch kit	Everything else sends buyers to the site and the pitch; wrong facts get repeated until they're fixed
Two	Content P1 guides and the benchmark; PR and reviews (references, first reviews); email for delivery and renewal; customers (handover and the results review)	The framing moments and the checking step, once there's a site to land on; every current client becomes a reference
Three	Website wave 2; advertising for the COO's trigger; social; creators; outbound to named firms; partners and channels for investors	Paid, outbound and partner attention only once the pages and proof can hold it
Four	Website wave 3; events (the first owners' round table); automation; the rest of the content plan	Events need clients willing to speak; the routines keep the rest running

9. Sales Strategy: how accounts are won and kept

In RocMode, the sales strategy lives in the GTM record: the win strategy, the messaging with its sales enablement (objections with proof, assets, kits for each person in the buying group), and how leads are covered and routed. Here it covers the whole account, from the first call to the referral, because no other plan owns what happens after the signature.

The motion, step by step

Motion step	The moments it serves	What Kestrel brings
Fit call	COO · comparing and buying; owner · comparing	Who it's for and who it isn't for; the six weeks; the price range
One-page scope	COO, plant manager, improvement lead and investor · scope	What the review covers in this plant, what it needs from the team, what it won't touch

Motion step	The moments it serves	What Kestrel brings
Committee proposal	Everyone · building the case	One page per person in the room, each answering their own question; forwardable without the pitch
Reference call	COO, owner, improvement lead and investor · checking	A client in the same sector, arranged within a week
Business case	COO and finance director · building the case	The kit, filled in with their numbers; the way out stated
Contract	Finance director, owner and procurement · buying	Terms that match the website word for word
Week zero	Plant manager and improvement lead · buying	The preparation checklist and the data request, sent automatically
The review	COO, plant manager and improvement lead · delivery	Weekly findings in the plant manager's words
Results and renewal	COO, finance director, plant manager and procurement · extend	Savings confirmed with the client's finance team; the programme or the next plant; re-engagement without a new tender
Referral	Owner and investor · extend	A reference the client is glad to give; the portfolio conversation

Objections, each with its proof

Chooser	The objection	The proof that answers it
Finance director	"Consultants always overstate savings."	Savings confirmed by clients' finance directors, months later; the way out
Owner	"Will the board trust a smaller name?"	References from owners of similar firms; the people who'll be on site
Plant manager	"They'll tell us what we already know."	A sample review report, in plant language; a visit to a floor like theirs
Improvement lead	"This is a verdict on my work."	Cases where the in-house team led the changes and kept the methods

Chooser	The objection	The proof that answers it
COO	"We were about to buy a new line."	The capacity found in lines like theirs, before the purchase
Procurement	"We can't compare a fixed price with day rates."	A brief template with the criteria that matter, and the terms in writing
Investor	"Can they move inside our hundred days?"	Start dates and portfolio cases

Every proof must be a signed claim in Credentials. If an objection has no signed proof, that's a gap in Credentials, not a sales problem.

Routing by buying path

- **Planned committee:** the full motion, with a proposal page for each person.
- **Owner-led:** fit call, a site visit and the contract, often in two meetings; references from other owners first.
- **Investor-led:** the operating partner first; the business case doubles as the board pack.
- **Urgent:** fit call to contract in days; the start date and week zero lead.
- **Repeat:** the account owner, not a new sale; procurement re-engages without a tender where the contract allows it.

What comes next

The plans say what to do. Part four does it: the fourteen build playbooks, each starting from its footprint ([the website](#), [being named](#), [content](#), [the pitch](#), [email](#), [automation](#), [PR and reviews](#), [customers](#), [partners and channels](#), [events](#), [outbound](#), [advertising](#), [social](#) and [creators](#)), with the nods inside them; then [the Campaign Engine](#), which turns the builds into twelve months of campaigns around moments and occasions; then [the Loop](#), which rescores and changes the plan when the numbers move.

Feeling a bit overwhelmed?

You should be, a little. Four plans: three tiers of moments, two segments, five buying paths, a sales motion in ten steps, seven objections each with signed proof, 47 pages, fourteen footprints in six categories that must agree on 149 moments, and a year's build order. Most firms write the marketing plan, the website brief and the sales playbook separately, in different years, by different people. That's how a website ends up selling one firm and a pitch another.

RocMode does the heavy part for you:

- **Your GTM strategy.** Your ideal customer, segments, win strategy, sales motion, messaging and how leads are covered and routed, all read from your strategy and your buyers' moments.
- **Your website and marketing strategy.** A website strategy with a reason for every page, and a marketing strategy from the growth thesis to the roadmap, with campaign concepts that become campaigns.
- **Your sales strategy.** Objections with their proof, kits for each person in the buying group, and routing, from the same moments as everything else.

[Join the RocMode beta](#)

PART FOUR

Builds, and running them

Fourteen builds in six categories, then steps 10 and 11



The Website Playbook

The idea

A website has two halves: a shape the category expects, and coverage of every moment that matters. Most sites have only the first. They take the standard sitemap for their category and fill it with copy, so the pages look right and say little. A purely question-led site has the opposite problem: every answer is there, but buyers can't find the structure they expect, and the brand doesn't hang together.

So the method runs both halves and reconciles them:


1. **Shape, top-down:** the page types buyers in this category expect, plus the brand's story spine.
2. **Coverage, bottom-up:** the moments, the stage gaps, the biases and the jobs. This is the new part.
3. **Reconcile:** every job gets a home in the shape, every page in the shape earns at least one job, and jobs with no home become new pages. That's usually where the differentiation shows.

Then each page gets a dual brief, the business check and its measures.

The shape matters to being chosen, not only to tidiness. "Familiar feels safe" (Chapter 8): when a buyer can't find the case studies or the pricing they expect, they start to doubt you before they've read a word.

The flow at a glance

Shape and coverage meet in the reconcile step




shape, coverage, reconcile · the website method

The shape and the coverage are worked out separately, then meet in the reconcile step. Everything after that applies to the reconciled site. Each quarter, coverage runs again, because moments change faster than the shape.

The website's footprint, and Kestrel in 60 seconds

The website's footprint: lit in every live cell, leading 31 of 149 moments

Kestrel's moment map. Rows: choosers. Columns: journey steps. Blank: that chooser isn't present at that step.



The website's footprint on Kestrel's grid

The website is the only build lit in every live cell. Whatever a buyer hears from an ad, a pitch or an assistant, they check it on the site. So the site has to answer all 149 moments: it leads 31 and supports the other 118.

Kestrel in 60 seconds. Kestrel Advisory is the invented firm used in every playbook: a 60-person operations consultancy that helps mid-sized manufacturers find waste and lift output. Its main offer is a fixed-price, six-week operations review, which can lead into an improvement programme. Its strategy is in [Part one](#), and its full moment map in [Where you stand, part one](#).

- **Who chooses:** usually a committee. The COO starts it, the finance director approves the spend, the plant manager lives with it and procurement runs the paperwork. The owner, the in-house improvement lead, an investor and AI assistants all have a say.
- **The reason to be chosen:** "We find the savings on the floor in six weeks, at a fixed price, and the review pays for itself or you don't proceed." (Invented. A real firm would need proof for every part of that sentence.)
- **The alternatives:** a large consulting firm, a specialist lean consultant, the in-house improvement team, buying new equipment, or doing nothing. Buying new equipment matters most: capital spending competes for the same budget.

Its slice of the Score

The website's own part of [the Selectability Score](#) is two stages, in every moment:

Answered (does the site answer this buyer's question?) and **Bought** (does the page they land on repeat the reason, show proof and make the next step easy?). It also feeds Found, because assistants and search can only send buyers to pages that exist.

Kestrel today	Answered	Bought	Moment score
The 31 moments the site leads	15	18	4.5 on average; 26 of the 31 score zero
Every moment outside a deal (99)	9	15	Almost all zero, because the gate is shut
Points the Score loses at these two stages	5.1	3.4	8.5 of the 82 lost

Read the last row carefully. The site's own stages cost Kestrel only 8.5 points, because most buyers never reach the site to be disappointed by it. The site's bigger effect is upstream: with no page answering the problem, sector or comparison questions, there's nothing for an assistant or a search engine to find. That's why the Score's first fix, "be there before they call", is mostly this build, shared with [Being named](#): a predicted 2 to 6 points, slowly.

The site's own audit, in step 3 below, starts from the same numbers: 14 pages, and 84 of the 149 moments with no page at all.

Step 1: the shape

Start with the page types buyers in this category expect to find. For B2B professional services sold to a committee, that's the list below. Each one exists because it answers a moment that comes up in almost every deal, which is why buyers have learned to look for it.

Page type	Why buyers expect it	Stage it mainly serves
Services	"What exactly do you do, and what do I get?"	Answered
Solutions (by problem)	"Do you solve the problem I have?"	Found, Answered

Page type	Why buyers expect it	Stage it mainly serves
Industries	"Do you understand businesses like mine?"	Believed, Preferred
Roles (who it's for)	"What does this mean for me, in my job?"	Answered, per person on the committee
Case studies	"Have you done this for someone like us, with what result?"	Believed
Approach or method	"How will you work, and what will it be like?"	Answered, Bought
Team and leadership	"Who will actually do the work?"	Believed
Insights	"Do you know more about this than we do?"	Known, Believed
About	"Are you stable, and are you a fit for us?"	Believed
Contact	"How do I start?"	Bought

The story spine runs through the shape. Before any page is written, set the narrative every page repeats in its own way:

- **The problem, in the buyer's words:** "Output is flat and costs are rising, and we can't see where the waste is."
- **The reason:** savings found on the floor in six weeks, at a fixed price.
- **The proof pattern:** every claim carries a named plant, a before-and-after number and a quote from the plant manager.
- **The way out:** if the review doesn't find more than it costs, the client doesn't proceed.

The navigation and the home page tell this story in order. Every page type echoes it.

Step 2: coverage, the full grid

Coverage comes from Kestrel's moment map: 8 choosers, 7 problems, 10 journey steps and 5 buying modes, giving 64 live cells and 149 moments, each traced to its source.

[Where you stand, part one](#) derives it in full: every chooser's question at every step, and what sits behind each one. Because the website is lit in every cell, this playbook works through all of it. Here is what the map asks of the site.

Chooser	What they need from the site
Chief operating officer	A diagnosis they recognise, and a case they can take upstairs
Owner or managing director	Trust in the firm, its people and its judgement
Finance director	Price, terms, the downside, and numbers that survive the board
Plant manager	Proof the work happens on the floor, and what it costs the team in time
Continuous improvement lead	To be worked with, not around, and to keep the methods afterwards
Procurement, with site safety	Supplier approval, insurance, inductions and payment terms
Private equity operating partner	Speed, portfolio experience and reporting the board can use
AI assistant	Quotable facts: who it's for, what it costs, what it returns, and who says so

Problem	What changes on the site
Flat output	Throughput proof, equipment effectiveness, scheduling
Rising costs	Cost-out proof, and where the money goes
Cash tied up in stock	Stock and lead-time proof, in the finance director's language
Scrap and rework	Root-cause proof, not more inspection
A new line won't reach speed	Ramp-up proof, and working alongside the equipment supplier
After an acquisition	Speed, portfolio experience, board reporting
Automation not paying back	Proof that fixing the process comes before the technology

Buying mode	What the site must do
Planned committee	Make every page work when forwarded; arm the champion for step 5, building the case
Owner-led	Trust first: the people, the cases, plain terms
Investor-led	Portfolio proof, speed to start, board-ready reporting
Urgent	Start dates, what's needed from the client, a fast way to talk
Repeat	Make the next review easy: what happens after the first, how to extend to another plant, and how procurement re-engages without a new tender

One question per cell isn't enough to brief a page from. Each cell also records a pain, a gain, a job to be done, an objection and what would move the buyer. The section briefs in Step 5 draw on those fields.

Step 3: gaps, biases and jobs

Score every moment on the three tests: worth winning, a true reason there, and open. The scores set the order of work. They don't decide which moments get a page. A moment worth winning that has no page is a buyer who gets their answer somewhere else.

Then audit the current site against all 149. Kestrel's old site (invented) had 14 pages and a twelve-post blog: home, services, approach, four generic solutions pages, one industries page, a wall of client logos, team, about, contact, careers and privacy. **84 of the 149 moments had no page at all. The other 65 landed on a page that touched the subject without answering the question.**

Step	Moments	On the old site	The stage that fails	The bias at work
1. Trigger	29	None answered	Known: nothing in the buyer's words for any of the seven problems	Familiar feels safe: buyers turn to names they've heard
2. Framing	23	7 touched, by blog posts	Found: assistants explain the problem and name other firms	Good enough beats best

Step	Moments	On the old site	The stage that fails	The bias at work
3. Options	24	14 touched, by services and approach	Answered: no comparison with buying equipment or doing it in-house	A hard choice stalls
4. Requirements	13	10 touched, by services and approach	Answered: nothing on what a good review covers, or how to brief and compare one	A hard choice stalls
5. Comparing	14	12 touched	Preferred and Believed: "we help organisations transform", and logos without numbers	Feeling first, reasons after; we follow others
6. Checking	11	7 touched, by the logos and the team page	Believed: no references to call, no one confirming the savings	We follow others
7. Building the case	13	9 touched	Believed: nothing the champion can forward or the board can check	Losses loom larger
8. Buying	11	4 touched	Bought: price "on application", terms only in the proposal, nothing on preparing for week one	Losses loom larger
9. Delivery	4	2 touched	Kept: nothing on what the plant's people do in week one, or how savings are counted	A hard choice stalls
10. Extend and refer	7	None answered	Kept: no results across reviews, nothing on keeping the gains or the next review	We follow others

Each gap becomes a job: a stage, a moment and a bias. Kestrel's grid produces dozens. These sixteen carry the most weight:

Moment	Stage	Bias	The job
COO · trigger · flat output: "We're running flat out and still shipping late."	Known	Familiar feels safe	Own the problem in the buyer's words: a problem page, and a benchmark report others quote
Assistant · framing: "Why is my factory's output flat?"	Found	Good enough beats best	Answer in quotable passages with sourced figures, so the assistant cites Kestrel
COO · options: "Fix it ourselves, buy equipment or bring someone in?"	Answered	A hard choice stalls	Compare all three fairly, including when buying the machine is the right answer
Finance director · options: "What does each option cost and return?"	Answered	Losses loom larger	Let them use their own numbers: a savings estimator with its assumptions shown
COO · comparing · by sector: "Who does this for plants our size?"	Preferred	Feeling first, reasons after	Say who it's for and who it isn't for, by size and sector, in words the COO can repeat
Assistant · comparing · by sector	Found	Good enough beats best	A page per sector stating size, sector and offer in quotable words; get named in the sources assistants cite
Owner · comparing: "Can I trust them with my business and my people?"	Believed	Familiar feels safe	Show the people who'll be on site, by name, with their plant experience
Finance director · comparing: "Fixed price or day rates? What if it finds nothing?"	Bought	Losses loom larger	Price range, scope and the way out, side by side
Plant manager · comparing: "On the floor, or in a meeting room?"	Answered	A hard choice stalls	The six weeks, day by day: who's on site and what they need

Moment	Stage	Bias	The job
Improvement lead · comparing: "Will they work with us, or replace us?"	Preferred	Losses loom larger	Show the in-house team leading the changes and keeping the methods
Finance director · building the case · by problem: "Will the numbers stand up?"	Believed	We follow others	Cases filtered by problem, with savings confirmed by the client's finance team
COO · building the case: "How do I get the board to back this?"	Believed	A hard choice stalls	A business case and board paper kit the champion can send as it is
Investor · comparing: "Have they done this in companies like ours?"	Believed	We follow others	A page for investors: portfolio cases, speed to start, board reporting
Procurement · buying: "Insurance, site safety, payment terms?"	Bought	A hard choice stalls	Everything procurement and site safety ask for, on one page they can forward
Owner · buying · urgent: "Who signs, and how fast can you start?"	Bought	Losses loom larger	Start dates and the first week's needs, so urgency doesn't go to whoever answers first
Finance director · extend: "Are the savings showing in the accounts?"	Kept	We follow others	Results across every review, and how gains are kept, so renewals and referrals have proof
Procurement · requirements: "What goes in the brief, and how do we compare quotes fairly?"	Answered	A hard choice stalls	Give procurement a brief template and the criteria that matter, so Kestrel is compared on what it does best
COO · checking · by sector: "Can I talk to a plant like ours that's done this?"	Believed	We follow others	A references page: a client to call in each sector, with their permission, and what they'll say
Finance director · checking: "Who confirmed those savings, and how?"	Believed	Losses loom larger	Say on the results page who confirmed each figure, and how it was counted

Moment	Stage	Bias	The job
Plant manager · buying: "What do we need ready before they arrive?"	Bought	A hard choice stalls	A week-zero checklist on "The first six weeks", so the start doesn't slip

Every job names a person, a stage and a bias. That's what lets a writer, or an AI, produce the right block rather than generic copy.

Step 4: reconcile

Put every moment into the shape. Go through the shape page type by page type, then find homes for what's left.

Page type in the shape	What coverage did to it	Result
Services	Split: the review, its price and terms, and the programme carry different moments	Three pages, plus "What if the review finds nothing?" and "Can you start this month?"
Solutions, by problem	Replaced: four generic pages out, seven real problems in	A hub and seven problem pages
Industries	Expanded: the assistant's shortlist question changes by sector	A manufacturing hub and four sector pages
Roles	Expanded from three to six: owners, improvement teams and investors have their own questions	Six role pages; procurement gets its own page under "Build the case"
Case studies	Rebuilt as a hub that filters by problem, sector, size and mode	A hub, a page per case, results across every review, and a references page
Approach	Kept, renamed	"The first six weeks"
Team	Kept, made specific	"The team on your floor": the people who'll be on site
Insights	Rebuilt around the trigger and framing steps	A guide, a benchmark report, questions to ask, and making gains stick

Page type in the shape	What coverage did to it	Result
About, contact, careers, privacy	Category basics	Kept, small

Pages the shape had no room for. This is where Kestrel's differentiation shows:

- **Kestrel compared:** a hub that includes doing nothing, and one page for each alternative: a large firm, a lean specialist, the in-house team, and new equipment. "Before you buy another line, find out what this one can do" is the comparison that matters most.
- **Build the case:** a business case and board paper kit, a savings estimator, a sample review report, and a page for procurement and site safety. Committee deals stall at step 5. Nothing on the old site helped there.
- **What if the review finds nothing?** The finance director's fear, answered head-on, with the terms.
- **Talk to a client like you:** a client to call in each sector, with their permission. The checking step had no home on the old site, because the logos couldn't be rung.

Page families. 21 of the pages come from four templates: seven problems, four sectors, six roles and four comparisons. Add one page per case study. Each family gets one template and one brief, filled in for each problem, sector, role or alternative. That's how a 60-person firm keeps 47 pages consistent.

Priority sets the order, not the list. A common mistake is to build the sitemap from the top few moments. The result looks tidy and leaves most buyers without an answer. Every moment worth winning gets a home. The scores decide what's built first.

The page anatomy falls out of the jobs. List the jobs a page carries in stage order, and you have its blocks:

1. the buyer's question, answered first;
2. who it's for and not for;
3. what it is and what it costs;
4. proof beside the claim;
5. the reason, against the alternatives;
6. the way out, next to the price;
7. one next step.

The buying mode sets how deep each block goes. In committee buying, every block has to work when forwarded to someone who never saw the pitch.

The example sitemap

Every page earns its place through at least one of three signals, and the signals set its priority. No page exists without a reason, and every rejected page is listed with its reason too.


- **Canon:** buyers in the category expect it. It's on the category template, or on at least three of the three to five leading competitors' sites.
- **Coverage:** it answers moments from the grid. The page table counts them.
- **Strategy:** the positioning requires it: where the firm chooses to play, who it chooses to win, or its big idea.

Signals present	Priority
Canon + coverage + strategy	P1
Coverage + strategy	P1 (strategy-mandated)
Canon + coverage, or canon + strategy	P2
Coverage only	P3
Strategy only	P4
Canon only	Hygiene

The menus follow the journey: what the review is, the problems it solves, who it's for and the proof. A second menu holds help for the champion, insights and the company basics.

47 pages, plus one per case study, each with a reason

Blue: P1, built first. P2 and P3 follow in waves. H: hygiene, the category basics.



Kestrel's sitemap · 47 pages in seven menu groups

47 pages, plus one per case study. The old site had 14 pages and a blog. Every page below traces to its signals, and the moments column counts the moments it answers.

Page	Type	Signals	Priority	Existing page	Moments	Choosers it serves
Home	Hub (hero page)	canon, strategy	Hero	Improve	Routes every moment	everyone
The operations review	Service page	canon, coverage, strategy	P1	Replace (was Services)	20	COO, finance director, assistant

Page	Type	Signals	Priority	Existing page	Moments	Choosers it serves
The first six weeks	Approach page	canon, coverage	P2	Improve (was Our approach)	11	plant manager, improvement lead, COO
Price and terms	Pricing page	coverage, strategy	P1 (strategy)	New	6	finance director, assistant, owner
What if the review finds nothing?	Question page	coverage, strategy	P1 (strategy)	New	2	finance director
The improvement programme	Service page	canon, coverage	P2	Improve (was a paragraph on Services)	4	improvement lead, COO, procurement
Can you start this month?	Question page	coverage	P3	New	4	COO, owner, procurement, investor
Problems we solve (hub)	Hub	canon	Hygiene	Replace (was Solutions)	—	category basics
When output is flat	Solution page, by problem	canon, coverage, strategy	P1	New	10	COO, owner, finance director, plant manager, assistant
When costs keep rising	Solution page, by problem	canon, coverage, strategy	P1	New	10	COO, owner, finance director, plant manager, assistant
When cash is tied up in stock	Solution page, by problem	canon, coverage	P2	New	10	COO, owner, finance director, plant manager, assistant
When scrap and rework climb	Solution page, by problem	canon, coverage	P2	New	10	COO, owner, finance director, plant manager, assistant
When a new line won't reach speed	Solution page, by problem	canon, coverage	P2	New	10	COO, owner, finance director, plant manager, assistant
After an acquisition	Solution page, by problem	canon, coverage	P2	New	11	COO, owner, finance director, plant manager, investor, assistant

Page	Type	Signals	Priority	Existing page	Moments	Choosers it serves
When automation isn't paying back	Solution page, by problem	canon, coverage	P2	New	10	COO, owner, finance director, plant manager, assistant
Mid-sized manufacturing (hub)	Industry page	canon, coverage, strategy	P1	Improve (was Industries)	8	COO, assistant
Food and drink manufacturing	Industry page, by sector	canon, coverage	P2	New	3	COO, assistant
Packaging manufacturing	Industry page, by sector	canon, coverage	P2	New	3	COO, assistant
Metal and engineering manufacturing	Industry page, by sector	canon, coverage	P2	New	3	COO, assistant
Plastics and rubber manufacturing	Industry page, by sector	canon, coverage	P2	New	3	COO, assistant
For COOs	Role page	canon, coverage	P2	New	5	COO
For finance directors	Role page	canon, coverage	P2	New	9	finance director
For plant managers	Role page	canon, coverage	P2	New	10	plant manager
For owners and managing directors	Role page	coverage	P3	New	4	owner
For in-house improvement teams	Role page	coverage	P3	New	9	improvement lead
For private equity investors	Role page	coverage, strategy	P1 (strategy)	New	9	investor
Case studies (hub, filter by problem, sector, size)	Case study hub	canon, coverage	P2	Improve (was logos)	15	owner, investor, COO, finance director, assistant
Results across every review	Proof page	coverage, strategy	P1 (strategy)	New	8	owner, finance director, assistant, investor

Page	Type	Signals	Priority	Existing page	Moments	Choosers it serves
Talk to a client like you (references)	References page	canon, coverage	P2	New	7	COO, owner, improvement lead, investor
The team on your floor	Team page	canon, coverage	P2	Improve (was Team)	3	owner, plant manager
Kestrel compared (hub, doing nothing included)	Comparison hub	coverage, strategy	P1 (strategy)	New	23	COO, owner, finance director, investor, assistant
Kestrel vs a large firm	Comparison page	coverage	P3	New	2	owner, finance director
Kestrel vs a lean specialist	Comparison page	coverage	P3	New	1	improvement lead
Kestrel vs in-house	Comparison page	coverage	P3	New	8	COO, improvement lead
Kestrel vs new equipment	Comparison page	coverage, strategy	P1 (strategy)	New	14	COO, finance director
Build the case: business case and board paper kit	Tool / download	coverage, strategy	P1 (strategy)	New	8	COO, finance director
Savings estimator	Tool	coverage	P3	New	8	finance director, owner
Sample review report	Download	coverage	P3	New	9	plant manager, finance director, investor
For procurement and site safety	Question page	coverage	P3	New	6	procurement
What an operations review involves, and what it returns	Long-form guide	canon, coverage	P2	Merge (the blog)	23	COO, assistant
Mid-sized manufacturing operations benchmark (annual)	Research report	coverage, strategy	P1 (strategy)	New	8	finance director, assistant
Questions to ask before hiring an operations consultant	Question page	coverage	P3	New	5	assistant

Page	Type	Signals	Priority	Existing page	Moments	Choosers it serves
Making the gains stick	Guide	coverage	P3	New	4	improvement lead, finance director, plant manager
About	Hygiene	canon	Hygiene	Preserve	—	category basics
Contact and book a fit call	Hygiene	canon	Hygiene	Preserve	—	category basics
Careers	Hygiene	canon	Hygiene	Preserve	—	category basics
Privacy and terms	Hygiene	canon	Hygiene	Preserve	—	category basics

Build order. Priority decides the waves, not whether a page exists.

- **Wave 1, at launch (18 pages):** the home page, the 12 P1 pages and the 5 hygiene pages.
- **Wave 2 (18 pages):** the P2 pages, mostly the problem, sector and role families, the case hub and the references page.
- **Wave 3 (11 pages):** the P3 pages.

Rejected, with reasons:

- **Four generic "Solutions" pages** (transformation, strategy, digital, change). Deprecated: no moment needs them, and they were copied from larger firms' sites.
- **The blog as a section.** Merged: the twelve posts answered no buyer question. Their useful parts moved into the guide and the problem pages.
- **28 sector-by-problem pages** ("flat output in food and drink", and so on). Rejected: the questions don't change enough to need a page each. The case hub filters by both instead.
- **Region and city pages.** Rejected for now: no moment in the grid asks where the firm is based. Revisit if assistants start shortlisting by region.
- **A lean glossary.** Rejected: two of five competitors have one, which isn't canon, and no chooser asked for it.
- **A separate site-safety page.** Merged into the procurement page, because the same person forwards both.

The coverage test, run before the sitemap is locked:

- **No orphans.** Every one of the 149 moments has at least one page that answers it.
- **No passengers.** Every page except the hubs and hygiene pages answers at least one moment. The thinnest is "Kestrel vs a lean specialist", with one moment. If it can't be written well, it merges into the comparison hub.

- **No dead ends.** Every chooser can get from their first question to the next step. The improvement lead's path, for example: the problem page, "For in-house improvement teams", "Kestrel vs in-house", the references page, "The first six weeks", then "Making the gains stick".
- **No destructive cuts.** Cutting a page whose moments nothing else answers is flagged, and demoting it is offered instead. Cutting "When output is flat" would orphan four trigger moments. Cutting the investor page would orphan two: the investor's scope and checking questions.

The anatomy library

Every page type has a standard order of sections, because people and machines both expect it. A buyer who lands on a comparison page looks for the side-by-side; an assistant reading a price page looks for the number. Variety goes inside the shape, never across it. This is the library Kestrel's 47 pages are built from. Required sections are in bold.

Page type	Kestrel's pages	Sections, in order	The stage it carries most
Home (hub)	1	The question and who it's for · Find your way in, by problem or role · What happens, and what it costs · Proof · The way out · Against the alternatives · Questions · One next step	Found, then routing
Problem page	7	The problem in the buyer's words · What's usually behind it · What to try first yourself · When outside help pays, and what it costs · A case with this problem · Against the alternatives · Questions · One next step	Known, Found, Answered
Service page	2	What it is, in one sentence · Who it's for and isn't · What happens, week by week · Price and terms · Proof · The way out · Questions · One next step	Answered, Bought

Page type	Kestrel's pages	Sections, in order	The stage it carries most
Industry or sector page	5	Who in this sector it's for, by size · The problems this sector brings · Cases from this sector · A reference to call · One next step	Found, Preferred
Role page	6	This person's question, answered first · What they'll be asked to do, and what they won't · Their fear, answered · Proof from someone in their role · What to forward · One next step	Answered, Believed
Comparison page	5	When the alternative is the right answer · Side by side, fairly · When Kestrel is the better choice · Proof · Questions · One next step	Preferred
Proof pages: cases, results, references	3, plus one per case	The plant, its size and sector · Before and after, with who confirmed it · The plant manager's words · What changed on the floor · One next step	Believed
Price, terms and question pages	5	The answer, in the first sentence · The detail and the terms · What happens next · One next step	Bought
Tools and downloads	3	What it does for the buyer · The tool, or what's inside · How the numbers work · What to do with the result	Believed, Bought
Guides and research	4	The answer, in the first paragraph · The detail, a section per question · Sources and method · Related problems · One next step	Found, Answered

Page type	Kestrel's pages	Sections, in order	The stage it carries most
Hubs and hygiene	6	What's here, and for whom · Links, each with a one-line answer	Routing

Four rules hold on every page type:

- **The question is answered in the first sentence of its section.** Machines read passages, and people skim. A section that warms up first loses both.
- **Proof sits beside the claim it supports**, not on a separate page the buyer has to find.
- **The way out sits beside the price.** A finance director should never read one without the other.
- **One next step.** A page with three calls to action asks the buyer to make a choice the page should have made.

Step 5: the page brief, then a brief for every section

Every section on a page carries one job, so every section gets its own brief. The page brief doesn't disappear; it holds what no single section can: the moments the page serves, its place in the shape, the order of the sections, and the arc from the first question to the next step.

Level	The brief holds
Site	The shape and the story spine
Page	The moments served, the place in the shape, the section order and the arc
Section	One job: one stage, one chooser and moment, one bias, one claim with its proof

Why brief each section:

- **One job each.** A page carries five to eight jobs. Brief them together and they blur into general copy. Brief them separately and each section can be checked against its own job.
- **The same test, one level down.** A section that can't name its job gets cut.
- **Better AI output.** A section generated from a tight brief is sharper than a page generated from a loose one, and it can be checked and regenerated on its own.

- **Machines read passages.** Assistants and search systems typically pull a passage, not a whole page. Each section should stand alone: its question answered first, its own facts, its own proof.
- **Reuse.** A section such as "What if the review finds nothing?" can sit on several pages, briefed once and kept consistent.

The page brief comes first. Here it is for Kestrel's operations review page:

PAGE The operations review (service page)
 MOMENTS (20) COO "who does this for plants our size, in our sector?"
 (comparing, by sector); COO "fix it ourselves, buy equipment or bring someone
 in?" (options, by problem); finance director "fixed price or day rates? What
 if it finds nothing?" (comparing); assistant "what does an operations review
 cost, and what does it return?" (building the case)
 SHAPE Service page; links to the seven problem pages, the manufacturing hub,
 the role pages, Price and terms, and Kestrel compared
 THE ARC Recognise the problem, see who it's for, see what you get and what
 it costs, believe it, feel safe, take the small next step
 SECTION ORDER 1 The question answered 2 Who it's for 3 The six weeks and
 the price 4 Proof 5 Compared with the alternatives 6 What if it finds
 nothing? 7 Next step
 THE REASON (echo in every section) Savings found on the floor in six weeks,
 at a fixed price
 FEELING TO LEAVE "These people understand a plant like ours."

Then each section gets its own brief. Two of the seven:

SECTION 6 What if the review finds nothing?
 JOB Bought · finance director, comparing · losses loom larger
 THE ONE CLAIM If the review doesn't find savings worth more than its fee,
 you don't proceed and you don't pay for the programme.
 PROOF BESIDE IT The clause, quoted from the standard terms; how many reviews
 have triggered it (from the firm's records)
 FACTS, WORD FOR WORD FROM THE ONE TRUTH Review fee range; the threshold;
 what happens to the findings
 THE HUMAN TAKEAWAY "The worst case is that we learn something and walk
 away."
 THE SENTENCE A MACHINE SHOULD QUOTE Kestrel Advisory's operations review is
 fixed-price, and clients don't proceed to the programme if the savings found
 are less than the fee.
 LEADS TO Next step (section 7)

SECTION 4 Proof

JOB Believed · finance director and COO, comparing · we follow other people

THE ONE CLAIM Mid-sized plants like yours have found savings in six weeks.

PROOF BESIDE IT Two named plants, each with a before-and-after number and the plant manager's own words (with permission)

FACTS, WORD FOR WORD FROM THE ONE TRUTH Plant size, sector and the period measured, for each case

THE HUMAN TAKEAWAY "They've done this for a plant like ours, and it showed in the numbers."

THE SENTENCE A MACHINE SHOULD QUOTE [Plant name], a [size] [sector] plant, cut [measure] by [amount] after Kestrel Advisory's six-week review.

LEADS TO Compared with the alternatives (section 5); the full cases in the case study hub

The flow check. Once every section is written, read the page top to bottom against the arc in the page brief. If it reads as a stack of parts, fix the joins, not the sections. Every field in every brief comes from the moment grid, the jobs and the one truth, so any section can be regenerated when a moment changes, without rewriting the page.

The homepage: wireframe and section briefs

The home page is the hero page: it gets full prose drafts, not just direction. It has one job above all: tell the story spine and send each chooser to their page.

Page brief. Type: hub. Positioning wedge: savings found on the floor in six weeks, at a fixed price, for mid-sized manufacturers. Target queries (five at most): "operations consultants for mid-sized manufacturers"; "fixed-price operations review"; "how to reduce waste in a manufacturing plant"; "operations review cost"; "Kestrel Advisory reviews". Navigation label: Home. The arc: recognise the problem, see it's for you, see how it works and what it costs, believe it, feel safe, take the small next step.

The home page, section by section, each with one job

The review · Problems · Who it's for · Proof

Check fit

1 HERO

Find the savings on your floor in six weeks, at a fixed price.

Operations reviews for manufacturers with 100–500 staff

Book a 30-minute fit call

Two named plants · fixed price · on site

2 YOUR PROBLEM, YOUR ROLE

Flat output

Rising costs

Cash tied up in stock

4 more problems · six role pages · not for global programmes

3 THE SIX WEEKS, AND THE PRICE

Week by week: who's on site, what they need · price range shown

4 PROOF

Plant A: measure, before and after

Plant B: measure, before and after

5 WHAT IF IT FINDS NOTHING?

If the savings found are less than the fee, you don't proceed.

6 COMPARED WITH THE ALTERNATIVES

Kestrel

Large firm

New equipment

In-house

7 QUESTIONS

How long does it take? · What does it cost?

How much of my team's time will it take?

8 NEXT STEP

Book a 30-minute fit call

Business case kit →

Footer: build the case, insights, about, careers, contact, privacy

Job: stage · chooser · bias

Found, Answered · COO

good enough beats best

Preferred · every chooser

feeling first, reasons after

Answered, Bought · plant manager, FD

a hard choice stalls

Believed · finance director

we follow other people

Bought · finance director

losses loom larger

Preferred · COO

feeling first, reasons after

Answered · every chooser

a hard choice stalls

Bought · COO

a hard choice stalls

Kestrel home page wireframe · eight sections, one job each

The wireframe shows structure, not design: what each section holds, in what order, and the job it does. Its brief is in the table below.

Page 24 of 39

#	Section	Job (stage · chooser · bias)	Human goal	AI goal	Entities	Guard (passes only if)
1	Hero	Found, Answered · COO · good enough beats best	Recognise their problem and see in five seconds that this is for a plant like theirs	One quotable sentence: who it's for, what it does, at what price model	Kestrel Advisory; operations review; mid-sized manufacturers	The headline answers a buyer question, not a slogan; the reason appears
2	Who it's for	Preferred, Answered · each committee member · feeling first	Each person finds their way in, by problem or by role	States the sizes, sectors and problems served	Each problem; each role; sizes; sectors	Seven problem links and six role links; one line on who it isn't for
3	The six weeks, and the price	Answered, Bought · plant manager and finance director · a hard choice stalls	See exactly what happens, and what it costs	States duration, scope and price range as facts	The review; six weeks; price range	The price range is on the page and matches the one truth
4	Proof	Believed · finance director · we follow other people	Believe it works for plants like theirs	Each case is quotable: plant, size, measure, result	Two named plants; plant managers	Every figure is sourced; no unsourced numbers
5	What if it finds nothing?	Bought · finance director · losses loom larger	Feel the downside is capped	States the terms plainly	The guarantee clause	Matches the contract wording word for word
6	Compared with the alternatives	Preferred · COO · feeling first, reasons after	Have a reason to repeat to the board	A comparison an assistant can repeat accurately	A large firm; a lean specialist; in-house; new equipment; doing nothing	Fair to the alternatives; no claim without proof
7	Questions	Answered · every chooser · a hard choice stalls	Get the remaining doubts answered	At least three answerable questions, each answered in its first sentence	The review; terms; procurement	Each answer is under 60 words and matches the one truth

#	Section	Job (stage · chooser · bias)	Human goal	AI goal	Entities	Guard (passes only if)
8	Next step	Bought · COO · a hard choice stalls	Take the smallest safe yes	States what happens after booking	Fit call; one-page scope	One action only; it resolves to a real page or form

Sections 1, 3, 4 and 8 are required. The rest are recommended for committee buying, and the order follows the arc. The flow check reads the page top to bottom before it's approved.

The finished example: "When output is flat", written in full

A brief is a promise. This is what one turns into. "When output is flat" is a P1 problem page in wave 1. It answers 10 moments, for the COO, the owner, the finance director, the plant manager and the assistant. Each section below carries its job in italics, then the copy as it would go live. Kestrel, its plants and its people are invented.

1. The problem in the buyer's words

Known, Found · COO · familiar feels safe

Output's flat, the lines are busy, and you're still shipping late

If your plant is running flat out and output hasn't moved in a year, you're not short of effort. You're short of a clear view of where the capacity goes. In most mid-sized plants we review, it's already in the line, lost to changeovers, small stops and waiting, long before anyone needs a new machine.

2. What's usually behind it

Answered · COO and plant manager · a hard choice stalls

Flat output in a busy plant usually comes from three places, in this order:

1. **Changeovers** that run long because nobody owns the clock. In the reviews we've run, they're the largest single loss more often than not.
2. **Small stops:** the jams and waits under five minutes that never reach the downtime log.
3. **Planning** that loads the bottleneck line with the wrong mix, so it waits while others run.

New equipment fixes none of these. It adds capacity to a line that's already losing its own.

3. What to try first yourself

Answered · plant manager and improvement lead · losses loom larger

Before you call anyone, including us: time five changeovers on your bottleneck line, start to first good part. Log every stop over 30 seconds for one week, by hand. If the two together cost more than a shift a week, the capacity is there to find. If they don't, your problem is probably demand or planning, and the guide below will help more than a review will.

4. When outside help pays, and what it costs

Answered, Bought · finance director · losses loom larger

The operations review is six weeks on your floor, at a fixed price agreed before we start: A\$45,000 to A\$90,000, depending on the number of lines. **If the review doesn't find more than it costs, you don't go ahead with the changes, and you don't pay for them.** That's in the contract, word for word.

5. A plant like yours

Believed · finance director and owner · we follow others

Brennock Packaging, 240 staff. Output on the bottleneck carton line rose from 41 to 47 a minute, a 15% rise, without new equipment. Changeovers fell from 52 minutes to 19. Brennock's finance director confirmed the figures four months after the review.

"We were three weeks from signing for a new line. We didn't need it." Plant manager, Brennock Packaging

[More cases with flat output →]

6. Against the alternatives

Preferred · COO and finance director · feeling first, reasons after

If you're weighing	Choose it when	Choose a review first when
A new line or machine	Demand is above what the line can make even at its best	You haven't measured what the line makes at its best
Your own improvement team	They have the time and the data, and someone outside isn't needed to settle the argument	They're stretched, or the same fixes keep coming undone

If you're weighing	Choose it when	Choose a review first when
A large consultancy	You need a programme across many sites at once	You need one plant fixed, by the people who pitched

7. Questions

Answered · every chooser, and the assistant · a hard choice stalls

How long until output rises? Most plants see changeover times fall within the six weeks; output follows as the schedule catches up.

What do you need from our team? About two hours a day from the plant manager, and access to the floor on every shift. "The first six weeks" has the detail.

What if our problem is demand, not capacity? Then we'll say so on the first call, and you won't need a review.

8. One next step

Bought · COO · a hard choice stalls

Book a 30-minute fit call. We'll ask about your lines and your numbers, and tell you plainly whether a review would pay. You'll have a one-page scope within two days.

What to notice. Every claim on the page is in the claims register: the six weeks, the fixed price, the way out and the Brennock figures, confirmed by the client. Section 3 tells some buyers they don't need Kestrel, which is why the rest of the page is believed. And every section answers its question in its first sentence, so an assistant can lift any one of them whole.

Templates

Copy these. Every field is filled from something decided earlier, and the source is named beside it, so nobody fills a field from memory.

The page brief

PAGE	[name]	TYPE	[from the anatomy library]
PRIORITY	[Hero / P1-P4 / Hygiene]	WAVE	[1 / 2 / 3]
SIGNALS	[canon / coverage / strategy]	EXISTING PAGE	[preserve / improve / merge / replace / new]
MOMENTS	[each chooser · step · need, from the moment map]		
THE ARC	[from the first question to the one next step, in one line]		
REASON	[the reason to be chosen, word for word, from Brand Strategy]		
CLAIMS USED	[each one from the claims register, with its state]		
TARGET QUERIES	[five at most, in the buyer's words]		
SECTIONS	[in the anatomy library's order; required ones marked]		
LINKS	[parent, children, related; where the next step goes]		
NAV LABEL	[two to four words]		
FRESHNESS	[evergreen / refresh-critical / time-bound; what triggers a refresh]		

The section brief

SECTION	[number and name]
JOB	[stage · chooser and moment · bias]
HUMAN GOAL	[what the reader should know or do after this section]
AI GOAL	[the one passage an assistant should be able to lift]
THE QUESTION	[answered in the first sentence]
CLAIM + PROOF	[one claim, its proof beside it, both from the register]
ENTITIES	[names, places, products, figures]
WORDS	[target length]
GUARD	[passes only if ...]

The page decision log (one row for every existing page)

Old page	Decision	Moments it served	Where they go now	Reason, in writing
	Preserve / improve / merge / replace / deprecate			

The measures sheet (one row per page, read monthly)

Page	The stage it carries	Its measure	Last month	This month	Predicted	What changed

Step 6: the business check

Before design starts, check the reconciled site against the business:

- ☐ **Revenue:** the review page, the problem pages and the manufacturing hub carry the moments behind most new work. If they don't, the moment list has drifted from the business. Check every strategy flag the same way: the investor page is P1 only because investor-owned manufacturers are a chosen target. If they aren't, it drops to P3.
- ☐ **One reason:** every page echoes "savings found on the floor in six weeks, at a fixed price". Pages that tell a different story get rewritten or cut.
- ☐ **One truth:** the price range, duration and terms are identical on the site, in proposals and on every listing.
- ☐ **The offer:** can the firm really honour "if it finds nothing, you don't proceed"? If not, it's an offer decision, not a copy decision.
- ☐ **Capacity:** 47 pages is a lot for a 60-person firm, which is why wave 1 is 18. Two P1 pages take time to earn: the benchmark report needs data from enough reviews, and the results page needs every figure sourced. Until they're ready, the guide and the case studies carry their moments. The order changes; the list doesn't.

Step 7: the measures

Each page is measured by the stage it carries, never by its traffic.

Page	Measure	Stage
Manufacturing industry page	Share of assistant answers to "best operations consultants for mid-sized manufacturers" that name Kestrel, and give its reason	Found, Preferred
The operations review	Share of qualified visitors who book the fit call	Bought
"What if it finds nothing?"	Deals that stall on risk, from sales notes (should fall)	Bought
Case study hub	Prospects who mention a case on the first call	Believed
Role pages	Pages forwarded within the committee, and how many people from one account visit	Answered, per person
"For procurement"	Days from verbal yes to signed contract (should fall)	Bought
Problem pages	First calls where the buyer describes their problem in Kestrel's words	Known, Found

Page	Measure	Stage
Kestrel vs new equipment	Deals where the client chose the review before a capital purchase	Preferred
Business case kit	Downloads followed by a second visitor from the same company	Believed
For in-house improvement teams	Deals where the improvement lead is named as a supporter in sales notes	Preferred

In the campaign plan

The Campaign Engine (step 10) turns the builds into a year of campaigns, each around a moment or an occasion. The website is in every one, because it supports every live cell. Its rule in the campaign plan is simple: **a campaign lands on a page that already answers its moment**. A campaign that needs a page the sitemap doesn't have has found a moment the map missed, so the page goes on the sitemap, or it's a time-bound page with a date it comes down.

Kestrel campaign	Moment or occasion	Lead build	Lands on	What the page adds for the campaign
"Before you buy another line"	COO · trigger · flat output; budget season	Advertising	"Kestrel vs new equipment", then "When output is flat"	The ad's line repeated in the headline; the five-changeover test from the problem page
The annual benchmark	Framing, for the finance director and the assistant; the report's release	Content	The benchmark page	A summary in quotable passages, so assistants cite the figures; the full report behind one form
Board season	Finance director · building the case	The pitch	"Build the case" kit	A version of the kit dated for this year's budget round

Kestrel campaign	Moment or occasion	Lead build	Lands on	What the page adds for the campaign
Owners' round table	Owner · options and checking; the event date	Events	"For owners and managing directors"	The speakers, the date, and a client owner to call afterwards
Investor 100 days	Investor · trigger and options	Partners and channels	"For private equity investors"	A one-page brief advisers can forward to their portfolio companies

Every campaign is measured on the page as well as in the channel: did the moment's stage move for the buyers it brought? An ad that brings clicks to a page that doesn't answer the question has bought traffic, not selection.

Running it

A site built from moments goes stale the way the moments do: a new rival, a new claim, a changed price. Running it means keeping the pages true to the map.

Who owns what. At Kestrel, the marketing lead owns the site and the sitemap. Each page family has an owner who knows the subject: the head of delivery owns the problem pages and "The first six weeks"; the managing partner owns price, terms and the way out. Nobody publishes a claim the claims register hasn't signed.

What happens	What it triggers	Pages Kestrel would touch
A price, term or the way out changes	Every page that states it, the same day	Price and terms, the review, the home page, "What if it finds nothing?", procurement, every problem page's section 4
A case is signed off	The case page, then every page whose proof it's better than	The case hub, its problem page, its sector page, its role pages
A new moment joins the map	The coverage check: which page answers it?	A section on an existing page, or a new page on the sitemap

What happens	What it triggers	Pages Kestrel would touch
A page is cut or merged	The orphan check: which moments did only it answer?	Whatever page takes them over, before the cut
The monthly scan shows an assistant misquoting a page	Fix the page's passage, then the source the assistant is reading	Usually the price, the sectors or the size range
A refresh-critical page reaches its date	Review against the claims register	The benchmark (annually), results across every review (quarterly)

The cadence. The measures sheet monthly. The page decisions quarterly, with the Score. The whole sitemap once a year, or when the map changes by more than a tenth of its moments.

The basics, every release. A plain-text copy of every page for machines to read, plus llms.txt and sitemap.xml. Structured data only for questions the page answers. Pages that load fast on a phone on a plant floor; text at least 16 pixels; headings in order; every image described; every form usable by keyboard; contrast that passes WCAG 2.2 AA. No dead links, and every next step resolves to a real page.

Common mistakes

Each mistake breaks a stage, and most work against a bias buyers already have. Kestrel's old site made nearly all of them.

The mistake	The stage it breaks	The bias it ignores	On Kestrel's old site	Instead
Pages about the firm, not the buyer's problem	Known, Found	Familiar feels safe	Services, approach, about: nothing in a buyer's words	A page per problem, headed in the words buyers use
Warming up before answering	Found, Answered	Good enough beats best	"In today's competitive landscape..."	The answer in the first sentence of every section

The mistake	The stage it breaks	The bias it ignores	On Kestrel's old site	Instead
Claims without proof beside them	Believed	We follow others	A wall of logos, no numbers	Named plant, before and after, who confirmed it, on the same page as the claim
Hiding the price	Bought	Losses loom larger	"Price on application"	The range and the way out, side by side
No reason against the alternatives	Preferred	Feeling first, reasons after	Nothing on new equipment or doing it in-house	A fair comparison, including when the alternative is right
A sitemap built from the top ten moments	Answered, for everyone else	A hard choice stalls	14 pages; 84 moments with no page	Coverage first; priority sets the order
Three calls to action per page	Bought	A hard choice stalls	"Contact us", "Download", "Subscribe" on every page	One next step, the smallest safe yes
Writing only for the champion	Answered, per person	Losses loom larger	Nothing a COO could forward to the finance director	Role pages, and a kit the champion can send as it is
Copying the competitors' sitemap	Preferred	Familiar feels safe	Four generic "solutions" pages copied from larger firms	Canon for what buyers expect; strategy for what makes you different
Measuring pages by traffic	All of them, unseen	Good enough beats best	A traffic report, and no idea which page won work	Each page measured by the stage it carries

The bias to watch in the team: **the curse of knowledge**. Everyone at Kestrel knows what an operations review is, so nobody thought to explain it. Read each page as the plant manager who's never hired a consultant.

The Loop for this build

Each quarter, the Loop (step 11) puts the site's slice of the Score beside what was predicted, and asks one question per page family: **did the stage this page carries move, for the moments it was built for?**

What the Loop finds	What it means	What changes
Answered rose, Found didn't	The page answers, but nothing sends buyers to it: a hidden moment	Being named and content take the moment; the page stays
Found rose, Bought didn't	Buyers arrive and leave	The page's price, proof or next step; check it against the anatomy library
Nothing moved after two quarters	The page doesn't answer the question buyers ask, or the moment is wrong	Rewrite from the section briefs; if it still doesn't move, check the moment with interviews
An assistant quotes the page wrongly	A passage is ambiguous, or another source contradicts it	Fix the passage; find the other source and correct it
Sales calls raise a question no page answers	A moment the map missed	Add it to the map, then run the coverage check

What Kestrel's first two loops would look for. After wave 1, Answered at the problem, options and comparing moments should rise first, from 9 towards half, because the pages now exist. Found should move a quarter or two later, once assistants and search have read them. Bought should rise wherever price and the way out now sit side by side. If Answered rises and the Score barely moves, that's expected in the first quarter: the gate opens slowly. If Answered doesn't rise, the pages aren't answering, and no amount of promotion will fix that.

The rules behind the method

Shape, coverage and reconcile rest on fifteen rules. They hold whether a team works in a spreadsheet or in software.

1. **Every existing page gets a decision:** preserve, improve, merge, replace or deprecate. Deprecating needs a written reason. Pages that already earn their keep are kept.
2. **Competitor canon:** a page type on at least three of the three to five leading competitors' sites is expected in the category.

- 3. **Coverage cells hold more than questions:** each chooser-by-step cell records the pain, the gain, the job to be done and the objection.
- 4. **Priority by signals:** canon, coverage and strategy. No page exists without a reason, and rejected pages stay visible.
- 5. **Priority sets the build order, never which pages exist.**
- 6. **Page families:** one template applied to a set of entities (each problem, sector, role, case), not hand-built pages.
- 7. **Freshness:** every page type is marked evergreen, refresh-critical or time-bound, with what triggers a refresh.
- 8. **Sitemap triage with a coverage check:** no orphaned moments, every chooser served end to end, and a warning before cutting the only page that covers a moment.
- 9. **Page brief fields:** type, positioning wedge, target queries (five at most), section list, hero flag and navigation label.
- 10. **Section brief fields:** the job (stage, chooser and bias), human goal, AI goal, entities, word target, and guard criteria.
- 11. **Canonical page shapes:** each page type has a standard order of sections, from the anatomy library.
- 12. **Fidelity by priority:** P1 pages get full design; P2 to P4 get content in a template; hygiene pages are templated.
- 13. **Machine legibility:** named entities, at least three answerable questions and a quotable passage per page; structured data only for questions the page answers; a plain-text copy of every page, plus llms.txt and sitemap.xml.
- 14. **No unsourced figures:** every number sourced, signed off, removed or bound to live data before it's published.
- 15. **Internal links built from the page list:** breadcrumbs, parent, children and related pages; every next step resolves to a real page.

Shapes by business type

The playbook keeps a library of starting shapes. Each one is a default to reconcile, not a template to fill. Expect each to gain and lose pages once coverage is run.

Business type	Typical mode	Expected page types
B2B professional services	Committee or owner-led	Services, solutions, industries, roles, case studies, approach, team, insights

Business type	Typical mode	Expected page types
B2B software	Committee	Product, features, integrations, pricing and plans, security and compliance, customers, comparisons, docs
Considered consumer (finance, health, education)	Considered personal	Products, eligibility, rates and fees, calculators, how it works, reviews, help
Ecommerce	Quick choice or habitual	Categories, product pages, reviews, delivery and returns, size or fit guides, offers
Local services	Owner-led or considered personal	Services, service areas, prices or typical costs, reviews, about the people, booking
Marketplaces	Varies on each side	Separate shapes for buyers and sellers, trust and safety, how it works, fees

The same flow for the other thirteen builds

Every build has a shape too. The method is the same everywhere: shape, coverage, reconcile, then brief, check and measure.

Build	Category	Its shape (top-down)	Its coverage unit (bottom-up)
Content	Owned	Pillars and formats the category expects	One piece per unanswered question
Social	Owned	Content pillars and a posting rhythm	One post per piece of evidence others should see
Advertising	Paid	Campaign architecture: brand platform, always-on, then promotions	One ad per need you must be remembered for
Being named: search and AI answers	Earned	The sources and listings the category is judged on	Each moment's questions, and who gets named today

Build	Category	Its shape (top-down)	Its coverage unit (bottom-up)
PR and reviews	Earned	Press office, awards, the review programme	One story or review per judgement that lacks outside support
Partners and channels	Partnered	The partner tiers the category runs on	One partner per audience you can't reach yourself
Creators	Partnered	The programme's tiers and rules	One partnership per claim that needs an outside voice
Outbound and account-based	Direct	The target list and its tiers	One sequence per named firm's likely moment
The pitch	Direct	The standard proposal structure	One section per fear, per person in the room
Events	Direct	The calendar of events the category attends	One conversation per moment a room decides better than a page
Email	Relationship	The lifecycle programme: welcome, nurture, onboarding, renewal	One message per question at each step
Automation	Relationship	The systems map	One trigger per moment that's lost for lack of a timely reply
Customers	Relationship	Onboarding, results reviews, community, referral	One touch per moment after the signature

Feeling a bit overwhelmed?

You should be, a little. That was one firm with one main offer: eight choosers, seven problems, ten steps, five buying modes, 149 moments and 47 pages, each with a page brief and a brief for every section. Leave a moment out and a buyer gets their answer somewhere else. Few teams can hold all of that in a spreadsheet, let alone keep it current as buyers and competitors change.

That's why we built RocMode's Website Engine.

- **Your buyer map, built for you.** RocMode maps your choosers, their journeys and their moments from what's known about your brand, with the pain, gain, job and objection in every cell. Change anything you like; your changes are kept.
- **Your website strategy and sitemap.** From that map it builds the strategy and the sitemap, with a reason for every page and a check that every moment has a home.

[Join the RocMode beta](#)

The Content Playbook

The idea

Content earns its place by answering a question a buyer actually has, before anyone is selling to them. Most content plans start from a calendar: two posts a week, a webinar a quarter, an ebook for the funnel. Then they hunt for topics to fill the slots. That gives volume, and buyers learn to ignore it.

This playbook starts from the other end. The coverage unit is **one piece per unanswered question**. Every question comes from a moment in the grid. Every piece answers a set of questions for named choosers, and the calendar is whatever it takes to write them.


Three things follow:

- **Fewer, deeper pieces.** One guide that fully answers ten related questions beats ten posts that each half-answer one.
- **Every piece feeds other builds.** A problem guide becomes the website's problem page, the passages assistants quote, the newsletter issue, the social posts and the ad's landing page. Content is the raw material; the other builds are its outlets.
- **The calendar comes last.** First decide what must exist. Then decide the order. Then set the rhythm the team can keep.

The footprint, and Kestrel in 60 seconds

Content: lit in 17 cells, leading 24 of 149 moments

Kestrel's moment map. Rows: choosers. Columns: journey steps. Blank: that chooser isn't present at that step.



The content footprint on Kestrel's grid

Content leads where buyers are still working out what's wrong: the COO's and finance director's framing, the finance director's first alarm about costs, and almost everything the in-house improvement lead asks. It supports the early moments other builds lead (the ad, the social post, the assistant's answer, the website's options pages) and two of the questions that come after the review.

Kestrel in 60 seconds. Kestrel Advisory is the invented firm used in every playbook: a 60-person operations consultancy that helps mid-sized manufacturers find waste and lift output. Its main offer is a fixed-price, six-week operations review. Its strategy is in [Part one](#), and its full moment map in [Where you stand, part one](#).

- **Who chooses:** usually a committee. The COO starts it, the finance director approves the spend, the plant manager lives with it and procurement runs the paperwork. The owner, the in-house improvement lead, an investor and AI assistants all have a say.
- **The reason to be chosen:** "We find the savings on the floor in six weeks, at a fixed price, and the review pays for itself or you don't proceed." (Invented. A real firm would need proof for every part of that sentence.)
- **The alternatives:** a large consulting firm, a specialist lean consultant, the in-house improvement team, buying new equipment, or doing nothing.

Its slice of the Score

Content's part of [the Selectability Score](#) is the early gate and the first answer: **Known** and **Found** at trigger and framing, and **Answered** wherever a buyer is still working out what's wrong.

Kestrel today, at the 23 moments content leads before a deal	Known	Found	Answered	Moment score
Average	10	0.7	5	0.1

One detail in those numbers is worth more than the averages. The only two moments where search finds Kestrel at all are both content: two old blog posts, on flat output and rising costs, that still rank on page one. They were written years ago, they answer part of a real question, and they still work. Twelve posts, two of them still ranking, and those two are the whole of Kestrel's presence before a buyer calls. That's the case for this build in one line: content that answers a buyer's question keeps working; content that fills a slot doesn't.

Step 1: the shape, pillars and formats

Content has two kinds of shape: the **pillars** (what the firm is known for talking about) and the **formats** buyers in the category expect.

The pillars come from the story spine. Kestrel's spine is: the problem in the buyer's words, the reason, the proof pattern and the way out. So its pillars are:

Pillar	What it's for	The spine it carries
The problems	Why output is flat, costs climb, stock piles up, and what to check first	The problem, in the buyer's words
The cost	What each problem costs, in money a finance director can use	The reason: savings found on the floor
The evidence	What reviews find, across plants, sectors and sizes	The proof pattern
Making change stick	Why improvements fade, and how in-house teams keep them	The way out, and what clients keep

Nothing goes outside the four. "The future of manufacturing" isn't a pillar; it's a slot-filler.

The formats come from the category. Count what the leading competitors publish. A format on at least three of five sites is canon: buyers expect to find it, and its absence costs belief.

Format	Competitors using it (of 5)	Canon?
Long-form guides	5	Yes
Articles	5	Yes
Case stories	5	Yes (built in the Website Playbook)
Research reports	4	Yes
Newsletters	3	Yes (run in the Email Playbook)
Webinars	3	Yes
Templates and checklists	2	No
Video	2	No
Podcasts	1	No

The counts are invented for the example. Canon tells you what's expected, not what's enough: a format only earns a piece when a question needs it.

Step 2: coverage, the question bank

Start from the footprint: 17 lit cells, split by problem where the question changes. That's 77 moments: 24 that content leads and 53 it supports.

A moment content leads holds more than its top question. The cell's pain, cost and objection each raise a question of their own, so lead moments get four. For the COO's framing moment on flat output:

1. **The top question:** "Is it people, process, equipment or planning?"
2. **What's causing it:** "Why does output stay flat when we're running flat out?"
3. **What it costs:** "What is the lost output worth to us each month?"
4. **What to do first:** "What should we check this week, before we spend anything?"

Supporting moments get one question each: the content that helps the lead build do its job.

Cell	This build	Moments	Questions each	Questions
COO · framing	Leads	7, one per problem	4	28
Finance director · trigger	Leads	7	4	28
Finance director · framing	Leads	7	4	28
Improvement lead · framing: "Why didn't lean hold?"	Leads	1	4	4
Improvement lead · options: "Can we do this ourselves?"	Leads	1	4	4
Improvement lead · extend: "Can we run the next review ourselves?"	Leads	1	4	4
COO · trigger	Supports advertising	7	1	7
Plant manager · trigger	Supports social	7	1	7
COO · options	Supports the website	7	1	7
Finance director · options	Supports the website	7	1	7
Assistant · framing	Supports search and AI answers	7	1	7
Assistant · options	Supports search and AI answers	7	1	7
Assistant · building the case	Supports search and AI answers	1	1	1
Improvement lead · delivery	Supports customers	1	1	1
Plant manager · extend	Supports customers	1	1	1
COO · scope: "What should the review cover?"	Supports the website	7	1	7

Cell	This build	Moments	Questions each	Questions
Assistant · scope: "What should I look for in an operations consultant?"	Supports being named	1	1	1

149 questions. That's the bank. Each one is written in the buyer's words, tagged with its cell, its problem, its stage and its bias, and kept in one list the whole team can see. A writer never starts from a topic. They start from a set of questions, and the piece is finished when every one is answered.

Step 3: what's answered today, then gaps, biases and jobs

Audit what exists against the bank. Kestrel's old blog (invented) had twelve posts. Five answer no question in the bank at all: company news, two event write-ups, an award and "the future of manufacturing". The other seven touch 14 questions in passing. **None of the 149 questions is answered fully.**

Pillar	Questions	Touched today	The stage that fails	The bias at work
The problems	77	9, by posts on equipment effectiveness, scheduling and changeovers	Known: nothing in the buyer's words; Found: posts too thin to quote	Familiar feels safe: buyers read the names they already know
The cost	56	2	Believed: no figures a finance director could use	Losses loom larger: without a cost, doing nothing feels safe
The evidence	1	0	Believed: nothing across reviews	We follow others
Making change stick	15	3, by one post on lean	Preferred: nothing for the improvement lead, who reads Kestrel as a threat	Losses loom larger

Each gap becomes a job:

Moment	Stage	Bias	The job
COO · framing, each problem	Known	Familiar feels safe	One guide per problem, in the COO's words, that names the likely causes and what to check first
COO · framing	Answered	A hard choice stalls	End each guide with the options, including fixing it in-house and buying equipment, and when each is right
Finance director · trigger, each problem	Known	Losses loom larger	Put a cost on each problem, in money and in a form a finance director can test with their own numbers
Finance director · framing	Believed	We follow others	Back every cost with sourced figures from the benchmark, never estimates dressed as facts
Assistant · framing and options	Found	Good enough beats best	Write each guide in passages that answer one question each, so assistants can quote them
Improvement lead · framing	Preferred	Losses loom larger	Explain why improvement programmes fade without blaming the people who ran them
Improvement lead · options	Believed	Familiar feels safe	Publish how to run a review in-house: candid, useful, and the best proof that Kestrel isn't a threat
Improvement lead · extend	Kept	We follow others	Show how client teams run the next review themselves
Plant manager · extend	Kept	Losses loom larger	A 90-day check that the changes stuck, which plant managers can run alone
Everyone · evidence	Believed	We follow others	An annual benchmark from Kestrel's reviews, published with clients' permission

The in-house guide will look risky to some. Giving away the method seems to give away the sale. It doesn't. The improvement lead was going to block Kestrel anyway; the guide turns a blocker into a reader. And a team that tries it alone and gets stuck knows exactly who to call.

Step 4: reconcile, the content plan

Put the questions into the pillars, and choose the format each set of questions needs. Group questions that one reader would want answered together. Kestrel's 149 questions fall into **19 core pieces**. This is the content build's equivalent of a sitemap.

Piece	Pillar	Format	Questions	Signals	Priority
When output is flat: causes, costs, what to check first, and what a review should cover	The problems	Long-form guide	11	Canon, coverage, strategy	P1
When costs keep rising	The problems	Long-form guide	11	Canon, coverage, strategy	P1
When cash is tied up in stock	The problems	Long-form guide	11	Canon, coverage	P2
When scrap and rework climb	The problems	Long-form guide	11	Canon, coverage	P2
When a new line won't reach speed	The problems	Long-form guide	11	Canon, coverage	P2
After an acquisition: where the operating margin hides	The problems	Long-form guide	11	Canon, coverage	P2
When automation isn't paying back	The problems	Long-form guide	11	Canon, coverage	P2
What flat output really costs	The cost	Article	8	Canon, coverage, strategy	P1
What rising costs are made of	The cost	Article	8	Canon, coverage, strategy	P1
What stock really costs	The cost	Article	8	Canon, coverage	P2
What scrap and rework really cost	The cost	Article	8	Canon, coverage	P2
What a slow ramp-up costs	The cost	Article	8	Canon, coverage	P2

Piece	Pillar	Format	Questions	Signals	Priority
What margin an acquisition can find on the floor	The cost	Article	8	Canon, coverage	P2
What idle automation costs	The cost	Article	8	Canon, coverage	P2
The mid-sized manufacturing operations benchmark (annual)	The evidence	Research report	1, and feeds every guide	Canon, coverage, strategy	P1
How to run an operations review in-house, and what a good one covers	Making change stick	Long-form guide	5	Canon, coverage, strategy	P1
Why improvement programmes fade	Making change stick	Article	4	Canon, coverage	P2
Running the next review yourselves	Making change stick	Article	5	Canon, coverage	P2
The 90-day check: did the changes stick?	Making change stick	Checklist	1	Coverage	P3

The questions add up to 149. Every piece traces back to named cells; every question has a piece.

Families. The seven problem guides share one template and one brief, filled in per problem. So do the seven cost articles. That's 14 of the 19 pieces from two templates, which is how a small team keeps them consistent.

One piece, many outlets. Each core piece is written once and used by the other builds. This is where content pays for itself:

Core piece	What it becomes	Assets
Each problem guide (7)	The website's problem page; a newsletter issue; four social posts; a one-page checklist	49
Each cost article (7)	A section of the finance director's page; a newsletter issue; two social posts; an ad for the COO's trigger moment	35
The benchmark report	A press release; seven problem cuts; four sector cuts; a webinar; ten social posts	24
The four improvement-lead pieces	A newsletter issue and two social posts each	12

19 pieces become 120 assets. The webinar finally earns its place: it presents the benchmark, rather than filling a quarterly slot.

Build order. Priority sets the order, never the list. Six P1 pieces first, then twelve P2, then one P3. At two core pieces a month, that's about ten months for one writer, with the benchmark refreshed each year.

Rejected, with reasons:

- **"The future of manufacturing" and other opinion pieces.** No moment asks for them. They fill a calendar, not a buyer's need.
- **News, events and awards as articles.** They answer no question in the bank. If they're worth sharing, they go in the newsletter or on social.
- **A weekly posting quota.** Volume set before questions is how the old blog ended up with twelve posts that answered nothing.
- **Gating the guides behind a form.** Assistants can't read behind a form, so a gated guide can't be quoted. The guides stay open. Only the benchmark's detailed data is gated, for buyers who want it.
- **28 sector-by-problem articles.** The questions don't change enough. Sector examples go inside each guide instead.
- **Mass-produced articles.** Thin pieces dilute the ones that do the work, and contradict the proof pattern.

The coverage test: every one of the 149 questions has a piece; every piece answers named questions; every lead moment is answered in full by at least one piece; and every piece has at least one outlet in another build.

The anatomy library

Each kind of piece has a standard order of sections. Readers know where to look, and assistants find each answer where they expect it. Required sections are in bold.

Piece	Kestrel's pieces	Sections, in order	Length
Problem guide	7	The problem in the buyer's words, answered in one paragraph · The likely causes, ranked · What it costs · What to check this week · The options, including doing nothing · What a good review covers · A case · Related questions · One next step	2,500 to 3,500 words

Piece	Kestrel's pieces	Sections, in order	Length
Cost article	7	The cost, in one sentence, with its source · How it's made up · Work out your own · What usually moves it · One next step	900 to 1,400 words
Research report	1	The headline findings · Method and sample · Findings by problem · Findings by sector and size · What the best plants do differently · Sources and permissions	20 to 30 pages; a 2-page summary open to all
Method guide (in-house)	1	What you'll be able to do · The steps, in order · The templates · Where teams get stuck · When to call for help	3,000 to 4,000 words
Practice article	2	The answer first · Why it happens · What works, with an example · One next step	1,000 to 1,500 words
Checklist	1	What it checks, and when to use it · The checks · What to do with the result	One page

Three rules hold for every piece:

- **Every section answers one question from the bank**, and answers it in its first sentence.
- **Every figure carries its source in the same sentence**, and comes from the benchmark, a named case or a cited study.
- **Every piece ends with one next step**, chosen for the reader's moment: the checklist after a guide, the business case kit after a cost article.

Step 5: the piece brief, then a brief for every section

The piece brief holds what no single section can: the questions, who they're for, the pillar, the arc from the first question to the next step, and where the piece will be used. Then every section gets its own brief, answering one question for one chooser.

PIECE When output is flat: causes, costs and what to check first (long-form guide)

PILLAR The problems

QUESTIONS (11) COO framing x4: "Is it people, process, equipment or planning?", "Why does output stay flat when we're running flat out?", "What is the lost output worth each month?", "What should we check this week?"; COO trigger; plant manager trigger; COO options; COO scope: "What should the review cover?"; finance director options; assistant framing; assistant options

READERS COO first; plant manager; finance director skims the cost section; assistants quote the passages

THE ARC Recognise the problem in your own words, see the likely causes, put a number on it, check three things this week, see your options, take one small next step

THE REASON (echo, never sell) Savings are usually found on the floor, not in new equipment

PROOF Benchmark figures for flat-output plants, sourced and dated; one named plant

NEXT STEP The checklist; then the business case kit for readers who want to act

OUTLETS Website problem page; 10 quotable passages; newsletter issue; four social posts; checklist

GUARD Every figure from the benchmark or a named case; no claim Kestrel couldn't defend to a plant manager

Then each section gets its own brief. Two of the seven:

SECTION 2 Why output stays flat when you're running flat out

ANSWERS COO: "Why does output stay flat when we're running flat out?"

JOB Known · COO · familiar feels safe (name the causes they half-suspect, in their words)

FIRST SENTENCE The answer: the usual causes, most likely first

THEN Three causes, each with how it shows up on the floor and a benchmark figure

PASSAGES Each cause stands alone in under 80 words, so an assistant can quote it

GUARD No cause blamed on the plant's people; every figure sourced

SECTION 6 Your options: fix it in-house, buy equipment or bring someone in
 ANSWERS COO and finance director: "Fix it ourselves, buy equipment or bring someone in?"
 JOB Answered · COO and finance director · a hard choice stalls (make the choice simple, and fair)
 FIRST SENTENCE When each option is right, in one line each
 THEN What each costs and how long it takes; when buying equipment is the right answer
 LINKS Kestrel compared; Kestrel vs new equipment; How to run an operations review in-house
 GUARD Fair to every option, including the ones that don't need Kestrel

The flow check. Read the finished guide top to bottom against the arc. Then read each passage alone, as an assistant would quote it. A guide has to work both ways.

The finished example: "What flat output really costs", written in full

The cost article is the finance director's piece: it leads the finance director's trigger on flat output, and supports the COO's ad and the website's finance director page. It answers eight questions from the bank. Each section carries its job in italics, then the copy. Kestrel, its benchmark and its clients are invented.

Headline: What flat output really costs a mid-sized plant

The cost, in one sentence

Known · finance director · losses loom larger

In a mid-sized plant, flat output usually costs between 4% and 9% of annual contribution margin, almost all of it from capacity the plant already owns but doesn't use (Kestrel operations benchmark 2025, 38 plant reviews).

How it's made up

Answered · finance director · a hard choice stalls

The cost doesn't show up as a line in the accounts. It hides in three places:

- **Lost sales and late penalties.** Orders turned away or shipped late because the bottleneck line can't keep up. In the benchmark, 21 of the 38 plants had paid late-delivery penalties in the year before their review.
- **Overtime and weekend shifts** to make up for capacity lost during the week. The median plant ran 9% of its labour hours as overtime (benchmark).

- **Capital spent too early.** A new line bought to add capacity the old one was losing. Six of the 38 plants had a purchase approved or under way when their review started.

Work out your own

Believed · finance director · we follow others

You don't need our figures. Use yours:

```
Lost hours a week    = (best rate - average rate) / best rate x running hours
Contribution lost    = lost hours x units per hour at best rate x contribution
per unit
```

A worked example. A carton line runs 120 hours a week. Its best sustained rate is 47 cartons a minute; its average is 41. That's 13% of its time lost: about 15 hours a week. At 2,820 cartons an hour and 9 cents' contribution a carton, that's about A\$3,900 a week, or A\$200,000 a year, from one line. If the line is the bottleneck, every carton it doesn't make is one the plant doesn't sell.

What usually moves it

Answered · finance director and COO · good enough beats best

In the benchmark, three changes recovered most of the lost capacity: changeovers timed and owned (the largest gain in 23 of the 38 plants), small stops logged and fixed at the root, and the bottleneck line loaded first in the schedule. None needed capital. In the plants that made all three, the median gain was 11% more output from the same line within three months.

One next step

Bought · finance director · a hard choice stalls

If your number from the worked example is larger than you expected, the business case kit puts it in a form the board will accept, with the assumptions shown. [Build the case →]

What to notice. Every figure carries its source in its sentence, so an assistant can quote any one of them correctly. The worked example lets a finance director trust the method without trusting Kestrel. And the piece sells nothing until its last line, which is why a finance director forwards it, and why it can feed the ad, the finance director's page, a newsletter issue and two social posts without being rewritten.

Templates

The question bank (one row per question; the team's single list)

ID	Question, in the buyer's words	Moment (chooser · step · need)	Lead or support	Stage	Bias	Pillar	Piece that answers it	Answered in full?

The piece brief

PIECE [title, as the buyer would search it]
 TYPE [from the anatomy library] PILLAR [one of the pillars]
 QUESTIONS [the bank IDs it must answer in full]
 READERS [choosers and moments]
 THE ARC [first question to one next step, in one line]
 FIGURES [each one, with its source]
 CLAIMS USED [from the claims register, with state]
 OUTLETS [which builds use it, and as what]
 PRIORITY [P1-P3] FRESHNESS [evergreen / refresh yearly / dated]

The section brief

SECTION [name]
 THE QUESTION [bank ID, answered in the first sentence]
 JOB [stage · chooser and moment · bias]
 FIGURE + SOURCE [in the same sentence]
 WORDS [target]
 GUARD [passes only if it reads correctly on its own]

The outlet map (one row per core piece)

Core piece	Website page	Passages for being named	Newsletter issue	Social posts	Ad or landing page	Sales asset	Live?

Step 6: the business check

- ☐ **Revenue:** the two P1 guides and cost articles cover the problems behind most of Kestrel's work. If the plan's effort sits elsewhere, the question bank has drifted from the business.
- ☐ **One reason:** every piece echoes "savings found on the floor", without selling. A guide that reads like a brochure loses the reader and the assistant.
- ☐ **The proof is real:** the benchmark needs enough reviews behind it and clients' permission. Until then, the guides use named cases only.
- ☐ **The in-house guide is a decision, not a draft:** the partners agree that publishing the method is the strategy. If they don't, it drops to the improvement-lead articles alone.

- **Capacity:** one writer, two core pieces a month, plus the outlets. If that isn't there, write the two P1 guides and the benchmark first, and let the other builds wait for their raw material.

Step 7: the measures

Measure each piece by the questions it was written to answer, never by page views alone.

Measure	What it tells you	Stage
Share of the 149 questions answered in full	Whether the plan is built	Coverage
Assistant answers quoting a Kestrel passage, per guide	Whether the passages are good enough to quote	Found
First calls where the buyer uses a guide's words for their problem	Whether the framing landed	Known
Readers who go on to the checklist or the business case kit	Whether the guide moves people to the next step	Answered
Press pick-ups of the benchmark	Whether the evidence is travelling	Believed
Deals where the improvement lead is named as a supporter, from sales notes	Whether the in-house guide turned blockers into readers	Preferred
Outlets live per core piece	Whether the other builds are using the raw material	Coverage

In the campaign plan

Content is the raw material of almost every campaign, so it's planned first and released on the campaign's date. The rule: **a campaign uses a core piece that already exists, or it commissions one from the question bank.** It never commissions a piece the bank doesn't ask for.

Kestrel campaign	The core piece	What content supplies
The annual benchmark (Q2)	The benchmark report	The report, seven problem cuts, four sector cuts, the webinar deck, the press release

Kestrel campaign	The core piece	What content supplies
"Before you buy another line" (Q3)	"When output is flat" and "What flat output really costs"	The ad's figure and line, the landing page's sections, two newsletter issues
Board season (Q3 and Q4)	The seven cost articles	One cost per finance director's problem, ready to drop into the business case kit
The improvement leads' series	"How to run a review in-house" and the two practice articles	A three-issue newsletter series and the talk for the owners' round table

The year's content calendar is this table sorted by date. Nothing goes on it that isn't here.

Running it

Who owns it. At Kestrel, one writer owns the question bank and the core pieces; the head of delivery checks every guide for accuracy; the managing partner signs any new claim. Outlets belong to the builds that use them, but the writer keeps the outlet map, so nothing is rewritten without the core piece changing first.

When	What happens
Every two weeks	One core piece drafted, checked and published, with its outlets scheduled
Monthly	The question bank updated from sales calls and the scan: new questions added, answered ones ticked
Quarterly	Pieces re-read against the measures; any figure older than a year flagged
Yearly	The benchmark refreshed; every guide's figures updated from it the same month
When the map changes	New moments become new questions in the bank, before anyone suggests a topic

Using AI to write. A section brief is tight enough for an assistant to draft from, and the guard criteria make the draft checkable. The figures, the cases and the claims never come from the draft: they come from the benchmark and the claims register, and a person who knows the floor reads every guide before it's published.

Common mistakes

The mistake	The stage it breaks	The bias it ignores	On Kestrel's old blog	Instead
Starting from the calendar	Answered	Good enough beats best	Twelve posts; five answered no buyer's question, and none answered one in full	Start from the question bank; the calendar comes last
Writing about the firm	Known	Familiar feels safe	Company news, an award, two event write-ups	News goes in the newsletter; pieces answer buyers
Opinion with no figure	Believed	We follow others	"The future of manufacturing"	Every claim with a sourced figure in the same sentence
Thin posts on big questions	Found	Good enough beats best	600 words on equipment effectiveness: too thin to quote, too vague to act on	One guide that answers ten related questions in full
No cost on the problem	Known, for the finance director	Losses loom larger	Problems described, never priced, so doing nothing felt free	A cost article per problem, with a way to work out your own
Treating the in-house team as a rival	Preferred	Losses loom larger	Posts implying in-house teams fail	Help them: the in-house guide turns a blocker into a reader
Gating the guides	Found	A hard choice stalls	A white paper behind a form that nobody filled in	Guides open; only the benchmark's detailed data behind a form
Writing each outlet from scratch	Believed	Familiar feels safe	The ad, the post and the page each said something different	One core piece; every outlet cut from it
Judging by page views	All, unseen	Good enough beats best	The award post had the most views and won no work	Each piece judged by the questions it answers and what readers do next

The bias to watch in the team: **the planning fallacy**. Every content plan assumes the writer will hit two pieces a month and the reviews will be quick. Plan for the P1 pieces and the benchmark first, so that if the year runs slow, the pieces that matter exist.

The Loop for this build

Each quarter the Loop (step 11) reads the content measures beside the Score, and asks for each core piece: **did the questions it was written for get answered, and did anyone act on the answer?**

What the Loop finds	What it means	What changes
Quoted by assistants, but readers don't take the next step	The piece answers but doesn't lead	Rewrite the last section; check the next step fits the reader's moment
Read, but not quoted	The passages don't stand alone, or the figures lack sources	Rewrite to the anatomy library; put each source in its sentence
Neither read nor quoted	Nothing sends anyone to it, or it answers the wrong question	Check its outlets are live; check the questions against sales calls
New questions in sales calls	The bank has drifted from buyers	Add them to the bank; extend a piece before starting a new one
A figure is a year old	The proof is ageing	Refresh from the benchmark, or remove it

What Kestrel's first year would look for. The two P1 guides and cost articles quoted in assistant answers to the flat output and rising costs questions within two or three scans of publishing, as the two old blog posts already are. First calls where the COO describes the problem in the guide's words by mid-year. And after the benchmark, press pick-ups and the first finance directors arriving with their own number already worked out. If the guides are read but the Score's Known and Found don't move at framing, the guides aren't reaching the sources assistants read, and the benchmark's press plan is the place to look.

Shapes by business type

Pillars always come from the story spine. The formats change with the category:

Business type	Formats that are usually canon
B2B professional services	Guides, articles, research reports, case stories, webinars, newsletters
B2B software	Documentation, how-to guides, templates, comparison pages, webinars, release notes
Considered consumer (finance, health, education)	Explainers, calculators, guides by life stage, plain-language FAQs, video
Ecommerce	Buying guides, how-to video, size or fit guides, care guides, gift guides
Local services	Pricing guides, before-and-after stories, seasonal checklists, local FAQs
Marketplaces	Guides for each side, safety and trust guides, success stories

Feeling a bit overwhelmed?

You should be, a little. That was one firm: 77 moments, 149 questions, four pillars, 19 core pieces, 120 assets across five other builds, and ten months of writing for one person. Most content plans skip the question bank and start from the calendar, which is why most blogs look like Kestrel's old one.

RocMode does the heavy part for you:

- **Your buyer map, built for you.** RocMode maps your choosers and their moments, with the pain, gain, job and objection in each, and you can change any of it.
- **The gaps, shown.** It shows which moments nothing answers yet, so you know what has to be written before anything gets scheduled.

[Join the RocMode beta](#)

The Social Playbook

What this build is for

Social is proof of life, memory among the people who follow you, and reach through people. It is no longer a free distribution channel.

The free broadcast has largely gone. Socialinsider measured 25 million posts from 130,683 Facebook business pages over 2024 and 2025, and found an average engagement rate of 0.15%; link posts did worst, at 0.05% ([Socialinsider](#)). A company page that posts its news to its own followers is mostly talking to its staff.

So social has three jobs left, once the goal is being selected:

- **Proof of life** (Believed). Buyers check a feed the way they check a website: is this firm real, recent, and working with people like us?
- **Memory among the people who follow you** (Known), but only among people you already reach.
- **Reach through people** (Known). An engineer, a partner or a client reaches people a company page never will, and is believed more. Anatoli Colicev and colleagues tracked 45 brands in 21 sectors and found brands' own posts raised awareness but not purchase intent; what customers said raised both ([Colicev and colleagues](#)).

Video adds a fourth: it's transcribed, so what's said aloud on a factory floor can be read by search engines and assistants. That's Found and Answered, and for Kestrel it may matter most.

Social is in the Owned category, with the website and content: your voice, in your own and your people's channels. Its rule is the same as every build's: **a post exists because a moment needs to be remembered or proved**. "We should post more" isn't a reason. Reach to new buyers at scale is paid, and belongs to Advertising. Kestrel's people posting about their own work is in this build; a formal employee advocacy programme is Partners and channels' nod.

The footprint, and what it reads from the chain

Social leads **1 cell and 7 moments**: the plant manager's trigger, once for each of the seven problems. The plant manager's question is "We're always firefighting, and nobody fixes the cause." It isn't yet a question about consultants. It's a complaint, made to colleagues, in forums and in comments, and often at the end of a long shift.

Moment	What the plant manager is firefighting
Plant manager · trigger · flat output	Lines stopping, changeovers running long, overtime every week
Plant manager · trigger · rising costs	Being asked to cut costs with no idea where they're going
Plant manager · trigger · cash in stock	Expediting parts that should already be there
Plant manager · trigger · scrap and quality	Rework benches, sorting, a customer's auditor on site
Plant manager · trigger · new line ramp-up	A new line that keeps stopping, and a supplier who says it's the operators
Plant manager · trigger · after an acquisition	A second plant to run, with different habits
Plant manager · trigger · automation not paying back	A new machine starved of parts

It supports **4 more cells (16 moments)**:

- **COO · trigger** for the seven problems (Advertising leads): an engineer's post from a client's floor reaches COOs through the people they follow.
- **COO · framing** for the seven problems, "Is it people, process, equipment or planning?" (Content leads): one-question videos, spoken and captioned, give assistants something to quote.
- **Plant manager · framing**, "Will they just tell us what we already know?" (Creators leads): Kestrel's engineers seen working on a floor, beside a plant's own team.
- **Improvement lead · framing**, "We've tried lean before. Why didn't it hold?" (Content leads): a client's own improvement lead saying what made it hold this time.

The plant manager is the chooser this build exists for. They live with the result, and can make or break it on the floor. They rarely search for consultants and never read a firm's brochure, but they watch other plant managers' videos and read their posts.

What it reads, and never rewrites:

From	It takes
Brand strategy	The choosers and their words; the reason
Brand identity	The voice, and the distinctive assets: the amber floor tape and the line

From	It takes
Credentials	Every number a post may use, and every client who's agreed to appear
Content	The guides and the benchmark, which posts point to and cut down
Customers	The clients and their people who'll appear, with permission
Partners and channels	The employee advocacy nod: how Kestrel's people post in their own names

Its slice of the Score, and its own audit

Kestrel today	Known	Found	Quality	Moment score
Plant manager · trigger, each of 7 problems (out in the market)	6	0	9	0
COO · trigger, each of 7 problems (out in the market)	14	0	12	0
COO · framing · flat output (asked of machines)	14	8	22	2
COO · framing, the other 6 problems (asked of machines)	14	0	9	0
Plant manager · framing (out in the market)	6	0	9	0
Improvement lead · framing (out in the market)	6	0	9	0

Of 23 moments, one scores above zero. Plant managers are among the choosers who know Kestrel least (Known 6). The single bright spot is the COO's framing question about flat output: assistants named Kestrel in a few answers, citing an old blog post. Nothing Kestrel has posted on social appears in any answer.

The audit. What Kestrel's social did last year:

Channel	What Kestrel did	What it produced
LinkedIn company page	1,100 followers, mostly staff, alumni and partners. About three posts a month: awards, the trade-show stand, "we're hiring"	Average engagement 0.3%; no enquiry anyone could trace
YouTube	Four videos: a firm overview and three event recordings. The newest is 14 months old	610 views in total
Partners' own posts	The managing partner posts every few weeks, mostly firm news	Some comments from clients
Engineers' own posts	2 of 35 engineers post at all; neither about their work	—
Clients' posts	Brennock's improvement lead once posted about the review, unprompted	The most-commented post about Kestrel last year

The feed passed none of the three tests. It proved Kestrel was alive, but not that it worked on floors. It wrote no moment into anyone's memory. And it reached no one through people, except once, by accident, through a client.

Step 1: the shape, where the choosers already are

The shape is **the places Kestrel's choosers already spend time, and what they do there.** It's set by them, not by Kestrel.

Platform	Who's there	What they do there	Who the machines read	Kestrel today
LinkedIn	COOs, owners, improvement leads; some plant managers	Read people they know; scroll the feed on the way to work	Little; posts are partly closed to crawlers	Company page, three posts a month
YouTube	Plant managers, improvement leads, engineers	Search how-to videos: changeovers, root cause, line balancing	Yes: transcripts and titles are widely read and cited	Four videos, none recent

Platform	Who's there	What they do there	Who the machines read	Kestrel today
Industry forums and Reddit's manufacturing communities	Plant managers and engineers, often anonymous	Ask each other what's gone wrong, and who's any good	Yes: forum threads are often cited by assistants	Absent
Trade-publication comment sections and newsletters	COOs and plant managers	Read, rarely comment	Partly	Absent
Facebook, Instagram, TikTok	Everyone, personally	Not work	Some	Absent, rightly
X	A few journalists and analysts	Follow news	Some	Absent

Two things stand out. **Plant managers learn on YouTube and ask in forums**, and both are places assistants read. And **LinkedIn reaches COOs through people, not pages**: a post from an engineer they've met, or a client they know, goes further than anything on the company page.

Forums, carefully. A firm posting about itself in a forum is spam, and the forum will say so. Kestrel's engineers can take part in their own names, saying who they work for, answering questions in their field and never linking to a sales page. That's slow, and it's the only version that works.

Step 2: coverage, one question per moment

The coverage unit is **one question, answered on a floor**: the question a firefighting plant manager would type into YouTube, answered in under three minutes by someone who's fixed it, standing where it happened. Each answer becomes a video, a post, and a short cut for the feed.

Moment	The question, in the plant manager's words	Who answers	Where
Flat output	"Why does our line stop every time we change over?"	Priya Raman, Kestrel's lead engineer	Brennock Packaging's line 3

Moment	The question, in the plant manager's words	Who answers	Where
Rising costs	"Where does the money go on a line that looks busy?"	A Kestrel engineer, with the benchmark's finding	A client's floor
Cash in stock	"Why are we short of parts with a full warehouse?"	A Kestrel engineer	A client's stores
Scrap and quality	"Why does more inspection never fix scrap?"	Calloway Foods' quality lead, with a Kestrel engineer	Calloway's floor
New line ramp-up	"Our new line won't reach rate. Is it the machine or us?"	A Kestrel engineer	A client's new line
After an acquisition	"How do you run two plants with different habits?"	A client plant manager who's done it	Their floor
Automation not paying back	"Why is our new robot standing idle?"	A Kestrel engineer	A client's cell
COO · framing (×7)	"Is it people, process, equipment or planning?" for each problem	The same videos, with the answer said aloud in the first 20 seconds	As above
Plant manager · framing	"Will they just tell us what we already know?"	A client plant manager, on what Kestrel found that they hadn't	Their floor
Improvement lead · framing	"We've tried lean before. Why didn't it hold?"	Brennock's improvement lead, who ran the next review himself	Brennock

The COO's framing moments need no separate videos. Each problem video answers the COO's question too, provided the answer ("it's usually process, not people") is said aloud early, so the transcript carries it. **Seven videos cover 14 moments, and reach the COO's 7 trigger moments through the people who share them.** The two framing moments for plant managers and improvement leads need clients on camera, and those two videos will be the most believed.

Alongside the videos, the feed needs a steady rhythm of smaller things: a number from the benchmark, a photo from this week's floor (with the client's permission), a client's own words. That's proof of life, and it's the part most firms skip.

Step 3: gaps, biases and jobs

Moment	Stage	Bias	The job
Plant manager · trigger, all seven	Known	We follow others	Seven one-question videos on YouTube, filmed on clients' floors, one a month for seven months; each cut down for LinkedIn
Plant manager · trigger	Known	Familiar feels safe	Kestrel's engineers post monthly in their own names: what they fixed this month and how, with the client's permission
COO · trigger	Known	We follow others	The managing partner and the head of delivery post weekly; clients' own posts about their results, encouraged and thanked, never scripted
COO · framing, all seven	Found, Answered	Good enough beats best	The answer said aloud in the first 20 seconds of each video; titles in the buyer's words; captions and chapters; a full transcript on the matching guide
Plant manager · framing	Believed	We follow others	A client plant manager on camera: "what they found that we hadn't"
Improvement lead · framing	Believed	Losses loom larger	Brennock's improvement lead on why it held this time: his team ran the next review
Anyone checking Kestrel's feed	Believed	Familiar feels safe	Proof of life: something from a real floor every week, never three weeks of silence

Moment	Stage	Bias	The job
Forums	Believed	We follow others	Two engineers answer questions in their field in their own names, a few times a month, saying who they work for. No links to sales pages

One pattern: **almost nothing here comes from the company page.** It comes from Kestrel's engineers, partners and clients, speaking for themselves. The company page collects it, so anyone who checks finds a firm that's busy on real floors.

Step 4: reconcile, the social plan

Social costs time more than money, so the plan is set in hours.

Line	What	How often	Time	Moments	Signals	Priority
One-question videos	Seven problem videos, then two client framing videos; on YouTube, cut down for LinkedIn	One a month	A day each: half a day filming, half editing (freelance editor)	Plant manager trigger; COO framing; plant manager and improvement lead framing	Coverage, canon	P1
Partners' posts	The managing partner and the head of delivery, in their own names	Weekly	30 minutes each	COO trigger	Strategy	P1
Engineers' posts	What they fixed this month, and how, with the client's permission	Monthly, for the 8 engineers who want to	30 minutes each	Plant manager trigger	Coverage	P1
Proof of life on the company page	A floor photo, a benchmark number, a client's words, a video	Weekly	An hour a week, the marketing lead	Anyone checking	Strategy	P2
Forums	Two engineers answer questions in their field	A few times a month	An hour a week each	Plant manager trigger	Coverage	P2

Line	What	How often	Time	Moments	Signals	Priority
Clients' posts	Thank and share clients who post about their results; never script them	When they do	—	COO trigger; improvement lead framing	Strategy	P2

The only spend is the freelance editor and a small camera kit: about A\$9,000 a year. The rest is about seven hours of Kestrel people's time a week, across the firm.

Rejected, with reasons:

Posting daily on the company page. It reaches mostly staff. Weekly proof of life is enough, and better.

Ready-written posts for staff to share. Ten engineers posting the same paragraph tells everyone who wrote it. Kestrel's people post their own work in their own words, or not at all.

Boosting posts. Paying to push a post is advertising, and Advertising has a plan and a test for it.

Instagram and TikTok. Kestrel's choosers aren't choosing there. The videos are filmed vertically as well, in case that changes.

Award logos and anniversaries. They prove Kestrel exists, and nothing else. One a year, at most.

Chasing a viral hit. Sharad Goel, Duncan Watts and colleagues studied a billion diffusion events on Twitter: moderately popular events came up about one in a thousand times, "viral hits" closer to one in a million ([Goel and colleagues](#)). A schedule Kestrel can keep for a year beats a hit it can't plan.

The anatomy library

Every post needs the first two parts, and at least one of the next three. The feed as a whole does the rest.

Part	Stage	What it does
A moment the buyer recognises	Known	"The line stopped again at changeover" stops a plant manager. "We're excited to share" stops nobody

Part	Stage	What it does
A distinctive look, the same every time	Known	The amber floor tape in frame; the same opening line: "One question, from the floor"
Something useful	Answered	The viewer leaves knowing one thing they didn't
A real person	Believed	An engineer, a partner or a client, in their own words
The reason, in passing	Preferred	"That's what six weeks on the floor is for." Said once, lightly, across fifty posts
Spoken and captioned	Found	The question, the answer and Kestrel's name said aloud, so the transcript carries them
The feed as a whole: proof of life	Believed	Something from a real floor every week

On "spoken and captioned": Ahrefs found YouTube mentions correlated with AI visibility at about 0.737 for ChatGPT, the strongest factor in its study, and warns that "correlation isn't causation" ([Ahrefs](#)). It's a reason to say things aloud, not a promise.

By format:

Format	Parts, in order	Length
The one-question video	The question, on screen and said aloud, in the first 3 seconds · "One question, from the floor", with the amber tape · The short answer, in the first 20 seconds · What it looked like at this plant, with the number · What to check on your own line · Who's speaking, and where	2 to 3 minutes; captions and chapters
The LinkedIn cut	The question · The short answer · The number · A link to the full video in the first comment	45 to 60 seconds, vertical and square

Format	Parts, in order	Length
The engineer's post	What was wrong, in the plant manager's words · What they tried · What worked, with a number if the client agrees · One photo from the floor · Who they work for, in their profile	Under 150 words
Proof of life	One thing from a real floor this week · Whose floor (with permission) · One line on why it matters	Under 60 words and a photo
The forum answer	The answer to what was asked · Why, from experience · "I work for Kestrel, an operations consultancy" if it's relevant · No link unless asked	As long as the question needs

Three rules hold for every post:

Clients appear only with permission, in writing, for that post. A client's floor is their property, and so is their scrap rate.

People post in their own words. Marketing can suggest a subject and check a number; it never writes someone else's post.

Every number is one Credentials has signed. A plant manager will ask how it was measured.

Briefs

One brief per video, agreed with the client before filming. This is the brief behind the finished example.

VIDEO BRIEF	One question, from the floor: no. 1
THE QUESTION	"Why does our line stop every time we change over?"
MOMENTS	Plant manager · trigger · flat output (lead) COO · framing · flat output: "Is it people, process, equipment or planning?" (answered aloud in 20 seconds)
WHO ANSWERS WITH	Priya Raman, Kestrel's lead engineer Brennock Packaging's improvement lead, who ran the next review himself (on camera, his choice of words)
WHERE	Brennock's line 3, at a changeover, with permission
THE ANSWER	It's usually process: the changeover waits for parts, tools and a sign-off that could all be ready before the line stops
THE NUMBER	Changeovers from 52 minutes to 19; line speed from 41 to 47

CHECK THIS	a minute. Signed by Credentials, confirmed by Brennock
ASSETS	Time your next changeover from last good part to first good part, and list everything the crew waited for
NOT ALLOWED	Amber floor tape in shot; "One question, from the floor"
OUTPUTS	Any Brennock figure not in the signed case; faces of staff who haven't agreed; a sales line at the end
	YouTube (3 min); LinkedIn cut (50 s); Priya's own post; the transcript on the flat-output guide

The finished example: "One question, from the floor", no. 1

The video, its YouTube listing, the LinkedIn cut, Priya's own post, and the company page's proof of life, for one moment. Kestrel, Brennock and the people are invented.

1. The video script, 2 minutes 50 seconds, filmed at Brennock Packaging's line 3.

Time	Picture	Said aloud
0:00	Line 3 stopped. A crew waiting. On screen: "Why does our line stop every time we change over?" Amber tape on the floor, in frame.	Priya: "Why does our line stop every time we change over? It's the question I get asked most. One question, from the floor."
0:08	Priya at the line, the Brennock sign behind her. Caption: Priya Raman, lead engineer, Kestrel Advisory, for manufacturers.	"The short answer: it's usually not the people, and it's rarely the machine. It's the process around the changeover. The line stops, and then everyone starts getting ready."
0:25	The crew walking to the stores; a forklift fetching a die; someone looking for a sign-off.	"Here, before the review, the crew waited for the next die, the right tools, and the quality sign-off. All of it could be ready before the line stopped."
0:55	A whiteboard: the changeover, split into "line stopped" and "could be done while running".	"We split every task into two lists: what has to happen with the line stopped, and what can happen while it's still running. About two-thirds was on the wrong list."
1:30	Brennock's improvement lead at the line, speaking in his own words.	Improvement lead: "We'd tried to fix changeovers before by pushing people to go faster. This time we moved the work, not the people. We ran the next review ourselves."

Time	Picture	Said aloud
2:00	The changeover now: a trolley with the die, tools and paperwork waiting beside the line. On screen: 52 minutes → 19 minutes. Line speed 41 → 47 a minute.	Priya: "Changeovers went from 52 minutes to 19. The line now runs at 47 a minute, up from 41. Same line, same people."
2:25	Priya, to camera.	"If you want to check your own: time your next changeover from the last good part to the first good part, and write down everything the crew waited for. That list is where your capacity is."
2:45	The amber tape; the name.	"I'm Priya, from Kestrel Advisory. That's one question, from the floor."

2. The YouTube listing

Title: Why does our line stop every time we change over? | One question, from the floor

Description: Changeovers at Brennock Packaging went from 52 minutes to 19, with the same line and the same people. Priya Raman, lead engineer at Kestrel Advisory, explains why changeovers usually run long (it's process, not people or the machine) and how to check your own line this week. With Brennock's improvement lead.

Chapters: 0:00 The question · 0:08 The short answer · 0:25 What the crew waited for · 0:55 Two lists · 1:30 Why it held · 2:00 The result · 2:25 Check your own line

Captions: edited by hand, not left to auto-captions, so "Kestrel" and "changeover" are spelt right in the transcript.

3. The LinkedIn cut, 50 seconds, on Kestrel's page and shared by Priya.

"Why does our line stop every time we change over?" At Brennock Packaging, two-thirds of the changeover work could be done before the line stopped. Changeovers went from 52 minutes to 19. Full video in the comments.

4. Priya's own post, in her words, the same week.

Spent a lot of this year at Brennock on line 3. The thing that fixed their changeovers wasn't a new die cart or a faster crew. It was splitting every task into "line stopped" and "line running" and finding most of it was on the wrong list.

Their improvement lead ran the next review himself, which is the bit I'm proudest of. We filmed a short video on it with him (link in comments).

If you run a line: time your next changeover from last good part to first good part, and write down what the crew waited for. Happy to compare notes.

5. Proof of life, on the company page, a fortnight later.

This week on Brennock's line 3: the changeover trolley, loaded before the line stops. Same crew, 19 minutes. (Shared with Brennock's permission.)

What to notice. The video opens on the plant manager's question, in their words, and answers the COO's framing question aloud inside 20 seconds, so the transcript carries it to search and assistants. The client's improvement lead speaks for himself, and what he says ("we ran the next review ourselves") answers another chooser's fear without anyone writing it for him. Priya's post doesn't mention Kestrel's services at all; her profile does that. The one number is the same number the website, the ads and the pitch use. And none of it asks anyone to get in touch.

Templates

1. The feed audit scorecard. Score the last 12 weeks of a feed, yours or a rival's, 0, 1 or 2 on each line. Under 8 means start with proof of life and real people before anything else. The scale is a working rule, not tested against sales. The last two columns show Kestrel's company page today and the target for six months from now.

Part	Stage	Question	0	1	2	Today	Target, 6 months
Moments	Known	Do most posts open on a situation a buyer would recognise?	Hardly any	Some	Most	0	2
A distinctive look	Known	Would a scroller know it's you before reading?	No	Sometimes	Always	0	2
Something useful	Answered	Does the viewer learn something in most posts?	Rarely	Sometimes	Usually	0	2

Part	Stage	Question	0	1	2	Today	Target, 6 months
Real people	Believed	What share of posts feature a client, an employee or a partner in their own words?	Under 10%	10% to 40%	Over 40%	0	2
The reason, in passing	Preferred	Is the reason there, lightly, across the feed?	Never	Once or twice, heavily	Often, lightly	0	1
Spoken and captioned	Found	Do videos say the question, the answer and the name aloud, with edited captions?	No video	Some	All	0	2
Proof of life	Believed	Was there something from real work every week?	Gaps of a month or more	Every fortnight	Every week	1	2
Total						1/14	13/14

2. The monthly rhythm. One month of Kestrel's feed, repeated with new subjects.

Week	Company page	People, in their own names
1	The month's one-question video (LinkedIn cut, linking to YouTube)	The engineer who answered it posts in their own words
2	A client's own words: a quote, a post shared, a photo of their floor	The managing partner: what clients asked about this month
3	A number from the benchmark, with what to check	Engineers' monthly posts (a few each week, not all at once)
4	Proof of life: this week on a floor	The head of delivery: one lesson from a review
Every week	—	Two engineers answer forum questions in their field

3. The one-question video brief.

VIDEO BRIEF	One question, from the floor: no. __
THE QUESTION	[in the chooser's words, as they'd type it into YouTube]
MOMENTS	[lead moment; any framing question answered aloud]
WHO ANSWERS	[the person who fixed it]
WITH	[a client, in their own words, if they agree]
WHERE	[the floor where it happened, with permission]
THE ANSWER	[one sentence, said within 20 seconds]
THE NUMBER	[signed by Credentials, confirmed by the client]
CHECK THIS	[what a viewer can check on their own line this week]
ASSETS	[tape in shot; the opening line]
NOT ALLOWED	[unsigned figures; unconsenting faces; a sales line]
OUTPUTS	[YouTube length; LinkedIn cut; people's own posts; transcript home]

In the campaign plan

Social runs on its monthly rhythm all year. Campaigns lend it subjects, and it lends campaigns real people.

Campaign	What social does	Moments	What other builds bring
The annual benchmark (Q2)	Four sector cuts as one-question videos ("Where does the money go in a food plant?"); the head of delivery's post on the finding	COO trigger and framing; plant manager trigger	Content: the benchmark; PR: coverage; Advertising: the benchmark ad
"Before you buy another line" (Q3)	The ramp-up and automation videos, timed for budget season; a client plant manager on what they fixed instead of buying	Plant manager trigger (ramp-up, automation)	Advertising: extra weight; Website: the "before you buy" page
Owners' round tables (spring and autumn)	Photos and the speaker's own post afterwards, with permission; nothing before	Owner options (Events leads)	Events: the room

Campaign	What social does	Moments	What other builds bring
Plant visits (twice a year)	Short video from the host plant, filmed with the host's plant manager	Plant manager framing	Customers: the host; Events: the visit
A client's results review	Ask the client if they'd like to post about it; share and thank them if they do	Improvement lead framing; COO trigger	Customers: the review

Running it

Who owns what. At Kestrel, the marketing lead owns the rhythm, the company page and the video schedule. The head of delivery owns which client floors are filmed, and asks the clients. Each engineer who posts owns their own posts; nobody approves them except for client names and numbers, which Credentials checks. A freelance editor cuts the videos and edits the captions.

When	What happens
Monthly, week 1	One question filmed and published; the engineer posts
Weekly	Proof of life on the company page; partners' posts; forum answers
Monthly	Engineers' posts, staggered
Before every post that names a client	Written permission for that post; the number checked against the signed case
Quarterly	The feed audit scorecard; new customers asked "had you seen Kestrel on LinkedIn or YouTube, or checked our feed?"
Yearly	The video list against the moments: which questions got watched, which got cited by assistants, which to film next

Helping people post, without writing for them. Engineers who want to post get an hour's session with the marketing lead: how to pick a subject, how to get permission, what not to say about a client. They're offered subjects each month ("what surprised you on a floor this month?"). They're never sent text to copy.

When a client says no. Many will, especially to filming. Ask once, accept the answer, and film on the next client's floor. Never film around them.

Common mistakes

The mistake	The stage it breaks	The bias it ignores	At Kestrel	Instead
The brand's voice alone	Believed	We follow others	The company page praising the company	Engineers, partners and clients, in their own words
News nobody asked for	Known	Feeling first, reasons after	Awards, anniversaries, the trade-show stand	Open on a moment a plant manager would recognise
Silence, then a burst	Believed	Familiar feels safe	14 months without a video	Something from a real floor every week
Counting followers	All, unseen	Good enough beats best	"Followers up 12%"	Feed audit score; new clients who saw or checked the feed
Same words, ten accounts	Believed	We follow others	—	People post their own work, or nothing
Video nobody can read	Found	Good enough beats best	Auto-captions spelling "Kestrel" three ways	Say the question and the name aloud; edit the captions
A sales line at the end	Preferred	Losses loom larger	—	End with what to check on your own line
A client filmed without asking	Believed	Losses loom larger	—	Written permission, for that post, every time

The mistake	The stage it breaks	The bias it ignores	At Kestrel	Instead
Chasing a viral hit	Known	Good enough beats best	—	A rhythm you can keep for a year

The bias to watch in the team: **the false-consensus effect**. We assume others care about what we care about. A firm is proud of its award, its new hire, its tenth birthday, and assumes its buyers will be too. Plant managers care about why their line stops. The feed should be about that.

The business check and the measures

Revenue: plant managers can make or break a review, and the improvement lead's fear is often the hidden blocker. A plant manager who's watched Priya on a floor arrives at the first meeting less worried. That shortens deals more than it wins them.

One truth: every number in a video or post is the signed number, word for word.

The proof is real: clients appear in their own words, with permission, saying what went wrong as well as what worked.

Capacity: about seven hours a week across the firm, and about A\$9,000 a year for editing and kit.

On the measurement ladder: followers, likes, views and reach stay on the diagnostic rung. Steer by these:

Measure	What it tells you	Stage
Feed audit score, quarterly	Whether the feed proves life, writes moments and shows real people	Believed, Known
Share of posts featuring a client, an engineer or a partner in their own words	Whether reach is through people	Believed
New clients who saw Kestrel on social, or checked the feed before calling	Whether social put Kestrel in the room	Known, Believed
Assistants citing or quoting Kestrel's videos for the framing questions	Whether spoken answers reach the machine	Found, Answered

Measure	What it tells you	Stage
Plant managers who mention a video in the first meeting	Whether the plant manager's fear was met before the deal	Believed
Watch time to the answer (past 20 seconds)	Whether the question was the right one	Answered

The Loop for this build

Each quarter, the Loop (step 11) asks: **did the people who watched and read move closer to choosing, and did the machines pick up what was said?**

What the Loop finds	What it means	What changes
Videos watched past the answer, but no assistant cites them	The machine can't read them, or the question isn't the one buyers ask assistants	Check captions and titles; match the title to the Score's question wording
Assistants quote a video for a framing question	The spoken answer is reaching the machine	Film the next problem the same way; add the transcript to the guide
Engineers' posts get more response than the company page	Reach through people is working	Fewer company posts, more support for people's own
Clients name the videos in first meetings	Social is meeting the plant manager's fear early	Film more client floors
The feed audit score slips	The rhythm has broken	Back to one video a month and proof of life every week

What Kestrel's first year would look for. Seven one-question videos and two client framing videos published. Assistants naming Kestrel for at least two of the COO's framing questions, against one today. A feed audit score of 10 or more every quarter. And a third of new clients saying they'd seen Kestrel's people on LinkedIn or YouTube before they called.

Shapes by business type

Business type	Typical mode	The social that usually matters most
B2B professional services	Committee or owner-led	Partners' and practitioners' own posts; short answers to buyers' questions on video; a company page that proves life weekly
B2B software	Committee	Founders and engineers posting their own work; customers in video series; product videos that answer one question each
Considered consumer (finance, health, education)	Considered personal	Real customers and practitioners answering real questions; strict care with claims and privacy
Ecommerce	Quick choice or habitual	A distinctive look in every post; customers' own posts and reviews; the platforms where the category is browsed
Local services	Owner-led or considered personal	Proof of life above all: this week's jobs, with the customer's permission; local groups, as a neighbour, not an advertiser
Marketplaces	Varies on each side	Sellers' or providers' own stories for the supply side; buyers' stories for the demand side

Feeling a bit overwhelmed?

You should be, a little. One firm: the platforms its choosers already use, a question per moment answered on a client's floor, captions edited by hand, written permission for every client name, eight engineers posting in their own words, two answering in forums, partners posting weekly, proof of life every week, and a feed scored every quarter. Most firms post their awards to their own staff and count the likes.

RocMode does part of the heavy lifting:

The questions worth answering. Your choosers and their moments, mapped from what's known about your brand, so every video starts from a question a buyer actually asks.

Posts drafted for approval, on a campaign calendar. RocMode drafts LinkedIn and Facebook posts from your strategy and lays them on a campaign calendar. Each draft waits for your approval, and your people's own posts stay theirs to write.

[Join the RocMode beta](#)

The Advertising Playbook

What this build is for

An ad has one job: to tie your name to a moment in the buyer's memory.

Most people who see an ad aren't buying, and won't be for months. About 95% of business buyers are out of the market at any one time ([Ehrenberg-Bass](#)). So an ad's work, for nearly everyone who sees it, is to leave Kestrel's name filed under a moment: the morning a key customer complains, the month the margin line turns red. When that moment arrives, the COO should already know who to ask. That's Known.

Its second, smaller job is to put Kestrel in front of someone at the moment they look: a search ad, a top listing. That's Found, bought rather than earned. It lasts only while you pay.

The book calls these **writing** the memory and **collecting** on it. Brand advertising writes. Search ads, retargeting and offers collect. Les Binet and Peter Field's rule of thumb, from 996 campaigns in the IPA's databank, puts about 60% of the budget into writing and 40% into collecting ([IPA](#)). Treat it as a starting point. Collecting can't work on a memory nobody wrote.

Advertising is the Paid category's only build: your voice, in space you buy. Its rule is the same as every build's: **an ad exists because a moment needs to be remembered**. Being in the trade magazine because rivals are isn't a reason.

Four channels don't need a build of their own, and live here as **nods**:

- **Sponsorship:** paying for your name on something the category already values, such as an association's report, an award or a conference session.
- **Retail media:** ads inside a retailer's or marketplace's own site, where buyers are already shopping.
- **Broadcast and outdoor:** radio, television, billboards and transport.
- **Direct mail:** printed pieces posted to many people at once. (A letter to a named person, for a reason, belongs to Outbound.)

Each is judged by the same test as any ad: does it write a moment into the right memory, for less than the alternatives?

The footprint, and what it reads from the chain

Advertising leads **1 cell and 7 moments**: the COO's trigger, once for each of Kestrel's seven problems. The COO's question is "Output's flat and a key customer is complaining. What's going on?", and the problem changes what they've just seen:

Moment	What the COO has just seen	How they'd say it
COO · trigger · flat output	A key customer's complaint about late deliveries	"We're running flat out and still shipping late."
COO · trigger · rising costs	The monthly accounts, margin down again	"Margins fell again and we can't see why."
COO · trigger · cash in stock	The finance director asking about the stock figure	"The warehouse is full and we're still short of the right parts."
COO · trigger · scrap and quality	A failed customer audit, or a quality claim	"We're paying to make things twice."
COO · trigger · new line ramp-up	The new line's counter, still below rate	"We spent the money and the line still isn't at rate."
COO · trigger · after an acquisition	The new owner's plan, with a margin target	"We need to show margin improvement within 18 months."
COO · trigger · automation not paying back	A new machine standing idle, waiting for parts	"We automated, and nothing got better."

It supports **2 more cells (14 moments)**: the owner's trigger for the same seven problems (Outbound leads, by name; advertising reaches the owners Outbound doesn't write to) and the finance director's trigger, "Margins fell again. What's driving the cost?" (Content leads, with the benchmark; advertising carries its finding to finance directors who'd never look for it).

That's 21 moments, all at the trigger, all out in the market. **Advertising is the only build whose whole footprint sits before the buyer starts looking.** It does its work months before anyone in the deal record has heard of Kestrel.

What it reads, and never rewrites:

From	It takes
Brand strategy	The choosers, the reason ("savings found on the floor in six weeks, at a fixed price") and the distinctive assets
GTM strategy	The segments, which become the audience: COOs, owners and finance directors at the 1,400 firms
Credentials	Every number an ad may use, and the client who said it
Content	The benchmark and the problem guides, which ads point to
Website	The page each ad sends people to, saying the same thing in the same words
Outbound	The 40 named accounts, for the account plan
Events	The round tables and talks, which a few ads invite people to

Its slice of the Score, and its own audit

Kestrel today (out in the market)	Known	Found	Quality	Moment score
COO · trigger, each of 7 problems	14	0	12	0
Owner · trigger, each of 7 problems	12	0	9	0
Finance director · trigger, each of 7 problems	8	0	9	0

All 21 moments score zero. Some COOs have heard of Kestrel (Known 14), but none of the searches or places the Score checked named it at the trigger, so the gate is shut. Finance directors know it least of all. Quality is low for the same reason as in Outbound: nothing answers the trigger question in the buyer's terms, and the next step is a contact form. Believed and Preferred weren't measured at the trigger, so they're left out of the average.

Advertising can lift Known directly, and it's the build that can do so at the scale of the whole market: the Score's Known is measured across the market, and so is an ad's reach. That makes this the build where the Score and the build's own measures should move together, if the ads work.

The audit. Last year's spend, A\$48,000:

Line	Spend	What it did	What the report said	What it actually showed
Search ads on "Kestrel"	A\$14,000	Ads on Kestrel's own name	31 conversions	Many clicks were people looking for Kestrel Mining Services; nobody else bids on the name
Search ads on "operations consultant"	A\$10,000	Ads on broad terms	2 enquiries	Clicks mostly from job-seekers and students; neither enquiry was a fit
Retargeting	A\$8,000	Ads to past site visitors	9 conversions	7 of the 9 were people already in a deal
LinkedIn sponsored posts	A\$10,000	Three bursts, around the trade show	410,000 impressions; 0.4% click rate	Nobody could say what changed. The posts opened on Kestrel, not on a moment
Trade magazine	A\$6,000	Two half-page ads: logo at the bottom, a list of services	—	Nothing measured

About two-thirds of the money went to collecting (A\$32,000), and a third to something like writing (A\$16,000), the reverse of the starting point. The collecting ads mostly paid for people who were coming anyway, or who wanted someone else. **No ad opened on a moment, and none was ever tested.**

Step 1: the shape, the category's media

The shape is **the media the category's buyers already spend time in**, and who advertises there. It's set by the buyers, not by Kestrel.

Medium	Who's there	Writes or collects	Who in the category uses it	Kestrel today
LinkedIn, targeted by role and firm	COOs, operations directors, owners, finance directors, at the 1,400 firms	Writes	Large consultancies, equipment makers, software firms	Three bursts a year
Trade newsletters and magazines (print and email)	COOs and plant managers; owners in smaller firms	Writes	Equipment makers mostly; few consultancies	Two half-page ads
Association newsletters and member reports	Owners and COOs, by sector	Writes	Banks, insurers, accountants	None
Search ads on problem terms ("reduce scrap rate")	Anyone already looking: COOs, improvement leads, students	Collects	Software firms, training providers	None
Search ads on the firm's name	People who already know the name, and people who mistake it	Collects	Everyone, by habit	A\$14,000
Retargeting	Past site visitors	Collects	Everyone, by habit	A\$8,000
Industry podcasts and video	Improvement leads, some COOs	Writes	Few	None
Sponsorship (nod)	Whoever the sponsored thing reaches	Writes, if the name is attached to something useful	Banks and insurers sponsor association reports and awards	None
Broadcast and outdoor (nod)	Everyone near the billboard or radio; few of them buyers	Writes, expensively	Equipment makers at trade shows	None

Medium	Who's there	Writes or collects	Who in the category uses it	Kestrel today
Direct mail (nod)	Everyone on a list	Writes, if kept	Rare	None
Retail media (nod)	Shoppers on a retailer's own site	Collects	Not used in this category	Not relevant

Share of voice. Kestrel is about 4% of the market for operations reviews at firms this size. Its share of the category's advertising is closer to 2%, most of it on its own name. The IPA's rule of thumb: brands whose share of voice is below their share of market "tend to shrink" ([IPA](#)). Kestrel can't outspend the large consultancies everywhere. It can outspend them in one place: in front of the 1,400 firms' COOs, owners and finance directors, on LinkedIn and in the trade and association titles, all year.

Retail media, as a nod. It's ads inside a retailer's or marketplace's own site, where shoppers are mid-purchase. Nobody buys an operations review from a retailer, so Kestrel has none. For brands that sell through retailers, it's collecting at its most direct, and the same warning applies: test whether it adds sales or only claims them.

Step 2: coverage, one ad per moment

The coverage unit is **an ad per moment**: the cue the buyer sees in the first three seconds, the feeling it leaves, and the one thing to remember. If an ad has to serve two moments, it's two ads. The one thing to remember is the same in every ad, because it's Kestrel's reason: *savings found on the floor in six weeks, at a fixed price*. What changes is the moment it's filed under.

Moment	The cue, in the first three seconds	The feeling	The proof (signed by Credentials)
Flat output	A phone at 6am: a customer's email, "Where's our order?"	Dread, then relief: it's findable	Brennock Packaging: line speed 41 to 47 a minute, no new equipment
Rising costs	The monthly accounts, margin line in red, a pen circling it	Frustration, then clarity	The benchmark: 38 plants lost 4% to 9% of contribution margin on the floor

Moment	The cue, in the first three seconds	The feeling	The proof (signed by Credentials)
Cash in stock	A full warehouse; a forklift driver searching for the one part that isn't there	Absurdity	(No signed stock result yet: use the benchmark finding, or don't run this ad)
Scrap and quality	A skip of scrapped parts behind the plant	Waste you can see	Calloway Foods: scrap from 4.1% to 2.6%
New line ramp-up	A new line's counter showing 41 under a sign saying "Target: 60"	Disappointment, then hope	Brennock: changeovers from 52 minutes to 19
After an acquisition	Two plant signs, two sets of keys on one desk	Pressure	(Outbound's "first 100 days" page; no client named in an ad)
Automation not paying back	A new robot arm, idle, waiting for parts	Embarrassment	The benchmark: 6 of 38 plants had approved a new line before fixing the process

Two moments have no signed proof of their own yet, so their ads use the benchmark or wait. Credentials, not the ad, decides.

The supporting cells get versions, not new ads. **The owner's version** shows the same cue from the owner's desk, and says "your plant", not "your line". **The finance director's version** opens on the margin line and points to the benchmark's finance cut. Same assets, same line, same proof.

Distinctive assets. An ad that isn't recognised as Kestrel's writes a memory with no name on it. Kestrel's assets, tested with twenty buyers before use:

- **The name, always as "Kestrel Advisory, for manufacturers".** "Kestrel" alone is confused with a mining services firm.
- **One colour: a deep amber,** on the floor-marking tape that runs through every ad, as it runs through a plant.
- **The line,** word for word from the reason statement.

The book's rule: the brand appears early and briefly, then again, rather than at the end. One study pooled 50,000 viewers' responses to 25 television ads: the later the brand appeared, the less often viewers recognised it ([Gerber and colleagues](#)).

Step 3: gaps, biases and jobs

Moment	Stage	Bias	The job
COO · trigger, all seven	Known	Familiar feels safe	Five ads (flat output, rising costs, scrap, ramp-up, automation), each opening on its cue, the brand in the first two seconds, run all year to COOs and operations directors at the 1,400 firms
COO · trigger, cash in stock and after an acquisition	Known	Losses loom larger	Hold until there's signed proof; meanwhile, the benchmark ad covers them
Owner · trigger	Known	We follow others	The owner's versions, in association newsletters and to owners at the 1,400 firms; the 40 named accounts see them too, as part of the account plan
Finance director · trigger	Known	Losses loom larger	One ad, the margin line, pointing to the benchmark's finance cut; to finance directors at the 1,400 firms
COO and improvement lead, looking (problem searches)	Found	Good enough beats best	Small search ads on problem terms ("reduce changeover time"), sending people to the matching guide, not to a contact form
Anyone searching "Kestrel"	Found	—	Keep a small ad on the name, with "for manufacturers" in it and mining terms excluded. Nobody else bids, so test it
Past site visitors	Bought	A hard choice stalls	Switch retargeting off and hold it out. If the holdout shows it adds nothing, it stays off
The benchmark's release	Known	We follow others	Sponsor one association's annual members' report, with the benchmark's sector finding inside it, not a logo on the back

One pattern: **every ad starts from a moment and is judged by the memory it leaves, not by what happens the same week.** The collecting ads shrink until a test shows they earn their place.

Step 4: reconcile, the media plan

The same A\$48,000, spent differently.

Line	What runs	Weight	Writes or collects	A\$	Moments	Priority
LinkedIn, COOs	The five COO ads and the benchmark ad, rotating	Steady, all year	Writes	18,000	COO trigger	P1
Trade newsletters	The same ads, as small units in the four sectors' newsletters	Every issue	Writes	6,000	COO trigger	P1
Owners	The owner's versions: association newsletters and LinkedIn	Steady, all year	Writes	6,000	Owner trigger	P2
Finance directors	The margin-line ad, to finance directors at the 1,400 firms	Steady, all year	Writes	3,000	Finance director trigger	P2
Sponsorship	One association's annual members' report, carrying the benchmark's sector finding	Once a year	Writes	4,000	Owner and COO trigger	P2
Problem search	Ads on problem terms, to the matching guide	Always on, capped	Collects	4,000	COO and improvement lead, looking	P2
Name search	"Kestrel Advisory, for manufacturers", mining terms excluded	Always on, then tested	Collects	2,000	Anyone searching the name	P3
Retargeting	Off, held out	—	Collects	0	—	P3
Production	Five COO ads, owner and finance versions, newsletter units	Once, refreshed yearly	—	5,000	—	P1

Of the A\$43,000 in media, about 86% now writes and 14% collects. That's further from the 60:40 starting point than it looks. Kestrel's collecting ads were tested by the audit and mostly paid for people already coming. And two other builds already collect for free: the website and being named put Kestrel in front of people who search or ask an assistant. The holdouts will say whether to move the split back.

Rejected, with reasons:

The trade magazine's half-page ad. A logo at the bottom and a list of services: nothing a reader could file under a moment. The newsletters reach the same people, every issue, for the same money.

LinkedIn in bursts. Three bursts around the trade show let the memory fade between them. Steady weight all year, at a level Kestrel can keep.

Search ads on "operations consultant". Job-seekers and students. The problem terms reach people with the problem.

Broadcast and outdoor, as a nod. A billboard on the road into an industrial estate reaches everyone who works there, and very few of them choose consultants. Radio is the same, wider. For Kestrel, the cost per decision-maker reached is many times LinkedIn's.

Direct mail, as a nod. A postcard to all 1,400 firms would reach about 70 in the market and 1,330 who aren't, with nothing they'd keep. Outbound's letters, to firms with a signal, do this job better. For firms whose buyers are households or very small businesses, mail can still write a memory cheaply, if the piece is worth keeping.

Sponsorship without a finding. A logo on a conference lanyard, or on the back of an awards programme, writes the name without a moment. The members' report works because the sponsored page is the benchmark itself.

The anatomy library

Every ad has two halves. The creative half decides what gets written into memory; the media half decides how many memories, and how often. The creative half has five parts, in the order the buyer meets them. Short ads carry the first three.

Part	Stage	What it does
1. The moment, in the first three seconds	Known	Shows the buyer's situation, not the product: the email at 6am, the red margin line
2. The brand, early and distinctive	Known	The name and the amber tape in the first two seconds, then again
3. A feeling	Known, Preferred	One that belongs to the moment: relief, clarity, the end of embarrassment
4. One thing to remember	Preferred	The reason, word for word: "savings found on the floor in six weeks, at a fixed price"

Part	Stage	What it does
5. Proof, then where to find you	Believed, Found	One signed fact, then a name plain enough to search or say to an assistant
6. Broad reach	Known	Every COO, owner and finance director at the 1,400 firms, not just last month's visitors
7. Continuity	Known	Steady weight all year; memory fades between bursts

By format:

Ad format	Parts, in order	Length
LinkedIn video	The cue, sound off, with the amber tape in frame · "Kestrel Advisory, for manufacturers" by second two · One line of what usually causes it · The proof, with the client named · The line, and the name again	15 to 20 seconds; captions always
LinkedIn single image	The cue, photographed in a real plant · The name and the tape · The line · The proof in the post text	One image, under 40 words of post text
Newsletter unit	The buyer's words as the headline · The proof · The line · Where to find it	Under 50 words
Search ad, problem terms	The problem in the searcher's words · What the guide shows · "Kestrel Advisory, for manufacturers"	The platform's limits; sends to the guide, not a form
Sponsored report page	The benchmark's finding for that sector, with its number · What to check this month · Where it comes from · Kestrel's name and line, once, at the foot	One page

Three rules hold for every ad:

Same words as the site. The line and the proof in an ad match the website, the pitch and what assistants say. An ad that promises what the site doesn't repeat writes a contradiction into two memories.

One moment per ad. If it needs to serve two, make two.

Never a number Credentials hasn't signed. A finance director will ask where it came from.

Briefs

One brief per moment, written before anyone writes a script. This is the brief behind the finished example below.

AD BRIEF	COO · trigger · flat output
MOMENT	"We're running flat out and still shipping late."
WHO	COO or operations director, manufacturer of 80 to 400 staff, metal and engineering, food and drink, plastics, packaging
THE CUE	A phone on a desk at 6am. A customer's email: "Where's our order?" The COO picks it up and walks onto the floor
THE FEELING	Dread, then relief: the capacity is there, it's findable
ASSETS	"Kestrel Advisory, for manufacturers" by second two, and at the end. Amber floor tape in every shot
REMEMBER	"Savings found on the floor in six weeks, at a fixed price."
PROOF	Brennock Packaging: 41 to 47 a minute on the same line. Source: Credentials, client confirmed
FIND US	Search "Kestrel Advisory"
MEDIA	LinkedIn, COOs and operations directors at the 1,400 firms; trade newsletters in the four sectors. Steady, all year
NOT ALLOWED	"Guaranteed"; any saving not signed; a client's name without their permission; anything about a firm's late deliveries that implies we know about theirs
JUDGED BY	The regional holdout (see the plan); searches for the name; moment recall in new fit calls. Not by clicks

The finished example: the flat-output ad, written in full

The video, the image, the newsletter unit, the search ad and the owner's version, for one moment. Kestrel and Brennock are invented.

1. The LinkedIn video, 18 seconds, captions on, made to work with the sound off.

Time	Picture	Caption
0.0–2.5 s	A phone on a desk, 6am. The screen: an email from a customer, "Where's our order?" A strip of amber floor tape runs across the bottom of the frame.	Kestrel Advisory, for manufacturers
2.5–6.0 s	The COO picks up the phone and walks onto the floor. The line is running. Everyone is busy.	Running flat out. Still shipping late.
6.0–10.0 s	Close on a machine waiting: an operator standing by for a changeover part. A clock on the wall.	The capacity is usually there. It's lost in the waits.
10.0–14.0 s	The same line, weeks later, moving steadily. The amber tape is on the floor beside it.	Brennock Packaging: 41 to 47 a minute. Same line. No new equipment.
14.0–18.0 s	The COO, at the line, reading an email that's good news for once. The name and tape again.	Savings found on the floor in six weeks, at a fixed price. Kestrel Advisory, for manufacturers

Post text: "When a key customer starts asking where their order is, the answer is usually on your own floor. Brennock found 6 more units a minute on the same line, with no new equipment."

2. The single image, for the same moment.

Picture: a phone on a stainless steel bench, 6am, showing "Where's our order?". Amber floor tape across the foot of the frame, with "Kestrel Advisory, for manufacturers" printed on it.

Post text: "Running flat out and still shipping late? Brennock Packaging got from 41 to 47 a minute on the same line. Savings found on the floor in six weeks, at a fixed price."

3. The newsletter unit, in the four sectors' trade newsletters.

"We're running flat out and still shipping late."

Brennock Packaging found 6 more units a minute on the same line, with no new equipment.

Savings found on the floor in six weeks, at a fixed price. **Kestrel Advisory, for manufacturers.**

4. The search ad, on terms like "increase line output" and "late deliveries manufacturing".

Shipping Late on a Busy Line? | Where the Lost Capacity Hides

Kestrel Advisory, for manufacturers. A one-page check you can run on your own line this week. Savings found on the floor, at a fixed price.

It sends people to the flat-output guide, not to a contact form.

5. The owner's version, in association newsletters.

Your biggest customer asked where their order is. Again.

The capacity is usually on your own floor. Brennock Packaging found 6 more units a minute on the same line, with no new equipment. Savings found on the floor in six weeks, at a fixed price. **Kestrel Advisory, for manufacturers.**

What to notice. Every version opens on the moment, not on Kestrel, and every version names Kestrel in the first two seconds or the first line. The amber tape and the line are the same everywhere, so five formats write one memory. The proof is one signed number, with the client named, and the same number the website uses. The search ad sends people to something useful rather than asking for a meeting. And none of it will be judged by its clicks.

Templates

1. The ad audit scorecard. Use it on any ad, yours or a rival's. Watch once at normal speed and once with the sound off, and score each part 0, 1 or 2. A creative total under 6 means rebuild the opening first. The scale is a working rule, not tested against sales. The last two columns show Kestrel's old trade-magazine ad and the new flat-output video.

Part	Stage	Question	0	1	2	Old magazine ad	New video
The moment	Known	Is a buyer's situation shown in the first three seconds?	No moment	Shown, but late or vague	Clear, in the buyer's terms	0	2
The brand, early	Known	Could a stranger name the brand at three seconds, sound off?	Only at the end	Early, once	Early, brief, repeated, with an owned asset	0	2

Part	Stage	Question	0	1	2	Old magazine ad	New video
A feeling	Known, Preferred	Does it leave a feeling that belongs to the moment?	None	A generic mood	A clear feeling, tied to the moment	0	2
One thing to remember	Preferred	One line, word for word with the site?	None, or several	One, but different from the site	One, word for word	0	2
Proof, then where to find you	Believed, Found	A checkable fact, and a name to search?	Neither	One of the two	Both	1	2
Creative total						1/10	10/10
Broad reach	Known	Does it reach people who aren't yet buyers?	Past visitors only	Some new buyers	The whole market of possible buyers	1	2
Continuity	Known	On air all year, near your share of market?	One burst	Several bursts	Steady, with share of voice known	0	2
Media total						1/4	4/4

A perfect score on paper means the ad is built right, not that it works. The holdout says whether it works.

2. The ad brief.

AD BRIEF	[chooser · step · problem]
MOMENT	[one only, in the buyer's words]
WHO	[role, kind of business, segment]
THE CUE	[what they see, hear or feel in the first three seconds]
THE FEELING	[what the ad leaves behind]
ASSETS	[name by second __; colour; line; where each appears again]
REMEMBER	[the reason, word for word]
PROOF	[one signed fact, and its source; leave out under 10 seconds]
FIND US	[what to search or ask for]
MEDIA	[who must be reached, including non-buyers; months on air; weight per week; share of voice against share of market]
NOT ALLOWED	[claims, words and implications to avoid]
JUDGED BY	[the holdout; the leading measures. Not clicks]

3. The holdout test plan, filled in for Kestrel's first year.

QUESTION	What would Kestrel lose without the COO and owner ads?
REGIONS	12 areas in Kestrel's region, paired by number of manufacturers. One of each pair, chosen by coin toss, gets no LinkedIn or newsletter ads: 6 on, 6 off
LENGTH	Six months, then swap. A brand ad's effect builds slowly; eight weeks would miss it
MEASURES	Searches for "Kestrel Advisory" by area; site visits by area; fit calls by area, and what callers say started it; replies to Outbound letters by area
RETARGETING	Off everywhere for the first three months; on in the 6 "on" areas for the next three
NAME SEARCH	Before switching it off anywhere, check who else bids on the name. Today nobody does
READ AS	A sketch. With 12 areas and few fit calls, a small effect can hide in the noise. A result that says no is still a result
AGREED WITH	The managing partner and the finance lead, before the first ad

In the campaign plan

Advertising has two layers. **The always-on layer** is the moment ads, running at a steady weight all year: that's the memory, and it doesn't stop for campaigns. **The campaign layer** sits on top for a few weeks at a time, around the category's occasions, using the same assets and the same line.

Occasion	Campaign	What the ads do	What other builds bring
The benchmark's release (Q2)	The annual benchmark	The benchmark ad to COOs, owners and finance directors; the sponsored page in one association's members' report	Content: the benchmark; PR: trade coverage; Outbound: the sector findings to 160 firms
Budget season (Q3)	"Before you buy another line"	The ramp-up and automation ads at extra weight, to COOs and finance directors while capital plans are written	Website: the "before you buy" page; the pitch: the comparison with buying equipment

Occasion	Campaign	What the ads do	What other builds bring
Board season	Board season	The finance director's margin ad at extra weight	Website: the finance pages; Email: the newsletter
Owners' round tables (spring and autumn)	Owners' round tables	A small owner ad in the sector's association newsletter: "Dinner with owners of engineering firms, 12 March"	Events: the room; Partners: the association co-hosting
Any month	The account plan	The problem page shown to people at the 40 named accounts, while Outbound writes	Outbound: the letters; Events: a place held

The campaign layer never replaces the always-on layer. A campaign that pauses the moment ads to pay for itself is spending the memory to buy a burst. At Kestrel, the extra weight comes from inside the media plan's lines: about a fifth of each line is kept back for its occasions.

Running it

Who owns what. At Kestrel, the marketing lead owns the memory: which moments the ads write, the assets and the line. A freelance media buyer owns cost, reach and continuity. The head of delivery signs every proof point with Credentials. The managing partner and the finance lead agree the holdout before the first ad runs, and read it with the marketing lead.

Advertising is run by the quarter and judged by the year, because memory takes repetition.

When	What happens
Weekly	Delivery, reach, frequency and cost, for fixing delivery only. Search terms checked for job-seekers and mining
Monthly	Searches for "Kestrel Advisory", by holdout area

When	What happens
Quarterly	Moment recall: every new fit call asked "what was happening when you first thought of calling someone?" and "where had you heard of Kestrel?"
Six-monthly	The holdout read; the swap
Yearly	Assets tested with twenty buyers for fame and uniqueness; the split between writing and collecting reset from the tests; production refreshed

Keep the cues fixed. The moments, the amber tape and the line change rarely. The executions can change every year. If the team is bored of the ads, buyers are only beginning to notice them.

Tie every ad to the one truth. Before an ad runs, the marketing lead checks its line and proof against the website page it links to. If they differ, the ad waits.

For agencies. Sell a memory plan, not a media plan: the moments to own, the assets to keep, the reach to buy, and the test that will judge it. Report holdout results and name searches, not impressions. Agree the test with the client's finance team before the first flight.

Common mistakes

The mistake	The stage it breaks	The bias it ignores	At Kestrel	Instead
Trusting attributed sales	All, unseen	Good enough beats best	"31 conversions" from ads on its own name	A holdout. In 15 Facebook experiments, common methods mostly overstated what ads did (Gordon and colleagues)
The name at the end	Known	Familiar feels safe	Logo at the foot of the magazine ad	Name and amber tape in the first two seconds, then again

The mistake	The stage it breaks	The bias it ignores	At Kestrel	Instead
Opening on the firm, not the moment	Known	Feeling first, reasons after	"We're a leading operations consultancy"	Open on the cue the buyer would recognise
Bursts	Known	Losses loom larger	Three bursts around the trade show	Steady weight all year; campaigns sit on top
Three messages in one ad	Preferred	A hard choice stalls	A list of seven services	One thing to remember, word for word with the site
Retargeting by default	Bought	Good enough beats best	7 of 9 "conversions" were already in a deal	Hold it out; keep it only if it adds
Dropping name search by default, or keeping it by default	Found	We follow others	Paying for mining job-seekers	Check who bids; exclude the wrong searches; test. Where rivals bid and a brand stayed off, they took 18% to 42% of the clicks (Simonov, Nosko and Rao)
A number Credentials hasn't signed	Believed	Losses loom larger	—	No proof, no number: use the benchmark, or wait
Changing the assets every year	Known	Familiar feels safe	A new look for each trade show	Keep the tape and the line for years; change the executions

The bias to watch in the team: **the streetlight effect**. It's looking for something where the light is, not where it was lost. Clicks and attributed sales are lit up on every dashboard, so teams measure them. The memory an ad leaves sits in the dark, months away, and only a holdout or a question to a new client finds it.

The business check and the measurement ladder

Revenue: the ads cost A\$48,000 a year, and one review brings A\$45,000 to A\$90,000 in fees. To earn their place, they need to bring one or two reviews a year that wouldn't have come otherwise. The holdout is how Kestrel will know whether they did.

One truth: every ad says what the website, the pitch and the assistants say, in the same words, with the same proof.

The proof is real: every number is signed by Credentials and confirmed by the client. Two moments have no ad yet because they have no proof yet.

Capacity: about two hours a week of the marketing lead's time, plus the freelance buyer and a yearly production budget. Less than the trade-show stand used to take.

The measurement ladder. Every build can be measured on three rungs. Advertising needs all three, because its effects arrive late and its reports overstate.

Rung	What it tells you	For Kestrel's ads	How often	Use it to
Diagnostic	Whether the ads ran and were seen	Delivery, reach, frequency, cost, clicks	Weekly	Fix delivery. Never as a goal
Leading	Whether memory and choice are moving	Searches for "Kestrel Advisory", by area; moment recall in new fit calls; asset fame and uniqueness; what assistants say; the Score's Known	Monthly to yearly	Steer between tests
Causal	What the ads added that wouldn't have happened anyway	The regional holdout	Twice a year	Decide the budget

Share of search is the leading measure to watch most closely: Binet found it sometimes ran a year ahead of market share ([IPA](#)).

Larger firms add a fourth tool on the causal rung: a marketing mix model, which uses years of weekly spend and sales by channel to estimate what each added. Kestrel doesn't have the years of data or the volume to build one. Holdouts read the short clock; mix models read the long one.

Measure	What it tells you	Stage
Searches for "Kestrel Advisory", ads areas against holdout areas	Whether the ads are writing the name into memory	Known
New fit calls who name a trigger moment and say they'd seen Kestrel's ads	Whether the memory was filed under the right moment	Known
Fit calls by area, on against off	Whether the memory turns into conversations	Found, Bought
Asset tests: who links the amber tape and the line to Kestrel, and only Kestrel	Whether the assets are doing their job	Known
Problem-search visitors who read the guide to the end	Whether the collecting ads send the right people to the right page	Found, Answered

The Loop for this build

Every six months, when the holdout is read, the Loop (step 11) asks: **did the ads write Kestrel's name into the moments they were for, and did that turn into conversations?**

What the Loop finds	What it means	What changes
Name searches and fit calls higher in ad areas	The memory is being written, and collected on	Keep the weight; extend to the holdout areas after the swap
Name searches up, fit calls flat	Known is moving; something later is stuck	Check the website page and the reply time before touching the ads
Nothing moves	The ads aren't remembered, or the effect is too small to see in 12 areas	Test the assets with buyers; run longer before deciding; if still nothing, move the money (as the book's bookkeeper did)
New clients name a moment, but not the one the ads were for	The memory is filed elsewhere	Shift weight to the moment buyers actually name

What the Loop finds	What it means	What changes
Retargeting adds nothing in its three months	It was paying for people already coming	It stays off
A rival starts bidding on the name	The defensive case has arrived	Keep a name ad on, and watch who's bidding

What Kestrel's first year would look for. Searches for "Kestrel Advisory" rising faster in the ad areas than in the holdout areas. At least a third of new fit calls naming a trigger moment and saying they'd seen the ads. Name-search clicks from people looking for the mining firm close to none. And the Score's Known for COOs moving up at the next rescore: the measure this build is best placed to move.

Shapes by business type

Business type	Typical mode	The advertising that usually matters most
B2B professional services	Committee or owner-led	One ad per chooser's moment, steady all year, in the few places those choosers read; a small, tested search layer
B2B software	Committee	One ad per role in the buying group, shared assets; search on problem terms; regional holdouts, and a mix model once there are years of data
Considered consumer (finance, health, education)	Considered personal	Ads tied to the life events that start the search; the one thing to remember identical to the site; strict rules on claims
Ecommerce	Quick choice or habitual	Broad reach, all year, with fixed assets; retail media and search, tested with holdouts; the brand in the first seconds of every ad
Local services	Owner-led or considered personal	One ad, one moment, run long; local radio, outdoor and mail can work here; an on-off test by town or by month

Business type	Typical mode	The advertising that usually matters most
Marketplaces	Varies on each side	Ads to each side's moment; the supply side often needs collecting, the demand side writing

Feeling a bit overwhelmed?

You should be, a little. One small firm: the category's media mapped, share of voice against share of market, seven moments with a cue and a feeling each, three distinctive assets tested with buyers, five formats per ad, every number signed, a media plan split between writing and collecting, and a regional holdout agreed with finance before the first ad runs. Most firms bid on their own name, run a few bursts around a trade show, and count the clicks.

RocMode does part of the heavy lifting:

The moments to write into memory. Your choosers and their moments, mapped from what's known about your brand, so every ad starts from a situation a buyer would recognise.

Ad variants and creative briefs. RocMode drafts ad variants and creative briefs from your strategy, and exports them as files for Google and LinkedIn ads. You review them, and you decide what runs.

Campaigns timed to your category's key dates. RocMode's campaign concepts are built around your category's calendar, so the campaign layer sits on the occasions that matter.

[Join the RocMode beta](#)

Being Named: the Search and AI Answers Playbook

The idea

In search and AI answers, the job isn't to rank. It's to be named, accurately, with your reason, in the answer. Ask an assistant for "the best operations consultants for mid-sized food manufacturers" and it gives a short list with a sentence on each. Many buyers take that list as it comes. If Kestrel isn't on it, Kestrel isn't in the running. If it's on the list with the wrong facts, that's worse.

Three things make this build different from the others:


- **The buyer doesn't read your site; the assistant does.** It reads passages from many sources and repeats what most of them agree on. So half the work is on your own pages, and half is in the sources assistants trust.
- **Answers vary.** Ask the same question twice and you can get different firms. So measure shares across many runs, never a single answer.
- **Traffic misleads.** A buyer who sees your name in an answer, then searches for you or calls, may never arrive through a link. Measure being named and being right, not clicks.

The method is the same as every build: the shape (the sources the category is judged on), coverage (the questions from the lit moments), reconcile, brief, check and measure.

The footprint, and Kestrel in 60 seconds

Being named: lit in 13 cells, leading 20 of 149 moments

Kestrel's moment map. Rows: choosers. Columns: journey steps. Blank: that chooser isn't present at that step.



The being named footprint on Kestrel's moment map

This build leads every question Kestrel's buyers put to an assistant before the pitch: what's wrong, what the options are, what a good review should cover, who does this, and what it costs. It supports the COO's and finance director's early searches, the investor's search for firms, and the reputation check every buyer runs before signing: "Is Kestrel Advisory any good?"

Measuring how assistants see you is diagnosis: it belongs to step five, the Selectability Score. Earning a place in their answers is this build.

Kestrel in 60 seconds. Kestrel Advisory is the invented firm used in every playbook: a 60-person operations consultancy that helps mid-sized manufacturers find waste and lift output. Its main offer is a fixed-price, six-week operations review. Its strategy is in [Part one](#), and its full moment map, every chooser's question at every step, in [Where you stand, part one](#).

- **Who chooses:** usually a committee. The COO starts it, the finance director approves the spend, the plant manager lives with it and procurement runs the paperwork. The owner, the in-house improvement lead, an investor and AI assistants all have a say.
- **The reason to be chosen:** "We find the savings on the floor in six weeks, at a fixed price, and the review pays for itself or you don't proceed." (Invented. A real firm would need proof for every part of that sentence.)

- **The alternatives:** a large consulting firm, a specialist lean consultant, the in-house improvement team, buying new equipment, or doing nothing.

Its slice of the Score

This build owns the gate at the 61 moments asked of machines: **Known** (do assistants recognise Kestrel, and tie it to this moment?) and **Found** (is it named when buyers ask without naming it?). It shares Believed and Preferred with PR and reviews, wherever Kestrel is named. The numbers come from [the Selectability Score](#); the scan below is the same one, read for this build.

Kestrel today, at the 61 moments	Known	Found	Moment score
Average	12	0.5	0.2
Points the Score loses here	11	31	42 of the 82 lost

Half of everything Kestrel loses, it loses at these 61 moments, and nearly all of it at Found. No other build has as much of the Score riding on it. That's also why it's slow: Found moves only when the sources assistants read have changed, and they read on their own schedule. Expect the first movement two or three monthly scans after the work starts, not the week after.

Step 1: the shape, the sources the category is judged on

For a website, the shape is the page types buyers expect. For search and AI answers, it's the **sources** assistants and search engines draw on when someone asks about the category. Every category has its own mix, and it's knowable: run the questions, count what gets cited. For B2B professional services sold to manufacturers, the mix usually looks like this:

Source type	What assistants and search use it for	In Kestrel's category
Firms' own sites	What a firm does, who for, at what price, with what proof	Problem, sector, price and results pages; structured data; a plain-text copy of each page
Trade press and research	Framing a problem, and which firms are worth knowing	Manufacturing magazines, industry reports, benchmark studies

Source type	What assistants and search use it for	In Kestrel's category
Comparison and "best of" articles	Shortlists	Directory round-ups, trade press lists, consultants' own comparison pages
Directories and member lists	Who exists, what size, which sector	Industry association member lists, consultant directories
Reviews and ratings	Whether a firm is any good	Review platforms for B2B services, map listings
Community discussion	What practitioners really think	Forums, practitioner groups, discussion threads
Government and supplier panels	Who is approved, and what support exists	Productivity programmes, approved supplier lists
Video and podcasts	How the work looks in practice	Plant walk-throughs, manufacturing podcasts
Professional profiles	Who the people are, and their experience	Company and personal profiles

The **entity** runs through the shape. Before anything is published, write down the facts every source must state the same way:

- **The name, with what makes it unique:** "Kestrel Advisory, the operations consultancy for mid-sized manufacturers". Never plain "Kestrel", which is shared by other firms.
- **Who it's for:** manufacturers with 100 to 500 staff, in food and drink, packaging, metal and engineering, and plastics and rubber.
- **The offer:** a fixed-price, six-week operations review, on site.
- **The way out:** if the review doesn't find more than it costs, the client doesn't go ahead.
- **The proof pattern:** a named plant, a before-and-after number, the plant manager's words.

Assistants repeat what most sources agree on. A fact stated five different ways in five places gets dropped, or mixed up with another firm's.

Step 2: coverage, the question set

The coverage unit for this build is **each moment's questions, and who gets named today**. Start from the footprint: every lit cell, split by problem and sector where the question changes. That's 61 moments: 20 this build leads and 41 it supports.

Lead moments get three versions of the question, because buyers ask the same thing three ways:

1. **As typed into search:** "operations review cost"
2. **As asked of an assistant, in full:** "What does an operations review cost for a 200-person food plant, and what does it usually return?"
3. **With a follow-up or a constraint:** "Is fixed price better than day rates for that?"

Supporting moments get one version each: the lead build carries the rest.

Cell	This build	Moments	Versions each	Questions
Assistant · framing: "Why is my factory's output flat?"	Leads	7, one per problem	3	21
Assistant · options: "How do I improve productivity in a factory?"	Leads	7, one per problem	3	21
Assistant · scope: "What should I look for in an operations consultant?"	Leads	1	3	3
Assistant · comparing: "Best operations consultants for mid-sized manufacturers?"	Leads	4, one per sector	3	12
Assistant · building the case: "What does an operations review cost, and what does it return?"	Leads	1	3	3
COO · framing: "Is it people, process, equipment or planning?"	Supports content	7	1	7
COO · options: "Fix it ourselves, buy equipment or bring someone in?"	Supports the website	7	1	7
COO · scope: "What should the review cover, and what does a good one look like?"	Supports the website	7	1	7

Cell	This build	Moments	Versions each	Questions
COO · comparing: "Who does this for plants our size, in our sector?"	Supports the website	4	1	4
Finance director · framing: "Where's the money going?"	Supports content	7	1	7
Finance director · options: "What does each option cost and return?"	Supports the website	7	1	7
Investor · options: "Which firms deliver operations value quickly?"	Supports PR	1	1	1
Assistant · checking: "Is Kestrel Advisory any good?"	Supports PR	1	1	1

101 questions. Each one runs on five assistants and one search engine, three times each, because answers vary. That's **1,818 answers** in a single scan, every one recording whether Kestrel is named, who's named instead, what's said about each, and which sources were cited. And because answers change as sources change, the scan runs again every month.

Three rules for writing the questions:

- **Use the buyer's words, not yours.** They come from the cell: the top question, the pain, the objection. Nobody asks for "operational excellence partners".
- **Leave your name out,** except in reputation questions. A question that names Kestrel tells you nothing about whether a buyer would find it.
- **Keep the set stable.** Change the questions and you can't compare one month with the next. Add new ones when the grid changes; retire old ones with a note.


Step 3: who gets named today, then gaps, biases and jobs

Run the question set before changing anything. The first scan gives two things: the sources the category is actually judged on, and where Kestrel stands in each moment. All the results below are invented for the example.

The sources. Count every citation across the 1,818 answers and group them by source type. A source type cited in at least one in ten citations is **canon** for this category: it's where assistants look, so it's where Kestrel has to be.

Four source types carry three in four citations. Kestrel appears in one old listing

Share of citations by source type, first scan of 101 questions. Canon (at least 10% of citations): 75% of all citations.



Invented for the Kestrel example · first scan, 101 questions, 1,818 answers

Where Kestrel stands. A question counts as "named" if Kestrel appears in at least one engine's answer, in any of the three runs. **Kestrel is named in 4 of the 101 questions, and in one of those the facts are wrong.**

Questions	Count	Kestrel named	Who's named instead, and what's cited	The stage that fails
Assistant · framing, seven problems	21	0	No firms; trade press, a large firm's articles, a government productivity guide	Found
Assistant · options, seven problems	21	0	A large firm and a lean specialist, from their own articles	Found
Assistant · scope: what to look for	3	0	A large firm's hiring checklist, which favours large firms	Answered
Assistant · comparing, four sectors	12	1	Two large firms, a lean specialist and an engineering firm; Kestrel once, in packaging, from an old directory listing	Found, then Preferred

Questions	Count	Kestrel named	Who's named instead, and what's cited	The stage that fails
Assistant · cost and return	3	0	"It varies widely"; one large firm's return claim, unsourced	Answered
COO and finance director · framing, options and scope	35	2	Two old Kestrel blog posts on page one for two problems; otherwise press and large firms	Found
COO · comparing, four sectors	4	0	The same firms as the assistant's shortlist	Found, then Preferred
Investor · options	1	0	Large firms with private equity practices	Found
"Is Kestrel Advisory any good?"	1	1	Named, but mixed up with a mining-services firm of a similar name; no reviews cited	Believed

Each gap becomes a job: a stage, a moment and a bias.

Moment	Stage	Bias	The job
Assistant · framing: "Why is my output flat?"	Found	Good enough beats best	Publish answer-first passages on each problem page, with sourced figures, so they can be quoted
Assistant · framing	Known	Familiar feels safe	Get the benchmark report covered in the trade press assistants already cite for these problems
Assistant · options	Answered	A hard choice stalls	A fair comparison of the options, including buying equipment and doing it in-house, in passages an assistant can lift

Moment	Stage	Bias	The job
Assistant · comparing, by sector	Found	Good enough beats best	One sentence on every sector page that states size, sector and offer; the same sentence in every directory and member list
Assistant · comparing	Preferred	Feeling first, reasons after	Give assistants the reason in words they can repeat: fixed price, six weeks, pays for itself or you don't proceed
Assistant · comparing	Found	We follow others	Be included in the comparison articles that make up a fifth of all citations, on merit, with a case per sector
Assistant · cost and return	Answered	Losses loom larger	Publish the price range, the typical return and the way out, side by side, with each figure sourced
"Is Kestrel Advisory any good?"	Believed	Familiar feels safe	Fix the mix-up: state the full name and what makes it unique everywhere, correct the confused listings, add structured data for the organisation
"Is Kestrel Advisory any good?"	Believed	We follow others	Reviews from client finance directors and plant managers, on the platforms assistants cite
Investor · options	Found	We follow others	Appear in the investor-facing press and portfolio operations articles, with portfolio cases
COO and finance director · search	Found	Good enough beats best	Move what the two ranking blog posts earned onto the new problem pages, so the rankings carry over
All moments	Found	Familiar feels safe	Keep sources current: date every figure, refresh the benchmark each year, update listings when facts change

Every job names a moment, a stage and a bias. That's what stops this build turning into "more content".

Step 4: reconcile, the source plan and the answer plan

Put the jobs into the shape. Two plans come out of it: which sources to work on, and what answer each group of questions should give.

The source plan

The same three signals as the website set each source's priority. **Canon:** the source type carries at least one in ten citations. **Coverage:** it carries jobs from Step 3. **Strategy:** the positioning needs it.

Source	Signals	Priority	The jobs it carries	What Kestrel does
Kestrel's own site	Canon, coverage, strategy	P1	Answer-first passages, price and return, the sector sentence, the options comparison	Built in the Website Playbook: problem, sector, price, results and comparison pages; organisation structured data; a plain-text copy of each page
Trade press and research	Canon, coverage, strategy	P1	Framing and options; being known for the problems	The annual benchmark report, offered to the titles assistants already cite; articles by Kestrel's people on each problem
Reviews and ratings	Coverage, strategy	P1 (strategy)	"Is Kestrel any good?"; belief at comparing	Reviews from client finance directors and plant managers, asked for after the results are in

Source	Signals	Priority	The jobs it carries	What Kestrel does
Comparison and best-of articles	Canon, coverage	P2	Being on the shortlist, by sector	A one-page fact sheet and a case per sector, sent to the editors of the lists that are cited; corrections where facts are wrong
Directories and member lists	Canon, coverage	P2	The sector sentence; fixing the mix-up	The same sentence in every listing; the old listing corrected; membership of the four sector associations
Community discussion	Coverage	P3	What practitioners think	Kestrel's people answer practitioner questions under their own names, saying where they work
Government and supplier panels	Coverage	P3	Procurement and investor checks	Approved-supplier listings where clients ask for them
Video and podcasts	Coverage	P3	How the work looks	One plant walk-through per sector; guest spots on the podcasts that are cited
Professional profiles	Coverage	P3	Who will be on site	Team profiles that match "The team on your floor"

Rejected, with reasons:

- **An encyclopedia entry.** A 60-person firm is unlikely to meet the notability rules, and writing your own entry breaks them. It would be removed, and the attempt would be visible.
- **Paid places in best-of lists.** A paid placement is an ad. It belongs in the advertising playbook, labelled as one.

- **Mass-produced articles.** No moment needs them. Thin pages dilute the passages that do the work, and contradict the proof pattern.
- **Questions outside the grid** (lean certification, factory software reviews). No chooser asks them on the way to buying a review.

The answer plan

For each group of questions, write the answer Kestrel wants an assistant to give, the passage that holds it, and the sources that must agree with it. This is this build's equivalent of a sitemap.

Questions	The answer we want	The passage that holds it	Sources that must agree	Priority
Comparing, by sector (16)	Kestrel Advisory, the operations consultancy for mid-sized manufacturers, runs fixed-price, six-week reviews in plants like theirs, with a named case	The first passage of each sector page; one case per sector	Directories and member lists; comparison articles; reviews	P1
Cost and return (3)	A fixed price range; the review typically finds more than it costs; if it doesn't, the client doesn't go ahead	Price and terms; results across every review	Benchmark coverage in the trade press; reviews	P1
Reputation (1)	The right firm, the right facts, and what clients say	About; results; reviews	Reviews; directories; profiles	P1, because wrong facts get repeated
Framing, by problem (35)	Kestrel's explanation of the cause, with a benchmark figure	Problem pages; the guide	Trade press articles drawing on the benchmark	P1 for flat output and rising costs, P2 for the other five

Questions	The answer we want	The passage that holds it	Sources that must agree	Priority
Options, by problem (35)	A fair comparison, including buying equipment and doing it in-house, and when each is right	Kestrel compared; Kestrel vs new equipment	Trade press; comparison articles	P1 for new equipment, P2 for the rest
Investor (1)	Kestrel works with portfolio companies, starts within the first 100 days, and reports in board terms	For private equity investors	Investor-facing press	P2
Scope, what to look for (10)	What a good operations review covers, how to brief one and how to compare quotes, fairly, with fixed price and a way out among the criteria	The review guide; Questions to ask; For procurement	Trade press; directories and member lists	P2

The coverage test: every one of the 61 moments has a target answer, a passage and at least one source. Every canon source type has a plan. And one check the website doesn't need: **every source states the entity facts the same way.** If one directory says "100 to 500 staff" and another says "mid-market", fix one before writing anything new.

The anatomy library

For a website, the anatomy is the page. For this build it's the **passage**: the few sentences an assistant lifts to answer one question. Each kind of question wants its own passage shape. These are the eight Kestrel's answer plan uses.

Passage type	Answers	Its shape	Length	Where it lives
The entity line	"Who is Kestrel Advisory?"	Full name · what it is · who it's for, by size and sector · the offer	One sentence, identical everywhere	Every page footer, every listing, every profile, the fact sheet
The cause	"Why is my output flat?"	The answer first · the three usual causes, ranked · one sourced figure · what to check first	80 to 120 words	Problem pages; articles in the trade press
The options	"How do I improve productivity?"	Each option in one line · when each is right, including the ones that aren't you · the order to try them	100 to 150 words	Comparison pages; the guide
The shortlist entry	"Best consultants for food manufacturers?"	Entity line · what makes it different, in the reason's words · one case from this sector	40 to 60 words	Sector pages; comparison articles; directories
The price and return	"What does a review cost and return?"	The range · what sets it · the typical return, sourced · the way out	50 to 80 words	Price and terms; the results page
The criteria	"What should I look for?"	Five criteria, fairly put · why each matters · questions to ask any firm	100 to 150 words	"Questions to ask"; the procurement page

Passage type	Answers	Its shape	Length	Where it lives
The proof	"Does it work?"	Named plant, size and sector · before and after · who confirmed it · the plant manager's words	40 to 70 words	Case pages; reviews; the benchmark
The reputation	"Is Kestrel Advisory any good?"	Entity line · years and reviews completed · what clients say, with the platform · how to check	60 to 90 words	About; review platforms; profiles

Three rules hold for every passage:

- **It stands alone.** No "as mentioned above", no pronoun pointing at the previous paragraph. An assistant lifts it without its neighbours.
- **Its facts match the fact sheet word for word.** "100 to 500 staff", never "mid-market" in one place and "SMEs" in another.
- **Every figure carries its source in the same sentence.** A figure whose source is two paragraphs away gets quoted without it, or not at all.

Step 5: the answer brief, then a brief for every passage

Assistants don't quote pages. They quote passages: a few sentences that answer one question on their own. So this build briefs at two levels. The **answer brief** holds what no single passage can: the questions, what's said today, the answer wanted, the facts that must match, and the sources to earn. Then each **passage** gets its own brief, like a section on a web page.

Here's the answer brief for one group of questions:

ANSWER BRIEF Comparing, food and drink

QUESTIONS (4) "Best operations consultants for mid-sized food manufacturers?" + two other versions; COO: "Who does this for food plants our size?"

MOMENT Assistant and COO · comparing · food and drink · stage: Found, then Preferred · bias: good enough beats best

TODAY Not named. Named: two large firms, a lean specialist. Cited: two comparison articles, a trade magazine, a large firm's site

THE ANSWER WE WANT Kestrel Advisory, the operations consultancy for mid-sized manufacturers, runs fixed-price, six-week operations reviews in food and drink plants, working on the floor on changeovers and hygiene downtime. A named plant, its before-and-after figure and its plant manager's words

FACTS THAT MUST MATCH Name and descriptor; 100 to 500 staff; the four sectors; fixed price; six weeks; on site; the way out

PASSAGES Food and drink page: "Who we work with" and "What a review finds in a food plant"; the food case page; Price and terms: "What it costs"

SOURCES TO EARN The two cited comparison articles (fact sheet + food case); the food manufacturers' association member list; two reviews from food clients

GUARD No claim an assistant can't check ("leading", "best"); every figure sourced and dated; the full name every time

Then each passage gets its own brief. Two of the four:

PASSAGE Food and drink page · "Who we work with"

ANSWERS "Who does operations reviews for mid-sized food manufacturers?"

FIRST SENTENCE The answer, with the entity facts: who, for whom, what, at what price model

THEN The two problems food plants bring most, in the plant manager's words; one named plant

LENGTH Under 80 words, so it stands alone when lifted

ENTITIES Kestrel Advisory; operations review; food and drink manufacturing; the named plant

GUARD Reads correctly with nothing around it; matches the directory sentence word for word

PASSAGE Price and terms · "What it costs"

ANSWERS "What does an operations review cost, and what does it return?"

FIRST SENTENCE The fixed price range, for which plant sizes

THEN What's included; the typical return across reviews, with its source and date; the way out, in the contract's own words

LENGTH Under 80 words

ENTITIES Kestrel Advisory; operations review; the price range; the results page

GUARD The range matches proposals and every listing; the return figure comes from the results page, never rounded up

The flow check is different here. A web page is read top to bottom; a passage is lifted on its own. So read each passage alone, as an assistant would quote it. If it needs the paragraph before it to make sense, rewrite it.

The finished example: one answer, written everywhere it has to be

One answer from the plan, done in full: the food and drink shortlist. The question is "best operations consultants for mid-sized food manufacturers". Today, assistants name two large firms, a lean specialist and an engineering firm, and never Kestrel. Here's every piece it takes to change that. Kestrel, its clients and its people are invented.

1. The entity line, fixed once and used everywhere:

Kestrel Advisory is the operations consultancy for mid-sized manufacturers: plants with 100 to 500 staff in food and drink, packaging, metal and engineering, and plastics and rubber. Its fixed-price, six-week operations review finds savings on the floor, and if it doesn't find more than it costs, the client doesn't go ahead.

2. The shortlist entry, the first passage on the food and drink page:

Operations reviews for food and drink manufacturers with 100 to 500 staff. Kestrel Advisory reviews food plants on the floor, over six weeks, at a fixed price. In 2025 it reviewed Calloway Foods, a 310-person sauce and dressings maker: scrap on the main filling line fell from 4.1% to 2.6% within three months, confirmed by Calloway's finance director. If a review doesn't find more than it costs, the client doesn't go ahead.

3. The directory listing, corrected in the four association member lists and the two consultant directories assistants cite:

Kestrel Advisory · Operations consultancy · Manufacturers with 100 to 500 staff · Food and drink; packaging; metal and engineering; plastics and rubber · Fixed-price, six-week operations reviews · Not affiliated with Kestrel Mining Services.

4. The correction request, sent to the directory whose old entry assistants mix up with a mining-services firm:

Our listing combines two firms. Kestrel Advisory is an operations consultancy for manufacturers; Kestrel Mining Services is unrelated. Please replace our entry with the attached text, and remove the mining sector and the old address. Our website states the same facts at kestreladvisory.example/about.

5. The fact sheet for editors, one page, sent with a sector case to the three comparison lists that make up a fifth of citations: the entity line; the offer, price range and way out; the four sectors and size band; one case per sector, each confirmed by the client; the two people editors can quote, with their plant experience; and the benchmark's headline figure, with its method.

6. The structured data on the about page, stating only what the page says:

```
{"@context": "https://schema.org", "@type": "ProfessionalService",
  "name": "Kestrel Advisory",
  "description": "Operations consultancy for manufacturers with 100 to 500 staff",
  "areaServed": "Australia",
  "knowsAbout": ["Food and drink manufacturing", "Packaging manufacturing",
    "Metal and engineering", "Plastics and rubber"],
  "sameAs": ["[association member page]", "[consultant directory page]",
    "[company profile]"]}
```

7. The answer Kestrel wants back, the test the monthly scan applies:

"...Kestrel Advisory works with food manufacturers of 100 to 500 staff, on a fixed-price, six-week review; it reports a 37% scrap reduction at Calloway Foods..."

If the assistant names Kestrel, gives the size band, the fixed price and a food case, the answer has landed. If it names Kestrel with mining in the same sentence, step 4 hasn't. Six pieces of work, in six places, for one question in one sector. There are three more sectors, and six more groups of questions in the answer plan.

Templates

The entity fact sheet (one per firm; every source copies from it)

FULL NAME	[the name, plus what makes it unique]
NOT TO BE CONFUSED WITH	[firms with similar names, and how they differ]
ENTITY LINE	[one sentence: what it is, who for, the offer]
WHO FOR	[size band, sectors, regions, word for word]
THE OFFER	[what, how long, how priced]
THE WAY OUT	[the guarantee or exit, as in the contract]
PRICE	[the range, and what sets it]
PROOF PATTERN	[what every case states]
PEOPLE TO QUOTE	[names, roles, experience]
LINKS	[site, profiles, listings: every place the facts appear]
LAST CHECKED	[date, and by whom]

The question set (one row per question; never edited, only retired)

ID	Moment (chooser · step · need)	Lead or support	Version (search / assistant / follow-up)	The question, in the buyer's words	Added	Retired, and why

The answer brief (one per group of questions)

QUESTIONS	[the IDs in this group]
TODAY	[who's named, what's said, what's cited]
THE ANSWER WE WANT	[in one or two sentences]
FACTS THAT MUST MATCH	[from the fact sheet]
PASSAGES	[type from the anatomy library, and where each lives]
SOURCES TO EARN	[from the source plan, with what each needs]
PRIORITY	[P1-P3, and why]
DONE WHEN	[what the scan must show, for how many months running]

The passage brief

PASSAGE TYPE	[from the anatomy library]
THE QUESTION	[answered in the first sentence]
JOB	[stage · chooser and moment · bias]
FACTS	[from the fact sheet, word for word]
FIGURE + SOURCE	[in the same sentence]
LENGTH	[from the anatomy library]
GUARD	[passes only if it reads correctly on its own]

The scan log (one row per question per month)

Month	Question ID	Engine	Runs naming Kestrel (of 3)	Facts right?	Reason given?	Named instead	Sources cited

Step 6: the business check

- ☐ **Revenue:** the sector and cost questions sit closest to new work. If the plan spends most effort on framing questions, it's chasing attention, not choices.
- ☐ **One truth:** the entity facts are identical on the site, in every listing, in proposals and in the fact sheet sent to editors.
- ☐ **The proof is real:** the benchmark report needs enough reviews behind it, and clients' permission to use their data. Without both, it isn't published, and the trade press plan waits.
- ☐ **Reviews are earned:** asked for after the results are in, from clients who agree, never paid for or written for them.
- ☐ **Capacity:** someone owns the monthly scan, the listings and the press relationships. If nobody does, start with the four P1 jobs: the sector sentence everywhere, the price passage, the mix-up fixed, and reviews.

Step 7: the measures

Measure shares across many runs, never single answers. Report by group of questions, by engine and by month.

Measure	What it tells you	Stage
Share of questions where Kestrel is named	Whether Kestrel is in the running at all	Found
Share of answers naming Kestrel that give its reason	Whether the reason is getting through	Preferred
Share of answers with Kestrel's facts right	Whether the entity is clear, or still mixed up	Believed
Position in shortlists	First named, or an afterthought	Preferred
Share of citations from sources Kestrel has worked on	Whether the source plan is working	Found
Buyers who say an assistant or search named Kestrel, from first-call notes	Whether being named turns into conversations	Bought

From the first scan, Kestrel is named in 4 of 101 questions, with wrong facts in one. A plan that works shows up in that order: facts corrected first, then named in the sector questions, then the reason given back. Clicks barely appear in this list. Most of the value happens before anyone clicks.

In the campaign plan

This build rarely leads a campaign, but every campaign leaves something behind for it: a page, an article, a review, a mention in a source assistants read. The rule in the campaign plan: **every campaign names the questions it should move, and the scan checks them a month or two later.**

Kestrel campaign	What it leaves for this build	Questions it should move
The annual benchmark	Coverage in the trade press assistants cite; a quotable summary on the site	Framing, all seven problems; cost and return
"Before you buy another line"	The comparison page, and the article it's based on	Options, by problem; the new-equipment comparison
Owners' round table	A write-up in an association newsletter; speaker profiles	The investor's and owner's options; reputation
A sector case each quarter	The sector page's case; the fact sheet updated for editors	Comparing, in that sector
Review requests after results	Reviews on the platforms assistants cite	"Is Kestrel Advisory any good?"

A campaign that leaves nothing for an assistant to read has reached only the people who saw it that week.

Running it

Who owns it. At Kestrel, the marketing lead owns the question set, the fact sheet and the monthly scan. PR owns the press, list and review relationships. The managing partner signs the fact sheet, because every source copies it.

When	What happens	Takes
Monthly	The scan: 101 questions, six engines, three runs. The scan log updated; any wrong fact flagged the same day	Half a day to read, with the scan automated
Monthly	Corrections: every wrong fact traced to its source, and the source asked to fix it	One to two hours
Quarterly	The source plan reviewed against what was cited; one sector case added; the fact sheet re-checked	A day
Yearly	The benchmark published; the question set reviewed against the map	Several weeks, mostly the benchmark
When a fact changes	Price, sectors, size band or the way out: the fact sheet first, then every place on its links list, the same week	Depends on the list; keep it short

Correcting an assistant. You can't edit an answer. You can fix what it reads. Find the sources it cites for the wrong fact, correct those you control, ask the owners of the rest, and make sure your own pages state the right fact in the first sentence. Then wait two scans before judging.

What not to do. No paid places presented as editorial, no reviews written for clients, no thin pages written to catch questions. Assistants and search engines are built to discount them, and the first two would break the claims rule: nothing said that Credentials can't sign.

Common mistakes

The mistake	The stage it breaks	The bias it ignores	What it would do at Kestrel	Instead
Chasing rankings, not being named	Found	Good enough beats best	Celebrate two blog posts on page one, while assistants name four other firms	Measure the share of answers that name you

The mistake	The stage it breaks	The bias it ignores	What it would do at Kestrel	Instead
Checking a handful of questions in one assistant	Found	We follow others	One good answer in packaging, read as a trend	The full question set, six engines, three runs, every month
Counting questions that name you	Found	Familiar feels safe	"Is Kestrel any good?" makes the firm look visible	Branded questions only for reputation, never for Found
Facts stated five ways in five places	Believed	Familiar feels safe	"Mid-market", "SMEs" and "100 to 500 staff": the assistant drops all three, or mixes Kestrel with a mining firm	One fact sheet; every source copies it
Being named without the reason	Preferred	Feeling first, reasons after	Kestrel on the shortlist as "a smaller consultancy"	Passages that carry the reason in words an assistant can repeat
Writing only about yourself	Answered	A hard choice stalls	No answer to "should I just buy a new line?"	Fair options passages, including when the alternative is right
Buying or faking the sources	Believed	We follow others	A paid list place or written-for-them reviews, found out by an editor	Earned places and real reviews, asked for after results
Judging it by clicks	All, unseen	Good enough beats best	"AI sends us 30 visits a month", while buyers search Kestrel by name after reading an answer	Being named, being right, and branded search and direct visits read together

The mistake	The stage it breaks	The bias it ignores	What it would do at Kestrel	Instead
Expecting it next week	Found	Losses loom larger	The work is dropped after one flat scan	Two or three scans before judging any change

The bias to watch in the team: **the illusion of control**. You don't write the answer. You write what the answer is built from, and you earn a place in the sources that build it.

The Loop for this build

Each quarter the Loop (step 11) reads three months of scans against the answer plan, one group of questions at a time, and asks: **is the answer we wanted coming back, from the sources we worked on?**

What the Loop finds	What it means	What changes
Named more, facts right, reason missing	The entity is clear; the passages don't carry the reason	Rewrite the shortlist entries and the price passage
Named more, facts wrong	A source is still contradicting the fact sheet	Trace the citations; correct the source, not just the site
Not named, sources unchanged	The source plan hasn't been worked, or the sources didn't take it	Check the P1 sources one by one; change approach, not effort
Not named, sources changed	Assistants haven't re-read yet, or the passages don't answer the question	Wait one more scan; then test the passages against the question
Named, but no change in branded search, direct visits or first calls	Being named isn't turning into conversations	Check Preferred and the reason; check the site's Bought

What Kestrel's first year would look for. Corrections first: the mining mix-up gone within two or three scans of the listings being fixed. Then the sector questions, as the sector pages and directories line up: named in some food and packaging answers by mid-year. The framing and options questions last, because they depend on the benchmark being

covered, and that comes once a year. If Kestrel is named in a quarter of the comparing questions by the year's end, with its reason, the build is working, even though the Score will still say "rarely selected" at those moments. Found moves first; the Score follows.

Shapes by business type

The source mix changes with the business. Run the questions and count the citations to find yours; these are the usual starting points.

Business type	Sources that usually carry the most weight
B2B professional services	Trade press, comparison articles, directories and member lists, firms' own sites
B2B software	Software review platforms, comparison and alternatives pages, documentation, developer communities
Considered consumer (finance, health, education)	Comparison sites, review sites, regulators' and government guidance, community discussion
Ecommerce	Product feeds and marketplaces, reviews, retailer listings, video reviews
Local services	Map listings and their reviews, local directories, community groups
Marketplaces	Reviews on both sides, community discussion, press about trust and safety

Feeling a bit overwhelmed?

You should be, a little. That was one firm: 61 moments, 101 questions, six engines, three runs each, 1,818 answers a month, nine source types and one set of facts that has to match everywhere. Most teams try a handful of questions in one assistant and call it done. The answers change every month, so a one-off check goes out of date fast.

RocMode does the heavy part for you:

- **Your questions, built from your buyers.** RocMode maps your choosers and their moments, then writes the questions they ask, in their words.

- **Run across leading AI assistants.** Each question runs on five assistants, three times each. It records whether you're named, who's named instead, and which sources the assistants cite.
- **Scored, with the biggest lever named.** Your selectability, moment by moment and stage by stage, with a range on the score, questions that name you kept out of Found, and the fix worth the most points.

[Join the RocMode beta](#)

The PR and Reviews Playbook

What this build is for

PR and reviews put someone else's voice behind your claims, at the moment a buyer stops taking your word for it.

Every other build is Kestrel talking about Kestrel. At some point in every purchase, that stops being enough. The COO wants to talk to a plant like theirs. The investor wants another portfolio company to vouch. The assistant, asked "is Kestrel Advisory any good?", goes looking for what others have written. This build is the Earned category's second build, beside Being named: the outside voices a buyer trusts more than yours, earned rather than bought.

It has three parts, and most firms run only one of them:

- **Press and research:** the trade titles and reports that frame the category's problems, and name the firms worth knowing.
- **References:** clients who'll take a call from a buyer like them, on request, quickly.
- **Reviews:** what clients write on the platforms buyers and assistants check.

The Earned category's three nods live here, each as a short section: **analyst relations** (rarely relevant to a firm of Kestrel's size, and said so), **awards** (worth entering only when the judges and the buyers overlap) and **word of mouth** (which this build can't buy, but can make easy).

The footprint, and what it reads from the chain

PR and reviews lead **4 cells and 7 moments**, all at the checking step, where buyers go looking for someone other than Kestrel:

Moment	Who	Their question
Checking, by sector (4)	COO	"Can I talk to a plant like ours that's done this?"
Checking	Improvement lead	"Have they worked alongside an in-house team before?"
Checking	Investor	"Will other portfolio companies vouch for them?"
Checking	AI assistant	"Is Kestrel Advisory any good? Reviews and results?"

But its reach is far wider than its lead. It supports **15 more cells**, the widest support of any build except the website: the triggers for the COO, owner, finance director and investor (where the benchmark's press coverage frames the problem), the owner's and investor's options, comparing for the COO, owner, investor and assistant, the finance director's numbers at the board, the owner's and finance director's checking, and the owner's and investor's follow-on. In almost every one, the job is the same: **Believed**, from a voice that isn't Kestrel's.

What it reads, and never rewrites:

From	It takes
Credentials	Every claim, and which ones outside voices could confirm
Content	The benchmark report, the only piece built to be quoted by others
Being named	The sources assistants cite, which decide where reviews and coverage matter
Email and Automation	The 90-day check, where a confirmed result becomes a reference or a review
The pitch	Which references each deal needs, by sector and role

Its slice of the Score, and its own audit

This build's part of [the Selectability Score](#) is **Believed**: the share of what's said about Kestrel that's accurate, and backed by someone other than Kestrel. It leads seven moments and touches Believed in 15 more cells.

Kestrel today, the 7 moments PR and reviews lead	Answered	Believed	Preferred	Bought	Moment score
COO · checking, by sector (inside deals)	60	65	55	50	58
Improvement lead · checking	30	30	35	40	33
Investor · checking	20	30	30	35	28
Assistant · "Is Kestrel Advisory any good?"	10	20	15	15	2

The COO's checking scores well because Kestrel's references, when they're found, are good. Everything else is thin. The assistant moment is the clearest: asked about Kestrel by name, assistants mix it up with a mining-services firm, find no reviews, and cite an old directory listing.

The audit.

Outside voice	What Kestrel has	What it misses
Press and research	No coverage in two years; one award (regional, 2023)	The trade titles that frame flat output and rising costs, which assistants cite
References	Four clients willing to take a call; none in food and drink or plastics; found by asking around, in two to three weeks	A reference per sector, ready within a week; an in-house improvement lead; an investor
Reviews	None on the review platforms assistants cite	Reviews from finance directors and plant managers, after confirmed results
Directories and lists	One old listing, mixed up with a mining firm	Correct entries in the four sector associations
Word of mouth	22 of 31 fit calls came by referral	Nothing that makes referring easy, or thanks anyone for it

That last row is the one to hold on to. Word of mouth is already Kestrel's biggest channel. It works despite the firm, not because of it.

Step 1: the shape, press, references, reviews

The shape is **the outside voices the category is judged by**: who buyers ask, and who assistants read. It's knowable. Being named's first scan already counted the sources assistants cite for Kestrel's questions; add the people buyers ask in person.

Voice	Who it speaks to	In Kestrel's category	Canon?
Trade press	COOs, finance directors, assistants	Three manufacturing titles cited in a tenth of answers	Yes
Research reports	Finance directors, assistants, the press	Industry and productivity reports; competitors' studies	Yes
References	Every chooser at checking	Expected by every buyer; offered by every competitor	Yes
Review platforms	Assistants, owners, finance directors	Two B2B review platforms cited in the scan	Yes
Association member lists	Assistants, procurement	The four sector associations	Yes
Analyst firms (nod)	Large buyers' procurement teams	Rarely, for firms of Kestrel's size	No
Awards (nod)	Owners	Two industry awards whose judges are plant and operations leaders	Optional
Peers and advisers (word of mouth, nod)	Owners, investors	Industry groups, accountants, operating partners	The biggest, and not earned by asking

Analyst relations, as a nod. For a 60-person firm selling to mid-sized plants, analysts don't shape the shortlist. Revisit only if large manufacturers' procurement teams start asking for analyst coverage.

Awards, as a nod. Enter only awards whose judges are the people Kestrel sells to, and only with a confirmed result. A business award judged by bankers moves no moment; an operations award judged by plant leaders can move an owner's checking.

Step 2: coverage, one story or review per judgement that lacks outside support

The coverage unit is **a judgement a buyer makes that needs someone else's voice**: not "get more press", but "the finance director needs someone other than Kestrel to say the savings were real". For each of the 22 cells this build leads or supports, write down the judgement, whose voice would settle it, and where the buyer or the assistant would find it.

Judgement	Moments	The voice that settles it	Where it's found
"Is this problem worth fixing, and is the capacity really in the line?"	Triggers: COO, owner, finance director, investor (22)	Trade press and research, quoting the benchmark	Trade titles; assistants' answers
"Is this firm worth a look?"	Options and comparing: owner, investor, COO, assistant (12)	Comparison articles; association lists; peers	Lists; assistants; industry groups
"Have they done it in a plant like mine?"	COO checking by sector (4)	A client in the same sector	A reference call
"Will they work with my team, not replace it?"	Improvement lead checking (1)	Another improvement lead	A reference call
"Will other portfolio companies vouch?"	Investor checking and follow-on (2)	An operating partner or portfolio COO	A reference call
"Are they any good?"	Assistant checking (1); owner checking (1)	Reviews; correct facts; peers	Review platforms; directories; people they trust
"Were the savings real?"	Finance director checking and building the case (8)	Another finance director	A review; a confirmed case
"Would I tell another owner?"	Owner follow-on (1)	The owner, to a peer	Word of mouth

That's eight judgements across 52 moments. Each one names a voice, which tells you whether the job is press, a reference or a review.

Step 3: gaps, biases and jobs

Moment	Stage	Bias	The job
Triggers, all four choosers	Known	Familiar feels safe	Place the benchmark's findings in the three trade titles assistants cite, with Kestrel's people quoted by name
COO · checking, by sector	Believed	We follow others	One reference per sector, briefed and ready within a week; recruit food and drink and plastics first
Improvement lead · checking	Believed	Losses loom larger	An improvement lead who kept the methods, willing to take calls (Brennock's)
Investor · checking	Believed	We follow others	An operating partner or portfolio COO as a reference
Assistant · "any good?"	Believed	Familiar feels safe	Reviews from finance directors and plant managers on the two platforms assistants cite; the mix-up corrected
Finance director · checking	Believed	Losses loom larger	Reviews and cases where the client's finance director says, in their own words, how the savings were counted
Owner · checking	Believed	We follow others	An owner reference; an operations award judged by plant leaders
Owner · follow-on	Preferred	We follow others	Make referring easy: a paragraph the owner can forward, and a thank-you that isn't a payment

Two rules come out of the jobs. **Ask only after a confirmed result:** a reference or review before the 90-day check is a favour; after it, it's evidence. And **never write the outside voice yourself:** an edited quote, a drafted review or a paid list place is Kestrel talking in a disguise, and assistants, editors and buyers are all getting better at noticing.

Step 4: reconcile, the outside-voice plan

Programme	What it produces	Judgements it settles	Signals	Priority
References	One briefed client per sector, plus an improvement lead, an owner and an investor; ready within a week	Plant like mine; my team; portfolio; owner's trust	Canon, coverage, strategy	P1
Reviews	Reviews from finance directors and plant managers on the two cited platforms, asked for after a confirmed result	Any good?; savings real?	Coverage, strategy	P1
Facts corrected	The mining mix-up fixed; four association entries correct (shared with Being named)	Any good?; worth a look?	Coverage	P1
The benchmark in the press	Findings placed in the three cited trade titles, once a year, with follow-up articles by Kestrel's people	Worth fixing?; worth a look?	Canon, coverage, strategy	P2
Comparison lists	A fact sheet and a sector case sent to the editors of the cited lists	Worth a look?	Canon, coverage	P2
Word of mouth	A forwardable paragraph after every confirmed result; a thank-you note from the managing partner	Would I tell another owner?	Coverage	P2
Awards	One operations award a year, entered with a confirmed case	Owner's trust	Coverage	P3

Seven programmes. Every judgement has one. The references and reviews come first because they're cheap, quick and decided by Kestrel's own clients. The press comes second because it depends on the benchmark, which comes once a year.

Rejected, with reasons:

- **A press release for every win.** No judgement asks for it, and trade editors stop opening them.
- **Paid placements in "best of" lists.** That's advertising; if Kestrel runs it, it's labelled as such in the Advertising playbook.
- **Incentives for reviews.** It breaks the platforms' rules and the claims rule, and a review a buyer suspects is worse than none.
- **A referral fee.** Owners refer because they were helped. Paying turns a favour into a sales job, and changes how the next owner hears it.

The anatomy library

Asset	Parts, in order	Length
The story pitch (to an editor)	The finding, in one sentence, with its number · Why their readers care now · The method, in two lines · Who can be quoted, and their plant experience · What's attached	Under 150 words
The reference brief (to the client giving the reference)	Who'll call, their plant, their question · What they're likely to worry about · What you're happy to say, and what you'd rather not · How long it will take	Half a page
The reference card (to the buyer)	The client: plant, size, sector, problem · What changed, confirmed by whom · Their role, and what they can speak to · How to book the call	A third of a page
The review request	Thanks, for a specific result · The ask: one platform, one link · What helps other buyers (what you were worried about, what happened) · An easy no	Under 120 words

Asset	Parts, in order	Length
The correction request	What's wrong, in one line · What's right, with a link to proof · The exact text to replace it	Under 100 words
The forwardable paragraph (word of mouth)	Who Kestrel is, in the entity line · What it did for them, in their figures · How to get in touch	60 words

Four rules hold for every outside voice:

- **Their words, never yours.** Brief a reference on what the buyer will ask; never on what to say.
- **After the result, not before.** No reference or review before savings are confirmed.
- **Matched to the reader.** A finance director's reference for a finance director; a plant manager's for a plant manager.
- **Easy to say no to.** A client who says yes reluctantly gives a reluctant reference.

Briefs

Two levels: the programme brief (once, for each of the seven programmes) and the asset brief (for each pitch, reference or request).

PROGRAMME	References
JUDGEMENTS	Plant like mine (by sector); my team; portfolio; owner's trust
MOMENTS	COO · checking (4); improvement lead · checking; investor ·
checking;	owner · checking
THE BENCH	Packaging: Brennock (plant manager; improvement lead)
	Food and drink: Calloway (finance director)
	Metal and engineering: a fabricator's managing director
(owner)	
	Plastics: none yet; Hartwell after its 90-day check
	Investor: none yet
READY WITHIN	One week of a request
ASKED	After the 90-day check, by the lead engineer who did the work
LIMIT	No client takes more than four calls a year
OWNER	Managing partner
MEASURE	Days from request to call; deals that move after the call

ASSET	Reference brief: Brennock's improvement lead, for Hartwell's
THE CALLER	Sam Ito, improvement lead, Hartwell Plastics (260 staff)
THEIR QUESTION	"Will they work with us, or replace us?"
LIKELY WORRY	That the review will be read by the COO as a verdict on Sam's team
THE REFERENCE	Brennock's improvement lead, who led the changes and ran the next review without Kestrel
AVOID	Asking them to say anything about price or the other bidders
TIME	20 minutes

The finished example: the benchmark pitch, a reference card and a review request, written in full

Three assets, one from each of the main programmes. Kestrel, its clients, its people and the editor are invented.

1. The story pitch, to the editor of a manufacturing title assistants cite.

Subject: Mid-sized plants are losing 4% to 9% of margin to capacity they already own

Hi Rachel,

Our 2025 benchmark of 38 plant reviews in mid-sized manufacturers found the same thing again and again: output was flat not because lines were too small, but because changeovers, small stops and scheduling were eating the capacity. In the median plant, that cost between 4% and 9% of annual contribution margin. Six of the 38 had already approved a new line when the review started.

Your readers are deciding capital budgets for next year now, and this is the finding most likely to change one.

The method: every figure comes from measurements taken on the floor during a review, and the method note says which savings each client's finance team has confirmed since. The full report and the data by sector are attached.

Our head of delivery, Mark Ellison, ran 14 of the reviews and spent 20 years running plants before that. He's happy to talk this week.

Jo Brady, Kestrel Advisory

2. The reference card, sent to Hartwell's improvement lead.

Your reference: the improvement lead at Brennock Packaging (240 staff, packaging)

What changed: warm-up and changeover losses cut on the carton line; output from 41 to 47 cartons a minute; confirmed by Brennock's finance director. Brennock's improvement team led the changes, and ran the next review, on a third line, without Kestrel. **What they can speak to:** what it was like to have outside engineers on the floor; how the findings were credited; what their team kept. **Booking:** a 20-minute call this week. Reply with two times and we'll confirm.

3. The review request, to Calloway Foods' finance director, after the 90-day check confirmed the savings.

Subject: Thank you, and one small favour

Michael, thank you for confirming the savings on the filling line. Scrap from 4.1% to 2.6% is a result your team made stick, and it's good to have it in your own figures.

If you'd be willing, a short review on [platform] would help other finance directors in your position. What helps most is what you were worried about before the review, and what you found afterwards. The link is below; it takes a few minutes.

And if you'd rather not, that's completely fine: the result speaks for itself.

Priya Raman

What to notice. The pitch leads with a number and a reason to run it now, and offers the method before the editor asks. The reference card tells the buyer exactly what the reference can speak to, so the call is about their question, not a general endorsement. The review request asks after the result, suggests what's useful without suggesting what to say, and makes saying no easy. Three assets; Kestrel needs about 40 of them a year.

Templates

The reference bench (one row per client willing to take calls)

Client	Sector, size	Person and role	Problem solved	Confirmed result	Can speak to	Calls this year (max 4)	Last called

The review tracker (one row per confirmed result)

Client	Person	Confirmed result and date	Asked on	Platform	Reviewed?	Thanked?

The press log (one row per pitch)

Title	Editor	Cited by assistants?	The finding pitched	Sent	Response	Published	Moments it should move

The reference brief

CALLER [name, role, plant, size]
 THEIR QUESTION [in their words]
 LIKELY WORRY [what they'll really want to know]
 THE REFERENCE [who, and what they can speak to]
 AVOID [what the reference shouldn't be asked to cover]
 TIME [minutes]

In the campaign plan

Most campaigns make a claim; this build supplies the outside voice that makes it believable. The rule: **no campaign makes a claim without naming the outside voice that backs it.**

Kestrel campaign	The claim	The outside voice
The annual benchmark	"Mid-sized plants lose 4% to 9% of margin to capacity they own"	Coverage in the three cited trade titles; a quote from a client finance director
"Before you buy another line"	"Find out what this line can do first"	A reference who didn't buy the line (Brennock); a review saying the same
Board season	"The numbers will stand up"	A finance director's review; a finance director reference
Owners' round table	"Owners like you chose Kestrel"	The speaker is the outside voice; the write-up in the association newsletter
Investor 100 days	"We start inside 100 days"	An operating partner reference, once there is one. Until then, the claim isn't made

The last row matters. A campaign can wait for its outside voice. It can't make one up.

Running it

Who owns what. At Kestrel, the managing partner owns the reference bench, because references are relationships, not marketing. The marketing lead owns the press log and the review tracker. Every request to a client comes from the person who did the work.

When	What happens
After every 90-day check that says "yes"	Thank the client; ask whether they'd be a reference, and separately whether they'd write a review
When a deal reaches checking	A reference card to the buyer within two days; the call within a week
After every reference call	Thank the reference; note what was asked, so the next brief is better
Monthly	The review tracker and the facts: any new review, any wrong fact in the scan, any listing out of date
Quarterly	The bench: gaps by sector and role; no client over four calls a year
Yearly	The benchmark to the press; one award entry, if a judged one fits

Word of mouth, as a nod. It can't be run, but it can be made easy and noticed. After every confirmed result, send the client the forwardable paragraph. When a new enquiry names who referred them, the managing partner thanks that person in writing the same week. Kestrel's 22 referrals last year went unthanked.

Common mistakes

The mistake	The stage it breaks	The bias it ignores	At Kestrel	Instead
Press releases about the firm	Known	Familiar feels safe	News of a new hire, sent to nobody in particular	Pitch findings the editor's readers need

The mistake	The stage it breaks	The bias it ignores	At Kestrel	Instead
References "on request"	Believed	We follow others	Two to three weeks to find one; deals cooled	A bench, by sector and role, ready within a week
The same reference for everyone	Believed	We follow others	A COO's reference offered to a finance director	Match the reference to the reader's role and question
Asking before the result	Believed	Losses loom larger	A review request in the final report	After the 90-day check, and only after a yes
Writing the outside voice	Believed	We follow others	Drafting quotes for clients to approve	Their words, unedited; brief them on the question, not the answer
Wearing out the bench	Believed	Losses loom larger	One loyal client taking every call	No more than four calls a year each
Paying for places or reviews	Believed	We follow others	A paid spot in a "best of" list, presented as editorial	Earn the place; if it's paid, it's labelled as an ad
Leaving wrong facts in place	Believed	Familiar feels safe	The mining mix-up, repeated by assistants for a year	Correct the sources, not just the site
Judging by mentions	All, unseen	Good enough beats best	"Featured in three titles" with no moment moved	Judge by the judgement each voice was meant to settle

The bias to watch in the team: **reticence about asking**. Most professional firms feel that asking a client for a reference or a review imposes on them. A client whose savings are confirmed usually wants to help; not asking leaves the next buyer to take Kestrel's word for it.

The business check and the measures

- **Revenue:** references and reviews sit closest to signed work; the press supports the earliest moments. Start with the bench.
- **One truth:** every result a reference or review mentions is a signed claim. If a client says something Kestrel can't sign, it isn't quoted.
- **The proof is real:** no outside voice before a confirmed result. Two of last year's nine reviews have one; the 90-day check is what changes that.
- **Capacity:** a reference call costs a client 20 minutes. Four calls a year each means the bench must grow with the deals.

Measure	What it tells you	Stage
Days from reference request to call	Whether checking stalls (today: two to three weeks)	Believed
Deals that move forward within a week of a reference call	Whether the reference answered the buyer's question	Believed
Reviews on the cited platforms, and their share from finance directors and plant managers	Whether the assistant's "any good?" has an answer	Believed
Assistant answers to "Is Kestrel Advisory any good?" with the right facts	Whether the mix-up is fixed (shared with Being named)	Believed
Coverage in the cited trade titles, and answers that quote it	Whether the benchmark is travelling	Known
Referrals thanked within a week	Whether word of mouth is noticed	Preferred

The Loop for this build

Each quarter the Loop (step 11) reads the bench, the reviews and the press log beside the Score's Believed, and asks: **for each judgement a buyer makes, is there now a voice other than Kestrel's to settle it?**

What the Loop finds	What it means	What changes
Believed rises at checking, but not at comparing	References work inside deals; nothing outside them	Reviews and comparison lists, so the voice reaches buyers before they call

What the Loop finds	What it means	What changes
Reviews appear, the assistant still says "no reviews found"	The platforms aren't the ones assistants cite	Check the scan's citations; move effort to the cited platform
Coverage published, framing answers unchanged	The coverage isn't in a source assistants read, or doesn't carry the finding	Pitch the cited titles; lead with the number
Reference calls held, deals still stall	The wrong reference, or the wrong question	Match by role; brief on the buyer's real worry
A sector with no reference	A gap the next win must fill	Ask the next client in that sector, after their 90-day check

What Kestrel's first year would look for. References within a week by the end of the first quarter, with a plastics reference once Hartwell's savings are confirmed. Reviews from finance directors on both cited platforms by mid-year, and the assistant's answer to "is Kestrel any good?" naming the right firm, with a review. Then the benchmark's coverage, which moves the slowest and the furthest: the moment a trade title quotes Kestrel's 4% to 9%, every assistant answer that cites it carries Kestrel's name into a trigger moment Kestrel has never reached.

Shapes by business type

Business type	Typical mode	The outside voices that usually carry the most weight
B2B professional services	Committee or owner-led	References by sector and role, trade press, peer referrals, a few reviews
B2B software	Committee	Software review platforms, analyst coverage for larger deals, customer case studies, community answers
Considered consumer (finance, health, education)	Considered personal	Review sites, comparison sites, consumer press, regulators' ratings
Ecommerce	Quick choice or habitual	Product reviews and ratings, creators' reviews, "best of" lists

Business type	Typical mode	The outside voices that usually carry the most weight
Local services	Owner-led or considered personal	Map reviews, neighbours' recommendations, local community groups
Marketplaces	Varies on each side	Reviews on both sides, press about trust and safety, sellers' stories

Feeling a bit overwhelmed?

You should be, a little. One firm: eight judgements across 52 moments, seven programmes, a reference bench by sector and role that can answer within a week, reviews asked for only after confirmed results, a benchmark placed in the titles assistants read, and wrong facts chased down wherever they've spread. Most firms send a press release when they hire someone and call it PR.

RocMode does part of the heavy lifting:

- **Where the outside voices count.** Its scan records which sources the assistants cite when your buyers ask, so you know which titles, lists and platforms matter in your category.
- **What's being said about you.** Across five AI assistants, it records whether you're named and who's named instead, moment by moment.

[Join the RocMode beta](#)

The Partners and Channels Playbook

What this build is for

Partners and channels reach the buyers who already trust someone else, through that someone.

Some buyers are hard to reach directly, and harder to persuade. An investor's operating partner doesn't search for consultants; they ask the transaction adviser who did the deal, or the lender who financed it, or a portfolio COO they trust. An owner asks their accountant. Procurement checks the panel the government programme approved. In each case, the buyer's question goes to someone who already has their trust, and that person names a firm, or doesn't.

This is the Partnered category: someone else's voice or audience, by agreement. It differs from PR and reviews in one word: agreement. A reference is freely given. A partner has a reason to recommend you that both sides understand, and the buyer should be able to understand it too.

For Kestrel, this build is small in moments and large in value: it leads the investor's path, where one relationship can bring a portfolio.

The Partnered category's four nodes live here: **affiliates** (paid referral arrangements, rarely right for professional services), **marketplaces** (directories where buyers compare and book), **associations** (industry bodies whose member lists and events reach the category) and **employee advocacy** (Kestrel's own people, speaking in their own names).

The footprint, and what it reads from the chain

Partners and channels lead **3 cells and 3 moments**, the investor's whole path outside the deal:

Moment	Who	Their question
Trigger	Investor	"The value-creation plan needs margin within 18 months."
Options	Investor	"Which firms deliver operations value quickly?"
Extend and refer	Investor	"Can we use them across the portfolio?"

It supports four more: the COO's options (where an automation integrator or equipment supplier can say "fix the line first"), procurement's comparing (where approved-supplier panels decide who can be compared at all), and the investor's comparing and checking.

Three moments led and ten supported, out of 149. But the investor moments carry an outsized prize. One operating partner with a portfolio of eight manufacturers is eight potential clients, each with a COO who'll take the call because their investor made it.

What it reads, and never rewrites:

From	It takes
Brand Strategy	The segments: investor-owned manufacturers are the second segment
GTM Strategy	The investor-led buying path and its motion
Credentials	The claims a partner may repeat, word for word
The pitch	The investor pack, which partners forward
PR and reviews	References a partner can offer their clients

Its slice of the Score, and its own audit

This build's slice of [the Selectability Score](#) is the gate at the investor's moments: **Known** and **Found**, through someone the investor already trusts.

Kestrel today	Known	Found	Quality	Moment score
Investor · trigger	6	5	27	1.5
Investor · options	6	0	12	0
Investor · extend (inside a deal)	Open	Open	35	35

The single point of Found comes from one adviser's list, which names Kestrel among a dozen firms. Investors don't find Kestrel because nobody they trust is pointing at it.

The audit.

Possible partner	Who they reach	Kestrel today
Private equity operating partners	Portfolio COOs and boards	No relationships; one past client was investor-owned

Possible partner	Who they reach	Kestrel today
Transaction advisers and lenders	Investors in the first 100 days after a deal	Named on one adviser's list, unprompted
Accountants to mid-sized manufacturers	Owners and finance directors	Two partners know one firm socially
Automation integrators and equipment suppliers	COOs weighing a new line	None; seen as competitors
Government productivity programmes	Owners; procurement	Not on the approved panel
The four sector associations	Every chooser in the sector	Member of one, inactive

None of these relationships has an owner, a reason for the partner, or a record of what it produced.

Step 1: the shape, the partner tiers the category runs on

The shape is **the partner types the category already uses**, sorted into tiers by how closely the partner works with you.

Tier	What it means	Partner types in Kestrel's category	Canon?
Referral partners	They introduce; you sell and deliver	Transaction advisers, lenders, accountants	Yes: most firms have a few
Strategic partners	They bring you into their own client work, by agreement	Private equity operating teams; automation integrators	Large firms only
Panels and programmes	You're approved to be chosen	Government productivity programmes; approved-supplier panels	Yes, for public and large buyers
Associations (nod)	Their members, their events, their lists	The four sector associations	Yes
Marketplaces (nod)	Buyers compare and book	Consultant directories	Rarely decisive in this category
Affiliates (nod)	Paid per referral	Rare; seen as a conflict by buyers	No

Every partnership is an exchange. The partner gets something they value: a better outcome for their own client, a reason to stay in touch, their own reputation enhanced. If you can't say what the partner gets, there is no partnership, only a list of people you hope will mention you.

Affiliates, as a nod. For professional services, paying per referral turns a trusted recommendation into a commission the buyer may never be told about. If you pay, say so to the buyer. Kestrel doesn't pay.

Step 2: coverage, one partner per audience you can't reach yourself

The coverage unit is **an audience you can't reach well on your own, and the partner who already has its trust**. For each of the seven cells this build leads or supports, name the chooser, who they ask, and what that person would need to recommend Kestrel.

Moment	Chooser	Who they ask	What the partner would need
Trigger	Investor	Their own operating team; the transaction adviser	Evidence that operations reviews move margin inside a hold period

Moment	Chooser	Who they ask	What the partner would need
Options	Investor	Operating partners at other funds; advisers; lenders	A firm they can name without risk to their own reputation
Comparing	Investor	Portfolio COOs	Portfolio cases, in board terms
Checking	Investor	Another operating partner	A reference (with PR and reviews)
Extend and refer	Investor	Their own portfolio	A way to roll out across companies without a new tender each time
Options	COO (7 problems)	Automation integrators; equipment suppliers	A reason to say "fix the line first": a better installation later, a client who trusts them more
Comparing	Procurement	The approved panel	Kestrel on it

Four partner types cover all seven: operating partners and their funds, transaction advisers and lenders, automation integrators, and approved-supplier panels. Accountants and associations add reach to owners, in support of Events.

Step 3: gaps, biases and jobs

Moment	Stage	Bias	The job
Investor · trigger	Known	Familiar feels safe	Be known to three operating partners as the firm for margin in the first 18 months, through one portfolio case and the benchmark's investor cut
Investor · options	Found	We follow others	Be on the shortlist transaction advisers and lenders hand to new owners in the first 100 days
Investor · comparing and checking	Believed	We follow others	Portfolio cases in board terms; an operating partner as a reference

Moment	Stage	Bias	The job
Investor · extend	Bought	A hard choice stalls	A portfolio framework: one agreement, a fixed price per plant, start inside 100 days
COO · options	Preferred	Losses loom larger	One automation integrator who refers "fix the line first" cases, and to whom Kestrel refers clients who do need new equipment
Procurement · comparing	Bought	A hard choice stalls	A place on the government productivity programme's approved panel

One rule follows from the exchange: **refer back**. Kestrel finds plants that do need new equipment, or a finance review, or a new hire. A partner who receives good work from Kestrel sends good work back.

Step 4: reconcile, the partner plan

Partner	Tier	What they get	What Kestrel gets	Moments	Priority
Three PE operating teams	Strategic	Margin found inside the hold period; board reporting in their terms; a portfolio framework	Introductions to portfolio COOs; references	Investor trigger, options, extend	P1
Two transaction advisers, one lender	Referral	A firm to name in the first 100 days that makes them look good	Introductions to new owners	Investor options	P1
The government productivity programme	Panel	A delivery partner that reports results their way	Eligibility; procurement's comparison	Procurement comparing; owner options	P2
One automation integrator	Referral, both ways	Clients sent to them who do need equipment; better-prepared installations	"Fix the line first" referrals	COO options	P2
The four sector associations (nod)	Association	Speakers, benchmark data for members	Member lists, events, newsletters	Owner options; triggers (with Events)	P2

Partner	Tier	What they get	What Kestrel gets	Moments	Priority
Accountants to manufacturers	Referral	A firm to name when a client's margins fall	Introductions to owners and finance directors	Owner and finance director triggers (support)	P3

Six partner groups, about a dozen relationships. Each has an owner at Kestrel, an agreement in writing (even if it's a one-page letter), and a review twice a year.

Rejected, with reasons:

- **Referral fees.** See the affiliates nod. Buyers trust a recommendation less when they learn it was paid, and they usually learn.
- **White-label work for large consultancies.** It earns fees and builds someone else's name at Kestrel's moments.
- **A listing on every consultant marketplace** (nod). Marketplaces aren't where Kestrel's buyers choose. One correct listing on the two that assistants cite is enough, and that's Being named's job.
- **Partnerships with equipment makers who'd want Kestrel to recommend their machines.** It breaks the reason to choose Kestrel: an independent view of the line.

The anatomy library

Asset	Parts, in order	Length
The partner one-pager (for the partner)	Who Kestrel helps, and who it doesn't · When to introduce us, in their client's words · What their client will get, and how fast · What the partner gets · Who to call	One page
The introduction note (the partner forwards it)	The entity line · Why this client, now, in the partner's terms · One relevant case · The smallest next step	Under 120 words
The partner agreement	What each side does · What each side gets · What's never done (fees, exclusivity, recommending products) · How it's reviewed	One page
The portfolio framework (investors)	One agreement for the portfolio · Fixed price per plant, and the way out · Start inside 100 days · Board reporting format	Two pages

Asset	Parts, in order	Length
The referral-back note	Who Kestrel is sending, and why · What the client needs from them · What Kestrel found that's relevant	Half a page
The partner review (twice a year)	Introductions each way, and what happened · What their clients said · What to change · Continue or end	One page

Three rules hold for every partnership:

- **The buyer always knows the relationship.** "We work with Kestrel" is said out loud.
- **The partner repeats only signed claims.** The one-pager is the only thing they need, and it comes from Credentials.
- **It runs both ways, or it ends.** A partnership with introductions in one direction only is a favour, and favours run out.

Briefs

PARTNER	Northgate Capital, operating team (invented)
TIER	Strategic
OWNER	Jo Brady, partner
WHO THEY REACH	8 portfolio manufacturers, 120 to 450 staff; COOs and boards
THEIR QUESTION	"Which firms deliver operations value quickly?"
WHAT THEY GET	Margin found inside the hold period; reporting in value-creation
	terms; one framework instead of eight procurements
WHAT WE GET	Introductions to portfolio COOs; an operating partner reference
CLAIMS THEY MAY USE	Fixed price; six weeks; the way out; start inside 100 days;
	Brennock and Calloway cases
NEVER	Fees; exclusivity; Kestrel recommending their other suppliers
REVIEW	March and September
MEASURE	Introductions, reviews started, portfolio companies served

The finished example: the investor channel, written in full

Three assets for one strategic partner: the one-pager the operating team keeps, the note they forward to a portfolio COO, and the framework that lets the portfolio buy without eight procurements. Northgate, its portfolio and its people are invented; so is Kestrel.

1. The partner one-pager, for Northgate's operating team.

Kestrel Advisory: margin from the lines you already own, inside the hold period

Who we help. Manufacturers with 100 to 500 staff in food and drink, packaging, metal and engineering, and plastics and rubber, where output is flat or costs are rising. **Who we don't:** businesses whose problem is demand, not capacity. We'll say so on the first call.

When to introduce us. When a portfolio company's value-creation plan needs operating margin within 18 months, and the first instinct is new equipment or a large programme.

What your company gets. A fixed-price, six-week review on the floor, starting inside your first 100 days. If it doesn't find more than it costs, they don't go ahead. Calloway Foods, an investor-owned business of 310 staff, cut scrap on its main filling line from 4.1% to 2.6%, confirmed by its finance director.

What you get. Findings reported in value-creation terms, monthly, for your board pack; one framework for the whole portfolio.

Call: Jo Brady, partner, Kestrel Advisory.

2. The introduction note, written by Kestrel, forwarded by Northgate's operating partner to the COO of a portfolio company.

Subject: An operations review for Ridgeway, before the line decision

Alex, as we discussed at the board, Ridgeway's output on the extrusion lines is flat and the capex request is on the table. Before it goes further, I'd like Kestrel Advisory to look at the lines. They do a fixed-price, six-week review on the floor, and if it doesn't find more than it costs, you don't go ahead. They're independent of the equipment makers. Jo Brady will call you this week, unless you'd rather not. — Claire

3. The portfolio framework, in summary.

One agreement, eight companies. Each review at a fixed price per plant, set by the number of lines, with the same way-out clause. A start within 100 days of a portfolio company's request. Monthly reporting to the operating team in their format. Either side can end the framework with 30 days' notice; any review under way finishes.

What to notice. The one-pager tells the partner when **not** to introduce Kestrel, which is why they trust it enough to use it. The introduction note is in the operating partner's voice, written for them, and names the decision it should change. The framework removes

the procurement step that would otherwise stall every introduction. Three documents, one partner. Kestrel's plan has about a dozen.

Templates

The partner map (one row per partner)

Partner	Tier	Who they reach	Moments	What they get	What Kestrel gets	Owner	Introductions in / out this half	Review date

The partner brief

```

PARTNER      [name]          TIER  [referral / strategic / panel /
association]
OWNER        [Kestrel person]
WHO THEY REACH [audience, size, roles]
THEIR QUESTION [what their clients ask them]
WHAT THEY GET  [the exchange, from their side]
WHAT WE GET    [the exchange, from ours]
CLAIMS THEY MAY USE [signed claims only]
NEVER          [fees, exclusivity, product recommendations, ...]
REVIEW        [dates]          MEASURE [introductions, both ways, and
outcomes]
```

The introduction log (one row per introduction, either direction)

Date	From	To	Client	Moment	What happened	Thanked?

In the campaign plan

Partners carry campaigns to audiences Kestrel can't reach, but only campaigns that help their own clients. The rule: **a campaign goes to a partner with a version made for their clients, and nothing they'd be embarrassed to forward.**

Kestrel campaign	What partners get	Which partners
The annual benchmark	An investor cut (margin found inside a hold period); a members' summary for associations	Operating teams; associations

Kestrel campaign	What partners get	Which partners
Investor 100 days	The one-pager and a dated note for new acquisitions	Operating teams; transaction advisers; lenders
"Before you buy another line"	The equipment comparison, for clients weighing capex	The integrator; accountants
Owners' round table	Seats for their owner clients	Associations; accountants

Running it, with the nods

Who owns what. At Kestrel, each partner has one owner among the partners; the managing partner owns the partner map. Introductions in either direction go in the log the day they happen.

When	What happens
Within a week of every introduction	Thanks to the partner, and a note on what happened
Monthly	The introduction log reviewed: who's introducing, who isn't, who Kestrel should be introducing to
Twice a year	Each partner review: introductions both ways, outcomes, continue or end
Yearly	The partner map against the moments: is every audience Kestrel can't reach covered?

Associations, as a nod. Join the four sector associations as an active member: speak at one event a year each, give members the benchmark's sector cut, and keep the member-list entry identical to the fact sheet. That's the whole programme.

Marketplaces, as a nod. One correct listing on the directories assistants cite, maintained by Being named. Not a channel to sell through.

Employee advocacy, as a nod. Kestrel's engineers know more about plant floors than anyone in marketing. Encourage them to answer practitioner questions and write under their own names, with no scripts, no required posting and no approval beyond the claims rule: nothing said that Credentials can't sign.

Common mistakes

The mistake	The stage it breaks	The bias it ignores	What it would do at Kestrel	Instead
A logo wall of "partners"	Believed	We follow others	Twelve logos, no introductions	A dozen real exchanges, each with an owner
Partnerships that run one way	Found	Losses loom larger	Asking advisers for introductions, sending them nothing	Refer back; log both directions
Paying for referrals quietly	Believed	We follow others	A fee the buyer never hears about	No fees; or, if paid, disclosed
Partners with a stake in the answer	Preferred	Feeling first, reasons after	An equipment maker's preferred consultant	Partners whose interest matches the client's
Giving partners nothing to forward	Answered	A hard choice stalls	"Mention us if it comes up"	The one-pager and the introduction note
Letting partners say more than is signed	Believed	Familiar feels safe	"They guarantee 20% savings", repeated by an adviser	Only signed claims, in the partner's materials
Procurement as an afterthought	Bought	A hard choice stalls	Eight portfolio introductions, eight procurements	The portfolio framework first
No review	All, unseen	The planning fallacy	Partnerships that quietly lapse	Twice a year: continue or end

The bias to watch in the team: **reciprocity, misread**. Partners don't introduce because they were asked nicely; they introduce because it was good for their client last time. The first introduction is a test. Make it go well.

The business check and the measures

- **Revenue:** investor-owned manufacturers are the second segment; one strategic partner can reach a portfolio. Start with operating teams.
- **One truth:** partners use only the one-pager's claims. If a partner says something Credentials hasn't signed, correct it with them the same week.
- **Independence:** no partnership may compromise Kestrel's independent view of the line. That rules out equipment makers and paid placements.
- **Capacity:** a portfolio framework could bring several reviews at once. Agree a maximum number of concurrent starts before signing one.

Measure	What it tells you	Stage
Introductions received, by partner	Whether partners are naming Kestrel	Found
Introductions that become fit calls, then reviews	Whether the introductions are the right buyers	Bought
Introductions Kestrel sends partners	Whether the exchange runs both ways	Preferred
Portfolio companies served per strategic partner	Whether the investor channel compounds	Preferred
Investor moments in the Score: Known and Found	Whether investors find Kestrel without being sent	Known, Found

The Loop for this build

Each half-year the Loop (step 11) reads the introduction log beside the investor moments in the Score, and asks: **which partners are bringing the right buyers, and which exchanges have stopped running?**

What the Loop finds	What it means	What changes
Introductions arrive, few become reviews	The partner is introducing the wrong buyers, or too early	Sharpen "when to introduce us" on the one-pager
A partner introduces nobody	No exchange, or no reason to remember Kestrel	Send them something useful; if nothing changes in two reviews, end it

What the Loop finds	What it means	What changes
Portfolio introductions stall at procurement	The framework is missing or unclear	Fix the framework before the next introduction
Investors start finding Kestrel without a partner	The channel is building Known	Keep it; it's working beyond the introductions

What Kestrel's first year would look for. One strategic partner with a signed framework and two portfolio reviews by the end of the year; two referral partners introducing at least once a quarter; a place on the government panel. Small numbers, but each one reaches buyers Kestrel's other builds can't.

Shapes by business type

Business type	Typical mode	The partners that usually matter most
B2B professional services	Committee or owner-led	Investors' operating teams, advisers and lenders, accountants, associations, panels
B2B software	Committee	Resellers and integrators, technology alliances and their marketplaces, consultancies that implement
Considered consumer (finance, health, education)	Considered personal	Brokers and advisers, comparison sites, employers and professional bodies
Ecommerce	Quick choice or habitual	Retailers and marketplaces, affiliates (disclosed), complementary brands
Local services	Owner-led or considered personal	Trades and suppliers who refer, estate agents, local business groups
Marketplaces	Varies on each side	Platforms and tools sellers already use, communities on each side

Feeling a bit overwhelmed?

You should be, a little. One firm: six partner groups, about a dozen relationships, each with an exchange both sides can name, a one-pager, an introduction note, an agreement, a log in both directions and a review twice a year, plus a portfolio framework so introductions don't die in procurement. Most firms call everyone they've ever had lunch with a partner, and wonder why nobody introduces them.

RocMode does part of the heavy lifting:

- **A channel training kit, if you sell through partners.** When your sales motion is partner-led, RocMode's sales kit includes a channel training kit, built from your strategy and the claims it can back up.
- **Who you win, and where they ask.** Your choosers and their moments, mapped from what's known about your brand, so you can see which buyers you'll only reach through someone else.

[Join the RocMode beta](#)

The Creators Playbook

What this build is for

A creator sells belief you can't make yourself, so you're renting it, and the rules of belief apply.

Buyers believe a creator because the creator has been right before, about things they care about. That belief is the creator's, not yours. You can rent it for a while. You can't own it, and a script spends it.

What a brand says splits into facts and judgements. Buyers accept facts from you, but not judgements from the party who gains from them. "Kestrel's engineers work on the floor with your people" is a claim from Kestrel. From a retired plant manager who's watched them do it, it's a recommendation.

So a creator does three jobs, in this order:

1. **Vouch** (Believed): would the audience trust this person on this subject, unpaid?
2. **Show the reason** (Preferred): does the work show why Kestrel, in a way someone could repeat?
3. **Put Kestrel in a moment** (Known): do the right buyers now link the name to their situation?

In business-to-business categories, the creators that matter are rarely called influencers. They're practitioners with a following: a retired plant manager who writes about the floor, a podcast host who interviews COOs, an engineer whose YouTube videos half the industry has watched. Their audiences are small and exact.

This build is in the Partnered category: other people's voices, by arrangement. Its rule is the same as every build's: **a creator is paid because a moment needs someone else's belief**. A big following isn't a reason. Affiliates, paid per sale, are Partners and channels' nod, not this build.

The footprint, and what it reads from the chain

Creators lead the smallest footprint of any build: **1 cell and 1 moment**.

Moment	Who	Their question
Plant manager · framing	Plant manager	"Will they just tell us what we already know?"

It's one moment, but it's the one where a deal is often lost without anyone noticing. A plant manager who expects consultants to tell them what they already know will be polite in the meeting and quietly against the review afterwards. Nothing Kestrel says about itself answers that fear. Someone the plant manager already trusts has to have seen Kestrel's people on a floor.

It supports **5 more cells (17 moments)**:

- **Owner · trigger**, for the seven problems (Outbound leads): owners listen to other owners, and to the few voices who speak to them.
- **Plant manager · trigger**, for the seven problems (Social leads): a practitioner's post about firefighting reaches plant managers who'd never follow a consultancy.
- **Owner · options**, "Do we need outside help at all?" (Events leads): a podcast conversation with an owner who brought someone in.
- **Plant manager · comparing**, "Will they work on the floor with my people, or from a meeting room?" (the pitch leads, inside a deal): the creator's account of a day on a floor is something a plant manager can check for themselves.
- **Owner · checking**, "What do people I trust say about them?" (Events leads, inside a deal).

What it reads, and never rewrites:

From	It takes
Brand strategy	The plant manager and the owner as choosers, with their fears
Credentials	The facts a creator must get right: what Kestrel does, what it costs, and the signed results
Customers	The clients willing to host a creator on their floor
Social	The one-question videos a creator may share or answer
Website	The page a creator links to, which must say the same things

Its slice of the Score, and its own audit

Kestrel today	Known	Found	Quality	Moment score
Plant manager · framing (out in the market)	6	0	9	0
Owner · trigger, each of 7 problems (out in the market)	12	0	9	0
Plant manager · trigger, each of 7 problems (out in the market)	6	0	9	0
Owner · options (out in the market)	12	10	28	3
Plant manager · comparing (inside a deal)	Open	Open	56	56
Owner · checking (inside a deal)	Open	Open	61	61

The pattern is the same as Events'. Inside a deal, when a plant manager or an owner meets Kestrel's people and its clients, Kestrel does well: 56 and 61. Out in the market, before anyone has met them, it scores nothing. A creator is one of the few ways to carry what happens inside a deal out into the market, in someone else's words.

The audit. Kestrel has worked with one creator:

What	What Kestrel did	What it produced
A manufacturing "influencer" with 85,000 LinkedIn followers	Paid A\$2,500 for one post about the trade-show stand, written by Kestrel's agency	40 likes, mostly from the creator's own network; no comment from anyone in a plant; the creator's followers were mostly other consultants and vendors
Practitioners with a following	None approached	—
Podcasts	None	—
Brennock's improvement lead	Posted about the review, unpaid and unprompted	Kestrel's most-commented post of the year (see the Social Playbook)

The paid post failed every test: the wrong audience, Kestrel's words in someone else's mouth, and a subject (a stand) nobody was deciding anything about. The unpaid post from a client did what the paid one was meant to.

Step 1: the shape, the category's voices

The shape is **the people Kestrel's choosers already listen to**, by type, with who actually follows them. Count by audience, not by followers.

Type of voice	Typical audience	What the audience trusts them on	In Kestrel's category	Kestrel today
Practitioners on LinkedIn: retired plant managers, operations directors	Plant managers, supervisors, COOs	What works on a floor; what consultants get wrong	A dozen with 5,000 to 30,000 followers each	None
Lean and improvement educators on YouTube	Engineers and improvement leads, worldwide; many students	Methods: changeovers, root cause, flow	Three or four large channels	None
Manufacturing podcasts	COOs, owners, improvement leads	Who's doing what, and how it went	Five or six regular shows	None
Owners' circles: business-owner podcasts and newsletters	Owners of mid-sized firms	Running a business, not running a line	A handful, by region	None
Industry "influencers" with large general followings	Consultants, vendors, students	Industry news and opinion	Many	One paid post
Trade-press columnists	COOs, owners	Opinion	A few; these belong to PR, not this build	PR's

Size isn't the measure. Simone Wies, Alexander Bleier and Alexander Edeling studied 802 Instagram campaigns and found engagement with sponsored posts rose with follower count, then fell: a larger following means a weaker tie with each follower ([Wies, Bleier and Edeling](#)). In a niche B2B category, a retired plant manager with 12,000 followers, most of them plant managers, beats a general manufacturing account with 85,000 followers, most of them selling something.

Step 2: coverage, one trusted voice per moment

The coverage unit is **a voice the chooser already trusts, on the subject of the moment**. For each moment, ask who the chooser would believe, unpaid, and whether that person's audience is really made of those choosers.

Kestrel checked five candidates the same way: by reading 50 recent commenters on each one's posts (or 50 listener reviews and LinkedIn followers, for a podcast) and counting who actually works where Kestrel's choosers work.

Candidate (invented)	Following	Of 50 sampled, in Kestrel's choosers	Fits which moments	Verdict
Ray Hollins , retired plant manager, writes on LinkedIn about running floors	12,000	33: plant managers, supervisors, two COOs	Plant manager framing, trigger and comparing	Yes
"Floor Time" , a manufacturing podcast; host Dana Mercer interviews COOs and owners	About 6,000 listeners an episode	28: operations leaders and owners in manufacturing	Owner options and trigger; COO trigger	Yes
An owners' newsletter , regional, for owners of mid-sized firms	3,500 subscribers	19 owners, 7 of them manufacturers	Owner trigger	Later, if the first two work
A large lean YouTube channel	240,000 subscribers	9: most commenters abroad, or students	Plant manager trigger, in theory	No: the audience isn't Kestrel's
The general manufacturing account Kestrel paid before	85,000	4: mostly consultants and vendors	None	No

Two voices cover the moments that matter. **Ray** answers the plant manager's fear, because he's been one, and plant managers know it. **Dana's podcast** reaches owners and COOs at the moment they're wondering whether outside help is worth it, in a conversation they'll listen to on a long drive.

The owner-checking and plant-manager-comparing moments, inside a deal, need no new creator work. What Ray and Dana publish is there for a buyer to find when they look Kestrel up, and the pitch can point to it.

Step 3: gaps, biases and jobs

Moment	Stage	Bias	The job
Plant manager · framing	Believed	We follow others	Ray spends a day on the floor of a review in progress (with the client's agreement) and writes what he saw, in his words, including anything he didn't like
Plant manager · trigger	Known	Familiar feels safe	Ray's series: four posts over six months, each on a firefighting problem he's seen solved, Kestrel named where it was involved
Plant manager · comparing (in a deal)	Believed	Losses loom larger	Ray's floor-day post linked from the pitch's "how we work on the floor" page, so a plant manager can read an outsider's account
Owner · options	Believed	We follow others	A "Floor Time" episode: Dana interviews Helen Dorsey, a client owner, on when an owner needs outside help, with Kestrel's head of delivery for the second half
Owner · trigger	Known	A hard choice stalls	A second episode, on the benchmark's finding, with Mark Ellison and a client COO
Owner · checking (in a deal)	Believed	We follow others	The episodes linked from the references page, where an owner checking Kestrel will look

One pattern: **the creator sees the work, then says what they think**. Kestrel provides access and facts. The creator provides the judgement, and the judgement is only worth renting if Kestrel doesn't write it.

Step 4: reconcile, the creators plan

Partnership	What	When	A\$	Moments	Signals	Priority
Ray Hollins	A day on a review floor, then four posts over six months in his own words; reviewed after three	Starting with the next review whose client agrees	7,000	Plant manager framing, trigger, comparing	Coverage, strategy	P1
"Floor Time"	Two paid episodes, disclosed at the start: an owner on outside help; the benchmark	One before the spring round table, one after the benchmark's release	5,000	Owner options and trigger; owner checking	Coverage	P2
The owners' newsletter	One piece by a client owner, if the first two show results	Next year	—	Owner trigger	Coverage	P3

A\$12,000 in the first year, plus about four days of Kestrel's time: hosting Ray's floor day, preparing the podcast guests, checking facts.

Rejected, with reasons:

The general manufacturing account again. 4 of 50 commenters were Kestrel's choosers.

The large lean YouTube channel. A big, loyal audience, mostly abroad or still studying. Kestrel's engineers can answer questions in its comments in their own names, as Social's forum work; no payment.

Scripted posts. The fee buys access to an audience's trust; a script turns that trust into an advertisement. Kestrel gives facts and access, never words.

Paying clients to post. A client who's paid to praise you is a creator, and has to disclose it, and their colleagues will discount it. Clients who post do so unpaid, as in the Social Playbook.

Affiliate codes or referral fees. That's Partners and channels' nod, and Kestrel doesn't pay for referrals.

The anatomy library

A partnership is a programme, not a post: choose the audience once, get four things right in each piece, then repeat over months and review.

Part	Stage	What it does
The right audience	Known	Chosen by who follows the creator, from a sample of 50, not by how many
Clear disclosure, first	Believed	"Paid partnership with Kestrel Advisory" in the first line, in plain words
The creator's own words	Believed	Their voice, their experience, their caveats, including what they didn't like
Something true to show	Believed, Preferred	A floor to walk, a client to talk to, a number they can check themselves
A path to check Kestrel	Answered, Bought	What they find when they look Kestrel up matches what the creator said: the same price, the same claim, the same proof
A series, reviewed	Known	The same creator, returning over six months; reviewed after three pieces; stopped if it isn't working

By format:

Format	Parts, in order	Length
The floor-day post	Disclosure · What the creator expected · What they saw, hour by hour or in three scenes · What surprised them, good and bad · What they'd tell a plant manager weighing it up	300 to 500 words, photos with permission
A series post	Disclosure · A firefighting problem in the creator's own experience · What they've seen fix it, Kestrel named where it was involved · What to check	Under 250 words
The podcast episode	Disclosure, in the host's words, at the start · The guest's story, told to the host's questions · Kestrel's person for the second half · The host's own view at the end	35 to 45 minutes

Disclosure lowers persuasion. Martin Eisend and colleagues pooled 61 papers and found it "reduced brand attitudes, credibility, and source evaluation" ([Eisend and colleagues](#)). The rules require it anyway, and an audience that finds out later trusts both of you less. Disclose early and plainly, and win the belief back with what the creator shows.

Briefs

The creator brief carries facts that must be right and leaves the words to the creator. This is the brief Kestrel gave Ray Hollins.

CREATOR BRIEF	Ray Hollins: a day on the floor, then a series
WHY YOU	Plant managers believe you about floors. We'd like you to see how we work and say what you think, good or bad
THE MOMENT	A plant manager wondering: "Will they just tell us what we already know?"
THE ACCESS	One day at a client's plant, week 3 of a six-week review, with the client's written agreement. Walk the floor with our lead engineer and the client's plant manager. Ask anyone anything
FACTS TO GET RIGHT	What a review is: six weeks, on the floor, at a fixed price (A\$45,000 to A\$90,000 by number of lines). Any result you quote: ask us for the signed figure, or ask the client
YOUR WORDS	Entirely. We won't see posts before they go out, except to check a client's name or number if you use one
YOU MAY SAY	What you didn't like. What you'd do differently. That someone else would suit some plants better
DISCLOSURE	"Paid partnership with Kestrel Advisory" in the first line of every post
THE SERIES	The floor-day post, then three posts over six months on firefighting problems you've seen fixed; reviewed after three
FEE	A\$7,000 for the day and four posts, paid whatever you say
WE'LL MEASURE	Who comments; plant managers who mention your posts to us; nothing else is asked of you

The finished example: Ray's day on the floor, written in full

Ray's floor-day post, as he wrote it, and the opening of the first "Floor Time" episode. Kestrel, Hartwell Plastics, Ray Hollins, Dana Mercer and everyone else are invented. Hartwell agreed in writing to the visit and to being named.

1. Ray's post, on LinkedIn, in week 3 of Hartwell's review.

Paid partnership with Kestrel Advisory. They paid me for a day on a floor and a few posts. They haven't seen this one.

I've spent 30 years on the receiving end of consultants. Most of them told me what I already knew, in slides, and left. So when Kestrel asked me to watch one of their reviews and write what I thought, I said yes mainly to find out if they were any different.

I spent Tuesday at Hartwell Plastics, on lines 2 and 4, in week three of six. Three things stood out.

They were on the floor, not in the meeting room. Their lead engineer, Priya, spent the morning at the presses with Dave Okafor, the plant manager, and two setters. I counted: in eight hours, about forty minutes in an office.

They didn't bring the answer. Dave told me this annoyed him in week one. He wanted to be told what to fix. What they did instead was measure scrap by shift and by tool for two weeks, with his setters doing the counting. By week three the setters had found the two tools behind most of the scrap. Dave's view now: "If they'd told us on day one, we'd have argued. We found it, so we believe it."

Not everything was good. The first morning's workshop had too many sticky notes and not enough time on the floor. Two of the operators I spoke to still weren't sure why the review was happening. I said so to Priya, and she wrote it down. I'll believe it's fixed when I see it on the next one.

Will they just tell you what you already know? On Tuesday, no. They told Hartwell's people how to find out what they didn't know, and let them find it. Whether it holds is the question I'd ask in a year, and I'll go back and ask it.

2. The opening of "Floor Time", episode 1, in Dana's words.

"Before we start: this episode is a paid partnership with Kestrel Advisory, an operations consultancy. They suggested the topic and the guests. They don't get to edit what anyone says.

My guest is Helen Dorsey, managing director of a 230-person fabricator. A few years ago she had to decide whether her problems were ones her own team could fix, or whether she needed someone from outside. Later on we'll hear from Kestrel's head of delivery, Mark Ellison, who worked with her. But first: Helen, when did you first think you couldn't fix it yourselves?"

Dana's questions for Helen, agreed in advance in outline only: what she tried first; what made her bring someone in; what it cost in time and disruption, not just fees; how her people reacted; what she'd do differently; when an owner shouldn't bring anyone in.

3. Where it's linked. Ray's post is linked from the website's "How we work on the floor" page, under "What an outsider saw", and from the pitch's answer to the plant manager's comparing question. The episode is linked from the references page, beside client references.

What to notice. Ray's first line is the disclosure, and his second says Kestrel didn't see the post: both make the rest more believable, not less. He says what annoyed the plant manager and what he didn't like, which is why a plant manager reading it will believe the part that's good. The judgement ("on Tuesday, no") is his; Kestrel couldn't have written it, and wouldn't be believed if it had. And the post answers the moment's question in its own words, which is the one thing the brief asked for.

Templates

1. The creator fit scorecard. Score each candidate 0, 1 or 2 on each line before talking money. Under 7 means no. The scale is a working rule, not tested against sales. Kestrel's scores for Ray and for the account it paid before:

Test	Stage	Question	0	1	2	Ray Hollins	The general account
Audience fit	Known	Of 50 sampled followers or commenters, how many are your choosers?	Under 10	10 to 24	25 or more	2 (33)	0 (4)
Trusted on this subject	Believed	Would the audience believe them on this subject, unpaid?	No	On related subjects	Yes, it's what they're known for	2	0
Their own voice	Believed	Do they write their own posts and say what they don't like?	Rarely	Sometimes	Always	2	1
Disclosure record	Believed	Do their past paid posts disclose plainly and first?	No, or late	Inconsistently	Always	2	1

Test	Stage	Question	0	1	2	Ray Hollins	The general account
Comments ask the buying question	Preferred	Do commenters ask "how did you fix it?" or "who did you use?"	Never	Sometimes	Often	2	0
Repeatable	Known	Would they do a series, and be reviewed?	No	Maybe	Yes	2	1
Total						12/12	3/12

2. The creator brief.

```

CREATOR BRIEF [creator: programme]
WHY YOU       [what their audience trusts them on]
THE MOMENT    [the chooser's question, in their words]
THE ACCESS    [what they'll see, with whom, with whose agreement]
FACTS TO GET  [what you do, what it costs, the signed results; where
RIGHT        to check any number]
YOUR WORDS    [entirely theirs; what you will and won't check]
YOU MAY SAY   [what didn't work; who'd suit some buyers better]
DISCLOSURE    [the exact words, and where they go]
THE SERIES    [pieces, over how long; when it's reviewed]
FEE           [amount, paid whatever they say]
WE'LL MEASURE [what you'll look at; nothing asked of them]

```

3. The disclosure checklist. For every paid piece, before it goes out. The rules differ by country: in Australia, the ACCC reviewed 118 influencers in 2023 and found 81% made posts that raised concerns under consumer law ([ACCC](#)). Check yours.

- ☐ The disclosure is in the first line, or the first seconds, not after "see more" or in a hashtag at the end.
- ☐ It uses plain words: "Paid partnership with [brand]", not "#sp" or "#collab".
- ☐ It's on every piece in the series, including reposts and cut-downs.
- ☐ Any gift, free service or access counts as payment if it was arranged.
- ☐ On podcasts, the host says it aloud at the start, and it's in the show notes.
- ☐ Every fact the creator states about you matches your website and Credentials.
- ☐ Any client named has agreed in writing to that piece.
- ☐ You've kept a copy of each piece and its disclosure.

In the campaign plan

Creators work slowly and on their own schedule, so they're timed loosely to campaigns, never tied to a launch date.

Campaign	What creators do	Moments	What other builds bring
Owners' round tables (spring)	"Floor Time" episode 1, with the round table's speaker, released a few weeks before	Owner options	Events: the dinner; Outbound: the episode as a follow-up to named owners
The annual benchmark (Q2)	"Floor Time" episode 2, on the finding, with a client COO	Owner and COO trigger	Content: the benchmark; PR: coverage
Any review whose client agrees	Ray's floor day, and the series that follows	Plant manager framing, trigger	Customers: the client's agreement; Social: the one-question videos Ray may share
"Before you buy another line" (Q3)	One of Ray's series posts on a new line that wouldn't reach rate, if he's seen one	Plant manager trigger	Advertising: extra weight; Website: the page

Running it

Who owns what. At Kestrel, the marketing lead owns the creator list, the fit scores, the briefs and the disclosure checklist. The head of delivery chooses which review a creator may visit, and asks the client. The lead engineer on that review hosts the day. Credentials checks any figure a creator asks for. Nobody at Kestrel edits a creator's words.

When	What happens
Before any approach	50 followers or commenters sampled; the fit scorecard; 7 or more to proceed
Before the floor day	The client's written agreement; the brief; a call with the creator to answer questions
On the day	The creator walks the floor freely; the host answers anything
Before each piece goes out	The disclosure checklist, done with the creator

When	What happens
After three pieces	Review: who commented, who mentioned it to Kestrel, whether to continue
Yearly	The voices in the shape re-checked; new practitioners sampled

Paying whatever they say. The fee is agreed before the creator sees anything, and paid whether they praise Kestrel or not. A creator who writes something critical and gets paid anyway is the best advertisement this build can buy.

When a creator gets a fact wrong. Ask them to correct it, with the source. Never ask them to soften a judgement.

Common mistakes

The mistake	The stage it breaks	The bias it ignores	At Kestrel	Instead
Choosing by follower count	Known	We follow others	85,000 followers, 4 of 50 in Kestrel's choosers	Sample 50; choose on fit first, size second
Writing the script	Believed	Feeling first, reasons after	The agency wrote the paid post	Give facts and access; leave the words to them
Disclosure at the end, or in a hashtag	Believed	Losses loom larger	—	First line, plain words, every piece
One post	Known	Familiar feels safe	One post, then nothing	A series over six months, reviewed after three
A subject nobody's deciding	Answered	Good enough beats best	The trade-show stand	The moment's question, in the chooser's words
Penalising criticism	Believed	Losses loom larger	—	Pay whatever they say; thank them for the criticism

The mistake	The stage it breaks	The bias it ignores	At Kestrel	Instead
Facts that don't match the site	Answered, Bought	A hard choice stalls	—	The brief carries the facts; the site says the same
Paying clients to post	Believed	We follow others	—	Clients post unpaid, or not at all

The bias to watch in the team: **the halo effect**. One impressive trait colours the whole judgement. A large following looks like influence, so teams assume the person must be trusted, relevant and worth the fee. Fifty commenters, read one by one, usually say otherwise.

The business check and the measures

Revenue: the plant manager's fear loses deals quietly. If Ray's post settles it for even two plant managers a year, before or during a deal, it has probably paid for itself: one saved review is worth more than the year's A\$12,000.

One truth: a creator's facts come from Credentials; their judgements are their own. The two are never mixed up.

The proof is real: creators see real work on a real floor, with the client's agreement. Nothing is staged.

Capacity: about four days of Kestrel's time a year, mostly the lead engineer's floor day and preparing podcast guests.

On the measurement ladder: reach, followers and likes stay on the diagnostic rung. Judge each creator on these:

Measure	What it tells you	Stage
Audience fit, from a sample of 50, before and after	Whether the right people are listening	Known
Share of comments that ask the buying question	Whether the post moved anyone toward a choice	Preferred
New clients, and plant managers in deals, who name the creator when asked how they first heard of Kestrel, or what reassured them	Whether the belief transferred	Believed

Measure	What it tells you	Stage
Searches for Kestrel's name in the weeks after each piece	Whether it sent people to check	Known, Found
Share of pieces whose facts matched Credentials exactly	Whether the facts held	Answered

The Loop for this build

After every three pieces, and each half-year, the Loop (step 11) asks: **did the creator's audience come to believe something about Kestrel it wouldn't have believed from Kestrel?**

What the Loop finds	What it means	What changes
Comments ask "who did you use?" and plant managers mention the post in deals	The belief is transferring	Continue the series; consider a second practitioner in another sector
Lots of likes, no buying questions, nobody mentions it	The audience enjoys the creator but isn't choosing	Re-sample the audience; stop if fit has slipped
The creator's criticism is repeated by prospects	The criticism is believed, and so is the praise	Fix what they criticised, and tell them you did
Podcast listens high, owners don't mention it	The episode reached the wrong owners, or the wrong moment	Pick a guest closer to Kestrel's owners: size, sector
A fact in a piece was wrong	The brief or the check failed	Correct it with the creator; tighten the brief

What Kestrel's first year would look for. Ray's floor day and three series posts published, each disclosed in the first line. At least three plant managers in deals mentioning Ray's post. Both "Floor Time" episodes released, and at least two owners at the round tables saying they'd heard one. And the plant manager's framing moment above zero at the next rescore, because something a plant manager trusts now names Kestrel.

Shapes by business type

Business type	Typical mode	The creators that usually matter most
B2B professional services	Committee or owner-led	A few practitioners with small, exact audiences; podcasts the buyers listen to; floor days, not scripts
B2B software	Committee	Practitioners who use the product in their own work; category podcasts and newsletters; analysts belong to PR
Considered consumer (finance, health, education)	Considered personal	Trusted specialists who explain, with strict disclosure and claims rules; never testimonials dressed as advice
Ecommerce	Quick choice or habitual	Many smaller creators whose audiences match; product in use, in their words; affiliate codes tested before scaling (Partners' nod)
Local services	Owner-led or considered personal	Local voices: a café owner, a community page, a local reviewer; paid and disclosed, or not at all
Marketplaces	Varies on each side	Creators who are themselves sellers or providers, showing how it works for them

Feeling a bit overwhelmed?

You should be, a little. One moment, and still: the category's voices counted, 50 commenters read for each, a fit scorecard, a client persuaded to let an outsider walk their floor, a brief with facts but no words, disclosure in the first line of every piece, a fee paid whatever they say, and a review after three posts. Most firms pay the biggest account they can find to post what the agency wrote.

RocMode does part of the heavy lifting:

The moment, and the question to answer. Your choosers and their moments, mapped from what's known about your brand, so you know which fear a creator needs to meet, and in whose words it's asked.

Choosing creators, briefing them and checking disclosure is still work for people. RocMode doesn't do it.

[Join the RocMode beta](#)

The Outbound and Account-Based Playbook

What this build is for

Outbound reaches a short list of named firms before they start looking, so you're the firm they remember when they do.

An owner whose margins slipped again this quarter rarely searches for help. They notice it, worry about it, and talk to their accountant. If Kestrel is going to be known at that moment, someone at Kestrel has to go to them first. That's the job of this build, and it's the only build that chooses its audience by name.

Two findings pull against a cold pitch. Gartner surveyed 646 B2B buyers in August and September 2025, and 67% said they'd prefer to buy without a sales rep ([Gartner](#)). And John Dawes of the Ehrenberg-Bass Institute estimates that only about 5% of business buyers are in the market for a category at any one time ([Ehrenberg-Bass](#)). So most cold messages reach people who aren't buying, and who'd rather not be sold to when they are.

Both findings point the same way. Outbound works when it **picks the 5% by their signals and sends them something worth keeping**, rather than a request for a meeting. It fails when it sends a sales pitch to the other 95%.

This build is in the Direct category: a person, reaching out. "Account-based" means the same work done for one named firm across several builds at once: the letter, the ads that firm's people see, the round table its owner is invited to. Its rule is the same as every build's: **a message exists because a named firm is likely to be in a moment**. A list bought because it was cheap isn't a reason.

Account management, for firms that are already clients, isn't here. It belongs to Customers.

The footprint, and what it reads from the chain

Outbound leads **1 cell and 7 moments**: the owner's trigger, once for each of Kestrel's seven problems. The owner's question is the same each time, "Margins are slipping. Is it the plant?", but what set it off differs:

Moment	What set it off	How the owner says it
Owner · trigger · flat output	A key customer complains about late deliveries	"We're running flat out and still shipping late."

Moment	What set it off	How the owner says it
Owner · trigger · rising costs	A bad quarter; labour, energy or material costs jump	"Margins fell again and we can't see why."
Owner · trigger · cash in stock	The bank asks about working capital	"The warehouse is full and we're still short of the right parts."
Owner · trigger · scrap and quality	A failed customer audit or a quality claim	"We're paying to make things twice."
Owner · trigger · new line ramp-up	Months after commissioning, the line is still below rate	"We spent the money and the line still isn't at rate."
Owner · trigger · after an acquisition	The owner bought a business, or sold a stake to an investor	"We need to show margin improvement within 18 months."
Owner · trigger · automation not paying back	New machines haven't moved the numbers	"We automated, and nothing got better."

It supports **3 more cells (9 moments)**:

- **COO · trigger, for each of the seven problems.** Advertising leads it. Outbound adds the COO as the second named person at each firm on the list.
- **Investor's operating partner · trigger** ("The value-creation plan needs margin within 18 months."). Partners and channels leads it. Outbound writes to the operating partners of funds that buy in Kestrel's sectors.
- **Investor's operating partner · comparing** ("Have they done this in portfolio companies like ours?"). The pitch leads it, inside a deal. Outbound's portfolio pack is often what's on the partner's desk when they compare.

The owner is why this build exists. Owners decide in owner-led firms, rarely search, don't read guides, and ask their peers. Events reaches them in a room; Partners reaches them through their accountant. Outbound is the only build that reaches them at their own desk, by name.

What it reads, and never rewrites:

From	It takes
GTM strategy	The segments: sector, size and ownership, which become the account list

From	It takes
Brand strategy	The owner, the COO and the operating partner as choosers, with their jobs and fears
Credentials	Every claim a letter may make, and every client who may be named
Content	The benchmark and the problem guides, which become the enclosures
Partners and channels	Who can introduce Kestrel to a firm, so a warm route is tried before a cold one
Customers	Past clients who've moved to new firms: the warmest outbound there is
The pitch	What happens the moment someone replies
Email and Automation	Sending, suppression, and the opt-out list everyone shares

Its slice of the Score, and its own audit

Kestrel today	Known	Found	Quality	Moment score
Owner · trigger, each of 7 problems (out in the market)	12	0	9	0
COO · trigger, each of 7 problems (out in the market)	14	0	12	0
Operating partner · trigger (out in the market)	6	5	27	1
Operating partner · comparing (inside a deal)	Open	Open	32	32

Fifteen moments touch this build, and fourteen score zero. At the trigger, nothing Kestrel does puts it in front of an owner or a COO: none of the searches or places the Score checked named Kestrel, so Found is 0 and the gate is shut. Quality is low too. Nothing Kestrel publishes answers "Is it the plant?" in the owner's terms (Answered 5), and the next step is a contact form (Bought 15). Believed and Preferred weren't measured for these moments, so they're left out of the average.

One limit to be clear about: **Known is measured across the whole market.** A letter to 200 firms out of 1,400 moves it slowly. So this build is judged first by what happens on its named list, and the Score shows the effect later, as those firms start to find Kestrel by other routes.

The audit. What Kestrel's outbound did last year:

Activity	What Kestrel did	What it produced
A bought list, with an agency's sequence	2,400 contacts, four emails each: 9,600 sent	22 replies, 15 of them "remove me"; 4 meetings; 0 reviews
LinkedIn requests, then a pitch	300 requests from three partners	2 meetings; 0 reviews
Letters to past clients who'd moved firms	11 letters, each from the partner who'd worked with them	6 replies; 3 meetings; 1 review
Investors' operating partners	Nothing	—
An account list	None; the agency chose who to send to	—

Eleven letters produced more than 9,600 emails. The difference wasn't the channel. Each letter went to a named person, for a reason, from someone they knew.

Step 1: the shape, the account universe

The shape is **every firm in your segments, counted**, before anyone decides who to write to. Kestrel's segments come from its GTM strategy: manufacturers with 80 to 400 staff, in four sectors, in the region it can reach in a day.

Sector	Firms	Owner-led	Investor-backed	Part of a group	Clients, now or past
Metal and engineering	640	400	80	160	14
Food and drink	410	220	70	120	9
Plastics and rubber	230	130	35	65	5
Packaging	120	80	25	15	3
All four	1,400	830	210	360	31

The 210 investor-backed firms belong to about 60 funds. 25 of those funds own two or more manufacturers in these sectors, and their operating partners are a short list of their own. Group-owned firms decide at group level, often abroad, so they stay in the shape but rank lower.

Now apply the 95:5 rule. About 5% of 1,400 firms, **around 70, are in the market at any one time**. Nobody knows which 70. The job of the next step is to find signals that make a good guess.

The shape splits into three tiers, by how much work each firm gets:

Tier	Firms	What it means	Who does the work
One-to-one	40 at a time	A firm with a signal from the last 90 days. A named owner and COO, an account brief, a letter from a partner	Outbound, with Advertising and Events
One-to-few	About 160	A firm in a segment with no signal yet. The benchmark's sector findings once a year, and an invitation to the round table	Outbound, Content and Events
One-to-many	The other 1,200 or so	Everyone else in the shape	Not this build: Website, Content, Advertising and PR

Firms move up when a signal appears, and back down when 90 days pass without a reply. **Forty is a ceiling, not a target.** It's the number three partners can write to well, and follow up personally.

Step 2: coverage, one signal per moment

The coverage unit is a **signal**: something you can see from outside that suggests a named firm is in a moment now. For each moment in the footprint, ask what would show it, where you'd see it, and how much you'd trust it.

Moment	The signal	Where you see it	How far to trust it
Flat output	Job ads for a third shift, or for production planners	Job boards, the firm's careers page	Medium: growth looks the same

Moment	The signal	Where you see it	How far to trust it
Rising costs	Filed accounts showing gross margin down two years running	The company register	High, but a year late
Cash in stock	Filed accounts showing stock days rising; a new bank facility registered	The company register	Medium
Scrap and quality	Job ads for quality engineers or inspectors; a product recall	Job boards; the food safety and product recall notices	Medium; high for a recall
New line ramp-up	A new line or site announced 6 to 18 months ago; a grant awarded for equipment	Trade press, the firm's news page, grant lists	High
After an acquisition	A deal announced: the firm bought another, or an investor bought in	Trade press, deal databases, the company register	High, for about 100 days
Automation not paying back	Automation announced 12 to 24 months ago, and now job ads for automation engineers	Trade press, job boards	Low: confirm before writing
COO · trigger	Any of the above, at a firm with a COO or operations director	The firm's site, LinkedIn	The COO is the second name on the account
Operating partner · trigger	A fund announces a new manufacturing business in its portfolio	Fund news pages, deal databases	High
Operating partner · comparing	A portfolio company asks Kestrel for a proposal	Inside the deal	Certain

Three checks before a signal puts a firm in the one-to-one tier:

It's recent. Within the last 90 days, or 100 days after a deal. Older signals mean the firm has probably acted already.

It's specific. "Hiring" isn't a signal. "Hiring a second production planner, eight weeks after a delivery complaint in the trade press" is.

Someone could explain it to the owner. If a partner couldn't say in one sentence why they're writing, the signal isn't strong enough.

When Kestrel first ran this across the 1,400 firms, 52 showed a signal that passed all three checks: 11 deals, 9 new lines, 14 margin falls, 8 stock rises, 6 recalls or quality hires and 4 third-shift hires. That's close to the 70 the 95:5 rule predicts. The top 40, ranked by how strong and recent the signal was, became the one-to-one tier.

Step 3: gaps, biases and jobs

Moment	Stage	Bias	The job
Owner · after an acquisition	Known	Losses loom larger	Within 30 days of the deal, a letter from a partner with "The first 100 days on the floor": what usually slips after a deal, and how to check
Owner · new line ramp-up	Known	Losses loom larger	A letter with the ramp-up guide's one-page checklist: the five reasons new lines stall, and which one to check first
Owner · rising costs	Known, Answered	A hard choice stalls	A letter with the benchmark's finding for their sector and size: where margin goes in plants like theirs
Owner · cash in stock	Known, Answered	A hard choice stalls	A letter with one page on stock days: their sector's median, and three causes that sit on the floor, not in the warehouse
Owner · flat output	Known	We follow others	A letter naming a client in the same sector (with permission) who found capacity without a new line
Owner · scrap and quality	Known	Familiar feels safe	A letter with the root-cause guide; for a recall, wait 30 days and say nothing about the recall itself

Moment	Stage	Bias	The job
Owner · automation not paying back	Believed	Feeling first, reasons after	Only after a conversation confirms it. Nobody wants a stranger saying their investment failed
COO · trigger	Known	Familiar feels safe	The same account, a second letter, a different page: the problem guide, written for the person who runs the plant
Operating partner · trigger	Known	We follow others	A portfolio pack to the 25 funds with two or more manufacturers: results across clients, the first six weeks, reporting a board can use
Operating partner · comparing	Believed	We follow others	Cases from investor-backed clients, with the operating partner's own words where they agree
After any reply	Bought	A hard choice stalls	The pitch takes over within one working day, with the smallest next step: a call to check one number together

One pattern: **every job starts with a reason the firm would recognise and a page it would keep**. Nothing goes to a firm without a signal from the last 90 days, and nothing asks for a meeting. The letter's only request is a reply, if the reader wants one.

Step 4: reconcile, the outbound plan

Motion	Format	Volume	Moments	Signals	Priority
Warm routes first	Before any cold letter, check Partners (accountants, associations, funds) and Customers (past clients at the firm) for an introduction	Every account in the one-to-one tier	All owner triggers	Strategy, coverage	P1

Motion	Format	Volume	Moments	Signals	Priority
Signal letters	A letter from a partner, by post and email, with one enclosed page; a second letter to the COO	40 accounts at a time; about 10 new a month	Owner trigger, all seven; COO trigger	Coverage, canon	P1
Past clients who've moved	A personal note from the partner who worked with them, at their new firm	As they move: about one a month	Owner and COO trigger	Strategy	P1
The account plan	For each one-to-one account: Advertising shows the problem page to its people; Events invites the owner to the round table	The same 40	Owner and COO trigger	Coverage	P2
Operating partners	The portfolio pack, from the managing partner, once a year and after each new deal	25 funds	Operating partner trigger and comparing	Coverage, strategy	P2 (Partners leads)
Sector findings	The benchmark's findings for their sector, once a year, with a short note	About 160 one-to-few firms	Owner trigger (rising costs, cash in stock)	Canon	P3

Rejected, with reasons:

The bought list and the agency sequence. 9,600 emails, four meetings, no reviews, and 15 people who now associate Kestrel with spam. The money pays for a researcher's time to find signals.

Connection requests followed by a pitch. It reads as a trick, because it is one. Partners can connect after a conversation, not before.

Cold calls to owners. No moment asks for one, and owners screen their calls. A reply to a letter is a better start than an interrupted afternoon.

Gifts in the post. They raise the question "what do they want from me?" before the reader knows who Kestrel is. In firms with a procurement team, a gift can be a problem in itself.

Bought intent data. Kestrel's category is too small for third-party intent signals to be reliable, and the public signals above are free. Review it if the shape grows past a few thousand firms.

Sending to all 1,400. About 95% aren't in the market. The other builds reach them, at lower cost, without spending the firm's name.

The anatomy library

Outbound asset	Parts, in order	Length
The account brief	The signal, with its date and source · The moment it suggests, and the owner's likely question · The named owner and COO · Any warm route (partner, accountant, past client) · What to enclose · Who signs	One page, before anything is sent
The letter	Why you're writing, in one sentence, naming the signal · What plants like theirs usually find, with one number · What's enclosed, and why it's worth keeping · An easy way to reply, and an easy way to say no	Under 200 words
The enclosure	The finding, in the owner's terms · Three things to check, this month, without anyone's help · What the numbers looked like at a client in the same sector · Where it comes from	One page
The COO letter	The same signal, in the COO's terms · The problem guide's first page · A line saying the owner has had a letter too	Under 150 words
The follow-up	Something new: a finding, a client story, an invitation · One line linking it to the first letter · The easy no, again	Under 100 words, three to four weeks later
The handover	What they replied, word for word · The signal and the brief · The smallest next step offered	Same day, to whoever runs the pitch
The close-out	Nothing sent. The account goes back to one-to-few, with a note of what was tried	After two letters and no reply

Four rules hold for every message:

A reason they'd recognise. Every letter names the signal it's responding to. If the reason would surprise or unsettle the owner ("we noticed your recall"), it isn't used.

Something to keep, nothing to ask. The enclosure is useful whether or not they ever speak to Kestrel. The only request is a reply.

Two letters, then stop. A third letter to someone who hasn't replied tells them Kestrel doesn't listen. The firm comes back only with a new signal.

Within the rules. The law on business emails and post differs by country. Write to a named person in their business role, say where you found their details, and make opting out one click or one line. Check your own country's rules before the first send, not after a complaint.

Briefs

One brief per account in the one-to-one tier, written before anything is sent. This is the brief behind the finished example below.

ACCOUNT owner-led	Harrowfield Precision: metal and engineering, 190 staff,
SIGNAL	Bought Kellow Machining (60 staff) on 14 January. Source: trade press, 16 January; company register, 21 January
MOMENT	Owner · trigger · after an acquisition "We need to show margin improvement within 18 months."
NAMED	Richard Hale, managing director and owner Nadia Shah, operations director (second letter)
WARM ROUTE	Checked: their accountant isn't one of Kestrel's partners; no past client works there. Cold, by letter.
ENCLOSE	"The first 100 days on the floor" (one page)
CLIENT STORY	A 150-person machining business, after buying a competitor (unnamed; figures confirmed by the client)
SIGNED BY	Jo Brady, partner (covers metal and engineering)
FOLLOW-UP	Three to four weeks later: the March owners' round table
IF THEY REPLY	The pitch, within one working day: a call to look at one number together
STOP	After the second letter, if no reply

And the brief for the operating partners' pack, which runs once a year:

AUDIENCE	Operating partners at the 25 funds owning two or more manufacturers in Kestrel's four sectors
MOMENTS	Operating partner · trigger: "The value-creation plan needs margin within 18 months." Operating partner · comparing: "Have they done this in portfolio companies like ours?"
WARM ROUTE	Partners and channels first: 6 of the 25 know a Kestrel partner
THE PACK	Results across investor-backed clients (signed by Credentials); the first six weeks; a sample of the board report
SIGNED BY	The managing partner
WHEN	Once a year, and within 30 days of each new manufacturing deal
MEASURE	Introductions to portfolio companies within 6 months

The finished example: one account, written in full

The letter, its enclosure, the COO's letter, the follow-up, and the handover when the owner replied. Harrowfield Precision, Kellow Machining, the people and Kestrel are invented.

1. The letter, to Richard Hale, managing director and owner of Harrowfield Precision, by post and by email, 12 days after the deal.

Dear Richard,

Congratulations on bringing Kellow Machining into Harrowfield. I read about it in the trade press in January.

I'm writing because the months after a purchase are when owners most often ask us for help, and it's nearly always about the same thing: the two plants' numbers don't add up the way the plan said they would. In the six engineering businesses we've worked with after an acquisition, first-year margin came in an average of four points below the plan, and most of the gap was on the floor, not in the overheads.

I've enclosed one page we wrote for owners in your position: what usually slips in the first 100 days, and three things you can check this month without anyone's help. It's yours to keep, whether or not we ever speak.

If any of it is useful and you'd like to compare notes, reply to this letter or email me at jo.brady@kestrel.example. If not, I'll write once more in a month or so, and then leave you in peace.

Jo Brady, Partner, Kestrel Advisory

2. The enclosure: *The first 100 days on the floor.*

The finding. After a purchase, the plan assumes both plants will run at the better plant's level. In the first year, they usually run nearer the weaker one's.

Three things to check this month:

1. **One number, both plants.** Output per labour hour on your most similar product. If the gap is more than 15%, the plan's savings are probably in the wrong place.
2. **Both floors at shift change, in the same week.** Time how long it takes to make the first good part. That's capacity your accounts won't show.
3. **One question to both planning teams:** "What stops us shipping on time?" If the answers differ, the plants run on different rules, and a merged schedule inherits the worst of both.

At a client. A 150-person machining business bought a competitor. Output per labour hour differed by 22% between the two plants. A year later the gap was 6%, and margin was three points up.

Where this comes from. Kestrel's reviews in six engineering businesses after an acquisition, with each client's figures confirmed by the client.

3. The COO's letter, to Nadia Shah, operations director, the same week.

Dear Nadia,

After an acquisition, the operations director usually inherits the hardest part of the plan: two plants, one set of targets, and a board expecting both to move at once.

I've enclosed the first page of our guide to running two plants as one. It's written for the person who runs the floor, not the board. I've written to Richard as well, with a shorter page for owners.

If anything in it is useful, reply and I'll send the rest. If not, no need to reply.

Jo Brady, Partner, Kestrel Advisory

4. The follow-up, to Richard, four weeks later. It brings something new, rather than "just checking in".

Dear Richard,

One more thing that may be useful, and then I'll leave you in peace.

On 12 March the Metal and Engineering Association is hosting a dinner for 14 owners. Helen Dorsey, who runs a 230-person fabricator, will talk about when an owner needs outside help, and when not. Nothing will be sold. If you'd like a place, reply and I'll ask the association to send you an invitation.

Either way, I wish you and the Kellow team well.

Jo

5. Richard's reply, and the handover, the same day.

Richard's reply: "Jo, I did your first check. 19% gap. Not what I was told when we bought it. Yes to the dinner. Can we talk before then?"

TO	Mark Ellison, head of delivery (runs the pitch for this account)
ACCOUNT	Harrowfield Precision; Richard Hale, owner
WHAT HE SAID	"I did your first check. 19% gap. Not what I was told when we bought it. Yes to the dinner. Can we talk before then?"
SIGNAL	Bought Kellow Machining, 14 January
MOMENT	Now: owner · options, "Do we need outside help at all?"
SENT SO FAR	Letter and enclosure; COO letter (no reply from Nadia); round-table follow-up
NEXT STEP	A 30-minute call to go through his 19% together. No proposal.
DINNER	Place requested from the association

What to notice. The letter names the reason for writing in its first line, and the reason is good news, not a weakness. The enclosure is useful to Richard even if he never replies, and it gave him a check he could do himself; the reply came because the check found something. The follow-up offered a room full of other owners rather than a meeting with Kestrel. And the reply moved Richard from the trigger to the options moment, so the next step comes from the pitch, not from Outbound. One account; about 40 at a time.

Templates

The account list (one row per firm in the one-to-one tier)

Firm	Sector, size, ownership	Signal	Signal date	Moment	Owner	COO	Warm route	Partner	Letter 1 sent	Letter 2 sent	Reply	Status

The account brief

ACCOUNT	[firm: sector, size, ownership]
SIGNAL	[what, when, and where you saw it]
MOMENT	[chooser · step · problem, and their likely question]
NAMED	[the owner; the COO or operations director]
WARM ROUTE	[partner, accountant, association, past client; or "cold"]
ENCLOSE	[the one page, and why it fits this firm]
CLIENT STORY	[a client like them, and whether they may be named]
SIGNED BY	[the partner for this sector]
FOLLOW-UP	[the new thing the second letter brings]
IF THEY REPLY	[who in the pitch, and the smallest next step]
STOP	[after the second letter, if no reply]

The signal log (one row per signal found, whether or not it's used)

Date found	Firm	Signal	Source	Recent?	Specific?	Explainable?	Tier

The signal log keeps the tier accountable. It shows how many signals were found, how many passed the three checks, and how many became letters.

In the campaign plan

Most builds follow the campaign calendar. Outbound mostly follows the signals, which arrive when they arrive. What the Campaign Engine gives it is the **enclosures**: each campaign produces a page worth posting to an owner, and Outbound takes it to the named firms whose signal fits.

Campaign	What Outbound takes from it	Which accounts	What other builds bring
The annual benchmark (Q2)	The sector findings, as a one-page enclosure	One-to-few: about 160 firms, once a year	Content: the benchmark; PR: trade coverage the letter can mention

Campaign	What Outbound takes from it	Which accounts	What other builds bring
"Before you buy another line" (Q3, budget season)	The capacity check, as an enclosure	One-to-one firms with a flat-output or ramp-up signal	Advertising: the same message to those firms' people; the pitch: the kit
Board season	The page for the board: what a review returns, in finance terms	One-to-one firms with a margin or stock signal; the COO letter	Website: the finance director's pages; Email: the newsletter
Owners' round tables (spring and autumn)	The invitation, as the follow-up letter	One-to-one owners in the dinner's sector	Events: the room; Partners: the association co-hosting
Each new manufacturing deal by a fund	The portfolio pack	The fund's operating partner	Partners and channels: the introduction, if there is one

The account plan joins the builds up for each of the 40 firms. While Jo is writing to Richard, Advertising shows the "after an acquisition" page to people at Harrowfield, Events holds a place at the March dinner, and the Website's acquisition page is the one the letter links to. **One firm sees one story from three directions in the same month.**

Running it

Who owns what. At Kestrel, a researcher (two days a week) owns the signal log and the account list. Three partners, one per sector group, own the letters for their accounts and sign them personally. The marketing lead owns the enclosures, with Content, and the account plans, with Advertising and Events. The managing partner owns the operating partners' pack. Whoever runs the pitch owns every reply from the day it arrives.

When	What happens
Weekly	The researcher checks the signal sources: trade press, deal news, the company register, job boards. New signals go in the log, and are checked against the three tests

When	What happens
Fortnightly	A 30-minute meeting of the three partners and the researcher: new accounts in, briefs agreed, letters signed, replies passed to the pitch
Within one working day of a reply	The handover to the pitch, with the reply word for word
Three to four weeks after letter 1	Letter 2, bringing something new
After letter 2, with no reply	The close-out: back to one-to-few, a note of what was tried
Quarterly	The account list against the measures: which signals produced replies, which enclosures were kept, which partners' letters worked
Yearly	The operating partners' pack refreshed; the one-to-few sector findings sent; the shape recounted

Partners make the letters their own. The researcher drafts the brief and the marketing lead drafts the letter. The partner rewrites it until it sounds like them, then signs it. An owner can tell a letter a partner wrote from one they were handed to sign.

One opt-out list for every build. Anyone who says no to Outbound is also off Email's lists and Advertising's account audiences. Automation keeps the list, so nobody has to remember.

Common mistakes

The mistake	The stage it breaks	The bias it ignores	At Kestrel	Instead
Buying a list because it's big	Known	Good enough beats best	2,400 contacts, chosen by an agency	Count the shape, then pick firms by signal
Asking for a meeting in the first message	Believed	Losses loom larger	"Do you have 15 minutes next week?"	Send something worth keeping; ask only for a reply

The mistake	The stage it breaks	The bias it ignores	At Kestrel	Instead
Writing about Kestrel, not their problem	Answered	Feeling first, reasons after	"We're a leading operations consultancy..."	Open with their signal; enclose a page about their problem
A reason that unsettles	Believed	Losses loom larger	—	Never write "we noticed your recall" or "your margins are falling". Use the good-news signal, or wait
A sequence that won't stop	Preferred	A hard choice stalls	Four emails, then a "breaking up" email	Two letters, then stop until there's a new signal
Cold when warm was available	Believed	We follow others	Writing to a firm whose accountant knows a partner	Check Partners and Customers first, every time
A partner's name on a letter they didn't write	Believed	Familiar feels safe	Agency copy, a partner's signature	The partner rewrites it until it sounds like them
A slow handover	Bought	A hard choice stalls	A reply sat in an inbox for nine days	The pitch has it within one working day
Counting sends	All, unseen	Good enough beats best	"9,600 emails out this quarter"	Count replies from the named list, and what those firms did next

The bias to watch in the team: **base-rate neglect**. About 95% of firms aren't in the market at any one time. Teams forget it in both directions: they read a big list as a big opportunity, and a quiet month as a bad letter. The base rate says to send fewer letters, to firms with a reason, and to judge them over a year.

The business check and the measures

Revenue: owner-led firms are 830 of the 1,400 in the shape, and their owners are the chooser least likely to find Kestrel on their own. A review won from the named list pays for a year of the researcher's time.

One truth: every number in a letter or an enclosure is one Credentials can sign, and every client story is confirmed by the client. Letters never guess at a firm's problems from outside.

The proof is real: the enclosures are built from Kestrel's own reviews. If a sector has no confirmed result yet, its letters use the benchmark instead of a client story.

Capacity: about 10 new accounts a month means around 30 letters, including COO letters and follow-ups. At about 45 minutes each, with the partner's rewrite, that's roughly 22 hours a month across three partners, plus the researcher's two days a week. That's less than the agency and the list cost last year.

Measure	What it tells you	Stage
Replies from the one-to-one tier, per 100 accounts	Whether the signal and the letter reached someone in a moment	Known
Replies that mention the enclosure	Whether what was sent was worth keeping	Answered
Named-list firms that later find Kestrel by another route: a search, a site visit, a referral, a round table	Whether the letter made Kestrel the firm they remembered	Known, Found
Replies handed to the pitch within one working day	Whether the team does the part that matters	Bought
Fit calls and reviews from the named list, per year	Whether the build earns its place	All
Opt-outs and complaints	Whether the letters read as welcome	Believed
Signals that passed the three tests, by type, against replies	Which signals to trust next year	All

The Loop for this build

Each quarter, the Loop (step 11) asks: **did the firms we wrote to move from the trigger to a conversation, and which signals told us the truth?**

What the Loop finds	What it means	What changes
One signal type produces most replies	It's the most reliable sign of a live moment	Look for it first; give it more of the 40 places
Replies, but none mention the enclosure	The letter worked; the page didn't	Rewrite the enclosure with the replies' own words
No replies, but named firms turn up later at the round table or on the site	The letter made Kestrel known, at its own pace	Keep going; count these as wins
One partner's letters get replies, another's don't	It's the voice, not the list	The partners compare letters; the quieter one rewrites more
Replies stall after the handover	The pitch asks too much, too soon	A smaller first step
Opt-outs rise	A reason or a tone that unsettles	Read every opt-out reply; drop the signal type behind them

What Kestrel's first year would look for. Replies from at least one in ten of the one-to-one accounts. At least four reviews won from the named list, against one last year. A third of the named firms turning up later by another route. And no complaints: the clearest sign that an owner found a letter welcome rather than intrusive.

Shapes by business type

Business type	Typical mode	What outbound usually looks like
B2B professional services	Committee or owner-led	Signal letters from partners to a short named list; past clients who've moved; operating partners at investors
B2B software	Committee	Account-based programmes across sales, ads and events for a few hundred named accounts; signals from hiring, funding and product usage

Business type	Typical mode	What outbound usually looks like
Considered consumer (finance, health, education)	Considered personal	Rarely cold; mostly follow-up to people who asked, and referrals from existing customers. Consumer rules on direct contact are stricter
Ecommerce	Quick choice or habitual	Almost never cold outbound to consumers. Outbound to trade buyers and stockists, if you sell wholesale
Local services	Owner-led or considered personal	Letters to local businesses or households after a visible trigger: a new site, a move, a planning application
Marketplaces	Varies on each side	Outbound to the supply side (sellers, providers) to seed the marketplace; rarely to buyers

Feeling a bit overwhelmed?

You should be, a little. One firm: 1,400 firms counted, a researcher checking signals every week, three tests for each, 40 account briefs, a letter and an enclosure for each, a second letter to the COO, a follow-up that brings something new, a handover within a working day, and a close-out when it's time to stop. Most firms buy a list, load a sequence, and count the emails.

RocMode does part of the heavy lifting:

The people you'd write to, and their moments. Your choosers and their moments, mapped from what's known about your brand, so every account brief starts from who's likely to be asking what, and why.

The segments your account list starts from. RocMode's GTM strategy sets out who you sell to, so the shape is counted from your segments, not from whatever list is cheapest.

What happens when someone replies. RocMode writes a demo script and an objection sheet for each buyer, and every answer cites its proof or names the proof still to collect.

[Join the RocMode beta](#)

The Pitch Playbook

What this build is for

The pitch wins the room, one person at a time, and keeps winning it after you've left.

In a committee sale, the decision isn't made in the meeting. It's made in the week after, when the champion forwards your proposal to people who never met you, and each of them reads it for their own question. The finance director reads for the numbers and the way out. The plant manager reads for what it will cost their team. Procurement reads for the paperwork. If the proposal answers only the champion, the others answer their own questions, usually with no.

Two findings make this harder every year. Gartner surveyed 646 B2B buyers in 2025, and 67% said they'd prefer to buy without a sales rep at all ([Gartner](#)). And the people who decide are more than the people you meet: Gartner puts a typical buying group at five to 16 people. So the pitch is less a performance than a kit: the assets that carry your case into rooms you're never in.

In Kestrel's map, this build leads more cells than any other: **19 cells and 28 moments**, almost all of them after the buyer has made contact. It's also where nearly all of Kestrel's current score comes from. That's the trap this playbook exists to avoid: a firm that's good in the room assumes the room is where it wins, and stops looking at the moments before anyone calls.

The pitch covers every step from the fit call to the signed contract, and the re-engagement after it. Account management and customer service are nods: they get short sections in the Customers playbook, because the same person who renews is the one who was served.

The footprint, and what it reads from the chain

The pitch leads where buyers judge Kestrel in person, or through what Kestrel put in their hands:

Step	The moments the pitch leads	Who
Scope	What the review needs from the floor, the team and the data; what fits the 100-day plan	Plant manager, improvement lead, investor

Step	The moments the pitch leads	Who
Comparing	Trust, working on the floor, working with the in-house team, portfolio experience	Owner, plant manager, improvement lead, investor
Checking	Seeing Kestrel's people on a floor like theirs	Plant manager
Building the case	The board's backing, the best use of the money, the numbers standing up (one moment per problem), team time, looking bad, the board pack	COO, owner, finance director (7), plant manager, improvement lead, investor
Buying	Start dates and what's needed, who signs, the contract and the way to stop, starting inside 100 days	COO (2), owner (2), finance director, investor (2)
Extend and refer	Re-engaging without a new tender	Procurement

It supports nine more: the COO's and finance director's comparing, the COO's and investor's checking (references, arranged by PR), the finance director's and procurement's scope, the investor's options, and the COO's and investor's follow-on.

What it reads, and never rewrites:

From	It takes
Credentials	Every claim and its state. Only signed claims go in a proposal; the way-out clause word for word
Brand Strategy	The reason to be chosen; the choosers, by name, with their jobs and fears
Where you stand	The moments above, each chooser's question in their words, and the Score inside deals
GTM Strategy	The motion, the segments, the win strategy against each alternative
Sales Strategy	The ten motion steps, the objections with their proof, the routing by buying path
Website	The proof pages the proposal links to, so the two never disagree

Its slice of the Score, and its own audit

Inside a deal the gate is open: the buyer is already talking to you. So the pitch's slice of [the Selectability Score](#) is quality: **Answered** (did the asset this person was given answer their question?), **Believed** (are the claims signed, and confirmed by clients?), **Preferred** (did Kestrel keep the deal at this step?) and **Bought** (did the step run without a stall?).

Kestrel today, the 25 deal moments the pitch leads	Answered	Believed	Preferred	Bought	Moment score
Average	39	43	52	46	45

The pitch earns 8.9 of Kestrel's 18 points, about half the Score. It also loses 12.4. The other three moments it leads, the scope questions from the plant manager, the improvement lead and the investor, score zero: they're asked before anyone has called Kestrel, and the one-page scope only answers them once a deal exists.

The audit. Before changing anything, lay out what the team actually sends, and the last year's deals.

What Kestrel sends today	Who it answers	What it misses
A 24-slide credentials deck	The COO, in the meeting	Everyone else; it can't be forwarded without someone to talk over it
One proposal template, 12 pages, the same for every client	The COO	The finance director's numbers and way out; the plant manager's team time; procurement's paperwork
Three case study PDFs	The owner, loosely	No figures confirmed by the client's finance team
References "on request"	Nobody, for two to three weeks	The checking step stalls while the reference is found
The contract	Procurement, at the end	The terms first appear after the decision, so they reopen it

The last twelve months	
Fit calls	31, of which 22 came by referral
Proposals	19
Won	9
Lost: no decision, stalled at the board	4
Lost: a large firm	3
Lost: bought equipment instead	2
Lost: the in-house team did it	1

The loss pattern points straight at the moments. Four deals died at building the case, where the finance director's numbers and the board pack were missing. Three went to a large firm at comparing, where the owner's trust and portfolio experience weren't shown. Two went to equipment because nobody priced the capacity already in the line. Not one was lost on the quality of the work.

Step 1: the shape, the motion and the standard proposal

The pitch has two shapes, and both come from outside the firm.

The motion comes from the Sales Strategy in [the Plans](#). It's set by how the category buys, never by what the sales team prefers. Kestrel's has ten steps:

#	Motion step	What the buyer needs from it
1	Fit call	Is this for a plant like ours, and roughly what does it cost?
2	One-page scope	What will it cover here, and what will it need from us?
3	Committee proposal	Something each of us can read alone, and forward
4	Reference call	Someone like us who's done it
5	Business case	Numbers the board will accept
6	Contract	Terms that match everything we've been told
7	Week zero	What to have ready before they arrive
8	The review	Weekly findings in the plant's words
9	Results and renewal	Savings confirmed; the programme or the next plant
10	Referral	A reference we're glad to give

Steps 8 to 10 belong to Customers; the pitch hands over at week zero and picks up again at re-engagement.

The standard proposal is canon: what buyers in the category expect to find, because every firm they've compared shows it. Count it the same way as website page types, across the proposals and tenders you can see.

Proposal section	In the category's proposals (of 5)	Canon?
Summary	5	Yes
Our understanding of your situation	5	Yes
Approach and timeline	5	Yes
The team	5	Yes
Price	4	Yes
Case studies	4	Yes

Proposal section	In the category's proposals (of 5)	Canon?
Terms	3	Yes
A page for each person in the buying group	0	No, and that's the opening
The way out	0	No: Kestrel's alone

The counts are invented for the example. Canon tells you what a proposal can't leave out. It never tells you what wins: no firm in Kestrel's category writes to each person in the room, and none states what happens if the work finds nothing.

Step 2: coverage, one section per fear, per person in the room

The coverage unit for the pitch is **a person's fear at a step**. Every lit cell holds a chooser's question; behind each question sits the thing they're afraid of, and the fear is what the section has to answer. A proposal that answers questions and ignores fears reads well and loses.

For each of the 37 moments the pitch leads or supports, write three things from the map and the Sales Strategy: **the question, the fear, and the signed proof that answers it**.

Chooser	Their question in the pitch	Their fear	The proof that answers it
COO	"How do I get the finance director and the board to back this?"	Championing something that fails in front of the board	A business case the board will accept, with the way out stated; a reference from a COO who did it
Owner	"Can I trust them with my business and my people?"	A smaller firm that lets them down where everyone can see	References from owners of similar firms; the names of the people who'll be on site
Finance director	"Will the numbers stand up at the board?" (one per problem)	"Consultants always overstate savings"	Savings confirmed by clients' finance directors months later; the cost worked out with their own figures; the way out

Chooser	Their question in the pitch	Their fear	The proof that answers it
Plant manager	"How much of my team's time will it take?"	"They'll tell us what we already know, and take my best people for six weeks"	The six weeks, hour by hour; a sample review report in plant language; a visit to a floor like theirs
Improvement lead	"Will this make my team look bad?"	A verdict on their work, delivered to their boss	Cases where the in-house team led the changes and kept the methods
Procurement	"Can we re-engage them without a new tender?"	A contract that can't be compared or renewed cleanly	A brief template with fair criteria; terms in writing from the start
Investor	"Can they start inside the first 100 days?"	A slow start that wastes the value-creation window	Start dates; portfolio cases; board reporting in their terms

That's seven people and 37 moments, each with a fear and a proof to write down: the finance director alone has seven, one for each problem's numbers. The standard proposal from Step 1 has a place for almost none of them. That's the gap Step 3 measures.

Step 3: gaps, biases and jobs

Hold today's kit against the 37 fears. Each fear the kit doesn't answer is a gap; each gap becomes a job with a stage, a moment and a bias.

Moment	Stage	Bias	The job
Finance director · building the case, each problem	Believed	Losses loom larger	A business case kit per problem: their own figures in, the cost of doing nothing out, the way out on the same page
Finance director · buying: "How do we stop?"	Bought	Losses loom larger	The way-out clause quoted in the proposal, word for word as in the contract

Moment	Stage	Bias	The job
COO · building the case	Believed	A hard choice stalls	A board paper the COO can send as it is, in the board's format, under two pages
Owner · comparing: trust	Believed	Familiar feels safe	The people on site, by name and plant experience, and a reference from an owner of a similar firm
Owner · building the case: best use of money	Preferred	Losses loom larger	Kestrel against new equipment, priced with their line's numbers
Plant manager · comparing and building the case	Answered	A hard choice stalls	The six weeks hour by hour: who's needed when, and for how long
Plant manager · checking	Believed	We follow others	A half-day visit to a client's floor, offered in the proposal, not on request
Improvement lead · comparing and building the case	Preferred	Losses loom larger	A page written to them: what they'll lead, what they'll keep, how the findings will credit their team
Investor · comparing, building the case, buying	Believed	We follow others	Portfolio cases; a start date inside 100 days; a board-pack page in value-creation terms
Procurement · re-engaging	Bought	A hard choice stalls	A framework option in the first contract, so the next review needs no new tender
Every chooser · checking	Believed	We follow others	A reference in each sector, ready within a week
Every chooser · after the meeting	Answered	A hard choice stalls	A proposal each person can read alone: one page each, forwardable without the pitch

One pattern runs through the table: **almost every job is about someone who wasn't in the meeting.** Kestrel's pitch was built for the COO across the table. Its losses came from the finance director at the board, the owner weighing a bigger name and the plant manager protecting their team.

Step 4: reconcile, the pitch kit

Put every job into the motion. The result is the pitch's equivalent of a sitemap: **the kit**, every asset the team can put in a buyer's hands, each traced to the moments it answers. Priority uses the same three signals as every build: canon, coverage and strategy.

Asset	Motion step	Moments it answers	Signals	Priority	Was
Fit-call guide: who it's for and isn't, the six weeks, the price range	1	4: comparing, for the COO, owner and finance director	Canon, coverage	P1	Improve (the 24-slide deck, cut to 6)
One-page scope	2	3: the plant manager's, improvement lead's and investor's scope	Coverage, strategy	P1	New
Committee proposal: a summary, then one page per person	3	6: building the case, every chooser	Canon, coverage, strategy	P1	Replace (the 12-page template)
Business case kit, one per problem	5	7: the finance director's numbers, by problem	Coverage, strategy	P1	New
Board paper	5	2: the COO's board, the investor's board pack	Coverage	P1	New
Terms summary, with the way-out clause word for word	6	3: who signs, the contract, how to stop	Canon, coverage, strategy	P1	New (was only in the contract)
Reference programme: one client per sector, within a week	4	4: checking, with PR	Coverage, strategy	P1	Replace ("on request")
The six weeks, hour by hour	2 to 3	3: the plant manager's comparing, case and team time	Coverage	P2	Improve (one slide)
A page for the improvement lead	3	3: comparing, building the case, looking bad	Coverage, strategy	P2	New
Investor pack	3 to 6	4: comparing, board pack, starting inside 100 days	Coverage, strategy	P2	New
Site visit offer	4	1: the plant manager's checking	Coverage	P2	New
Start dates and week-zero checklist	6 to 7	3: how soon, what's needed, starting this month	Coverage	P2	New; handed to Automation
Procurement pack and framework option	6, then 10	2: re-engaging without a tender; the brief	Coverage	P3	New

13 assets. Every one of the 28 lead moments has at least one; every asset answers at least one. Three come from other builds and aren't rewritten here: the case pages (Website), the cost articles (Content) and the reference calls (PR and reviews). The kit links to them.

Rejected, with reasons:

- **A capability statement.** No moment asks for it. It answers "what do you do?", which the fit call already settled.
- **The full credentials deck in every meeting.** Cut to six slides for the fit call. After that, the buyer needs answers to their own questions, not the firm's history.
- **A proposal per chooser, as separate documents.** One proposal with a page each keeps the numbers identical; seven documents drift.
- **Discounting to close.** It answers no fear the way out doesn't answer better, and it tells the finance director the price was soft.

The anatomy library

Each asset in the kit has a standard shape. Buyers learn where to look; sellers stop reinventing. Required parts are in bold.

Asset	Parts, in order	Length
Fit-call guide	Who it's for, and who it isn't · Three questions to ask them · What the review is, in two sentences · The price range and the six weeks · What happens next, and when	Six slides or one page
One-page scope	What the review will cover in this plant · What it needs from the team, by role and hours · What it won't touch · The start date	One page
Committee proposal: summary	Their situation, in their words · What changes, by when · The price and the way out, side by side · Who's on site · One decision to make	One page

Asset	Parts, in order	Length
Committee proposal: a person's page	Their question, answered in the first line · Their fear, named and answered · The proof, from someone in their role · What they'll be asked to do · Who to call	One page each
Business case kit	The cost of the problem, with their figures · What the review typically finds, sourced · The fee, and the way out · The payback, low and high · The assumptions, editable	A spreadsheet and a one-page summary
Board paper	The decision asked for · The problem and its cost · Options considered, including doing nothing and buying equipment · The recommendation, the risk and the way out	Two pages, in the board's format
Terms summary	Price, duration, what's included · The way-out clause, word for word · Payment terms · Insurance and site safety · Who signs	One page
Investor pack	Portfolio cases, with timing · Start dates inside 100 days · A board-pack page in value-creation terms · Roll-out across the portfolio	Four pages

Four rules hold for every asset:

- **It works without you in the room.** Read it as someone who never met Kestrel. If it needs talking over, it isn't finished.
- **Every figure is a signed claim.** If Credentials can't sign it, it doesn't go in front of a buyer.
- **The price never appears without the way out.**
- **One decision per asset.** The fit-call guide asks for a scope, the scope for a proposal meeting, the proposal for a signature.

Briefs: the kit brief, then a brief for every section

A pitch is written for one account, so the brief has three levels: the kit (once, for every deal), the deal (once per account) and the section (once per person's page).

Level	The brief holds	Written
Kit	The motion, the assets, their shapes, the claims each may use	Once, updated quarterly
Deal	This account's buying path, the people in the group, their problem and sector, the figures they've shared, the alternatives they're weighing	After the fit call
Section	One person, one question, one fear, one proof	For each page of the proposal

The deal brief for one of Kestrel's live accounts (invented):

ACCOUNT	Hartwell Plastics, 260 staff, plastics and rubber
BUYING PATH	Planned committee; the owner chairs the board
PROBLEM	Rising costs: energy and scrap on two injection lines
THE GROUP	COO (champion); finance director; plant manager; improvement lead;
	procurement; the owner (decides)
ALTERNATIVES	A new energy-efficient press (quoted: A\$1.4m); the improvement team
THEIR FIGURES	Scrap 5.8%; energy 11% of cost of goods; overtime 12% of hours
CLAIMS	Fixed price; six weeks; way-out clause; Calloway scrap case (signed)
WHAT WINS	The finance director accepting the business case before the board

Then each person's page gets its own brief:

SECTION	The finance director's page
JOB	Believed · finance director · building the case · losses loom larger
QUESTION	"Will the numbers stand up at the board?"
FEAR	"Consultants always overstate savings."
PROOF	Calloway Foods: scrap 4.1% to 2.6%, confirmed by its finance director
THEIR FIGURES	after three months; the way-out clause
ASK	Scrap 5.8%, energy 11% of cost of goods
GUARD	Review the business case kit with their own numbers before the board
	Every figure signed or theirs; the fee and the way out on the same page

The flow check. Read the proposal as each person, one at a time, reading only their page and the summary. Could they say yes from that alone? Then read it as the board, which sees only the summary and the board paper.

The finished example: Kestrel's committee proposal to Hartwell Plastics, written in full

This is the proposal that follows the deal brief above: a summary, then one page for each person in the group. Each page carries its job in italics. Hartwell, its people and its figures are invented; Kestrel's claims are the signed ones from Credentials.

Summary: for everyone, and for the board

Believed, Bought · the whole group · a hard choice stalls

Hartwell's costs are rising on two injection lines, and the first fix probably isn't a new press. Scrap is running at 5.8% and energy at 11% of the cost of goods. Before committing A\$1.4 million to a new press, a six-week operations review on both lines will show how much of the cost is recoverable from the equipment you have.

- **What changes:** scrap and energy losses traced to their causes on both lines, with fixes your team runs, and the savings confirmed with your finance director three months on.
- **The fee:** A\$60,000, fixed, for two lines. **If the review doesn't find more than it costs, you don't go ahead with the changes, and you don't pay for them.**
- **Who's on site:** two Kestrel engineers for six weeks, both with injection-moulding plant experience.
- **The decision:** a start date. We can start on 3 November.

For the COO

Believed · COO · a hard choice stalls

You asked how to get the finance director and the board behind this. The attached board paper is written to be sent as it is. It sets out the problem, its cost in Hartwell's figures, the options including the new press and doing nothing, and the way out. A COO at a packaging plant of a similar size will take your call this week, if you'd like to hear how the review went from someone who championed one.

For the finance director

Believed · finance director · losses loom larger

You'll want to know whether these numbers survive a board. They're your numbers, not ours. At 5.8% scrap on roughly A\$14 million of materials, scrap costs Hartwell about A\$810,000 a year. Plants that fixed scrap at the root reduced it by between a quarter and a third; Calloway Foods did better, from 4.1% to 2.6%, confirmed by its finance director three months later. For Hartwell, a quarter to a third is A\$200,000 to A\$270,000 a year, before energy. The business case kit lets you change every assumption. And the fee comes with the way out, written into the contract.

For the plant manager

Answered · plant manager · a hard choice stalls

You asked how much of your team's time it takes. About two hours a day from you, and an hour a shift from each line lead. The review runs on your floor, on every shift, alongside your people. Week by week: weeks one and two, measuring and listening; weeks three and four, testing fixes with your crews; weeks five and six, the changes your team will run, written in your plant's words. You're welcome to spend half a day at a client's plant first, to see what that looks like.

For the improvement lead

Preferred · improvement lead · losses loom larger

This review won't be a verdict on your team's work, and it won't replace it. Your team leads the changes; we bring the measurement and a second pair of eyes for six weeks. Every finding is credited to the people who make it work, and the methods stay with your team when we leave. At Brennock Packaging, the in-house lead ran the next review on a third line without us.

For procurement

Bought · procurement · a hard choice stalls

Everything you'll need is on one page, attached: the fixed fee and what it includes, payment in two stages, insurance and site safety documents, and the way-out clause word for word. The contract includes a framework option, so a review of a further line can start without a new tender.

For the owner

Believed · owner · familiar feels safe

You'll be weighing whether a firm of our size can be trusted with your plant. Ask an owner who did. The managing director of a 230-person metal fabricator will take your call; she chose Kestrel over a large firm two years ago and has used us twice since. The two people on your floor are named on the summary, with their plant histories. And if the review doesn't find more than it costs, it costs you nothing to find out.

What to notice. Each page opens with that person's own question, answered in its first line. Each names their fear before they have to raise it. Every figure is either Hartwell's or a signed claim. The price never appears without the way out. And the whole proposal can be forwarded in pieces: the finance director can send their page and the kit to the board without the rest. Seven pages, for one deal. Kestrel writes 19 proposals a year.

Templates

The deal brief

ACCOUNT	[name, size, sector]
BUYING PATH	[planned committee / owner-led / investor-led / urgent / repeat]
PROBLEM	[from the seven, in their words]
THE GROUP	[each person: role, champion / decider / blocker, what they've said]
ALTERNATIVES	[what else they're weighing, with figures if known]
THEIR FIGURES	[what they've shared; never estimate them]
CLAIMS	[signed claims this deal will use]
WHAT WINS	[the one moment that decides this deal]

The person's page brief

SECTION	[person's page]
JOB	[stage · chooser · step · bias]
QUESTION	[in their words, answered in the first line]
FEAR	[named before they raise it]
PROOF	[a signed claim, from someone in their role]
ASK	[the one thing they're asked to do]
GUARD	[passes only if they could say yes from this page and the summary]

The deal map (one row per person, kept live through the deal)

Person	Role in the decision	Their question	Their fear	Asset sent	Answered?	Last contact	Next step

The win-loss record (one row per proposal)

Account	Buying path	Won or lost	The moment it was decided	Who decided	Lost to	What the kit lacked

In the campaign plan

Most campaigns end in a conversation, so the pitch is where they pay off or leak. The rule: **every campaign hands the pitch the asset its buyers will ask for next, before the campaign starts.**

Kestrel campaign	What it sends to the pitch	The asset the pitch must have ready
"Before you buy another line"	COOs and owners weighing equipment	The business case kit with the equipment comparison; a reference who didn't buy the line
Board season	Finance directors building the case	The business case kit, dated for this budget round; the board paper
The annual benchmark	Finance directors and COOs with a figure in hand	A fit-call guide that opens with their benchmark gap
Owners' round table	Owners who've just heard from an owner	The owner's page, and a reference from the speaker

Kestrel campaign	What it sends to the pitch	The asset the pitch must have ready
Investor 100 days	Operating partners with a deadline	The investor pack with start dates

When a campaign is planned, the deal team reads its brief and confirms the asset exists. A campaign that fills the calendar with fit calls the kit can't answer has wasted the buyers it found.

Running it

Who owns what. At Kestrel, the managing partner owns the kit and the win-loss record. Each deal has one owner, usually the partner who ran the fit call, who keeps the deal map. The head of delivery signs anything that describes the work on the floor. Nobody edits a signed claim in a proposal; they ask Credentials for a new one.

When	What happens
After every fit call	The deal brief, within a day; the one-page scope within two
Weekly	The deal review: each live deal's map, person by person. Who hasn't been answered? Which fear is still open?
After every decision, won or lost	The win-loss record: the moment it was decided, by whom, and what the kit lacked. A call to the buyer for lost deals, by someone who didn't run the deal
Quarterly	The kit reviewed against the win-loss record; one asset improved, never all of them
When a claim changes	Every asset that uses it, the same week; live proposals reissued

Using AI to write. A person's page brief is tight enough for an assistant to draft from. What it can't do is know the account: the figures, the people and what was said in the room come from the deal brief, which a person writes. Every draft is checked against the claims register before it leaves.

The deal review's one question. For each person in each live deal: could they say yes today, from what we've sent them? If not, what are they missing? That question, asked weekly, does more than any pipeline report.

Common mistakes

The mistake	The stage it breaks	The bias it ignores	In Kestrel's old pitch	Instead
Pitching to the person across the table	Answered, for everyone else	A hard choice stalls	One proposal, written for the COO	A page for each person in the group
A proposal that needs you to explain it	Believed	We follow others	24 slides, made to be talked over	Assets that work forwarded, without anyone in the room
Leaving the fear unnamed	Preferred	Losses loom larger	The finance director's doubt about savings never addressed	Name each person's fear before they raise it, and answer it with proof
Price without the way out	Bought	Losses loom larger	The fee on page 11, the clause in the contract	The price and the way out, always side by side
Proof from the wrong person	Believed	We follow others	A COO's quote shown to the finance director	Proof from someone in the reader's own role
References on request	Believed	Familiar feels safe	Two to three weeks to find a reference, while the deal cooled	One reference per sector, ready within a week
Terms that appear after the decision	Bought	Losses loom larger	Procurement saw the contract last, and reopened the price	The terms summary in the proposal
Discounting to close	Believed	Feeling first, reasons after	A 10% discount offered to a stalled board	The way out answers the fear a discount only hides

The mistake	The stage it breaks	The bias it ignores	In Kestrel's old pitch	Instead
Blaming lost deals on price	Every stage, unseen	Feeling first, reasons after	"Lost on price" in the CRM for six of ten losses	Ask the buyer, after the decision, which moment decided it

The bias to watch in the team: **the illusion of transparency**. A partner who has run fifty reviews can't easily imagine not knowing what one involves, and assumes the plant manager can picture the six weeks. Write every page as if the reader has never hired a consultant, because often they haven't.

The business check and the measures

Before the new kit goes out, check it against the business:

- **Revenue:** the kit is built around planned committees, where most of Kestrel's work comes from. If most new work starts to come through investors, the investor pack moves to P1.
- **One truth:** the fee, the six weeks and the way-out clause read identically in the fit-call guide, the proposal, the terms summary, the website and the contract.
- **The offer:** the way out is in the contract, signed by the managing partner. A seller who softens it in a meeting changes the offer, and that's a partner's decision, never a seller's.
- **The proof is real:** every case figure is confirmed by the client in writing; every reference has agreed to take calls this quarter.
- **Capacity:** a seven-page proposal per deal takes a day with the kit, and three without it. If the team can't give a day per proposal, cut the page count, not the people: the finance director's and the owner's pages come first.

The measures. Each asset is judged by the moment it serves, never by how many were sent.

Measure	What it tells you	Stage
Proposals that reach a decision within 30 days	Whether the committee has what it needs	Bought
Deals lost to "no decision"	Whether building the case is answered (Kestrel: 4 of 10 losses)	Believed

Measure	What it tells you	Stage
Business case kits opened by a second person at the account	Whether the champion is forwarding	Answered
Reference calls held within a week of the request	Whether checking stalls	Believed
Days from verbal yes to signed contract	Whether the terms were clear before the decision	Bought
Win rate by buying path and by alternative	Where Kestrel wins, and against whom (today: 9 of 19)	Preferred
Losses where the improvement lead opposed the review	Whether the improvement lead's page works	Preferred

The Loop for this build

Each quarter the Loop (step 11) puts the win-loss record beside the Score's deal moments, and asks: **at which moment are deals being decided, and did the kit answer the person who decided them?**

What the Loop finds	What it means	What changes
Losses cluster at building the case	The finance director or the board isn't answered	The business case kit and board paper; check they're being forwarded
Losses to one alternative	The win strategy against it isn't in the kit	The comparison page and proof for that alternative
Wins slow down after the verbal yes	Terms are arriving late	The terms summary moves earlier in the motion
A person's page is never opened by a second reader	The champion isn't forwarding it, or it isn't worth forwarding	Rewrite it to their question, or send it to them directly
Win rate rises but fit calls don't	The pitch is working; the market moments aren't	Nothing in the kit. The work is in Website, Being named and Content

What Kestrel's first year would look for. "No decision" losses falling first, from four in ten, as the business case kit and board paper reach the board. Then fewer losses to large firms as owner references and named teams arrive at comparing. Win rate could move

from 9 in 19 towards 12, which would lift the pitch's moments by several points. But the last row of the table is the one to watch: a better pitch only converts the buyers who already find Kestrel. Without the market-side builds, the Score stays "rarely selected", however good the room gets.

Shapes by business type

The motion and the kit change with how the category buys:

Business type	Typical mode	The kit that usually matters most
B2B professional services	Committee or owner-led	Scope, a proposal with a page per person, a business case, references, terms with a way out
B2B software	Committee	Demo script per role, a trial or pilot plan, security and compliance pack, a business case, procurement pack
Considered consumer (finance, health, education)	Considered personal	A consultation guide, a plain-language summary of costs and risks, cooling-off terms, someone like them to talk to
Ecommerce	Quick choice or habitual	Rarely a pitch: live chat and the product page do the job; a sales team only for trade and bulk buyers
Local services	Owner-led or considered personal	A quote that explains itself, before-and-after examples, reviews from nearby customers, a clear guarantee
Marketplaces	Varies on each side	The seller pitch: what it costs, what it earns, how fast it starts paying, who else has joined

Feeling a bit overwhelmed?

You should be, a little. One firm: a ten-step motion, seven people in the room, 37 moments, 13 assets, a deal brief and a seven-page proposal for every account, a weekly deal review and a win-loss record that has to tell the truth about why deals die. Most firms send the same deck to everyone and call the losses price.

RocMode does the heavy part for you:

- **A script and an objection sheet for each buyer.** For every person in the buying group, RocMode writes a demo script and an objection sheet, and every answer cites its proof, or says what proof you still need to collect.
- **Battlecards and seller copy you can stand behind.** It drafts them only from what your strategy and your signed claims can back up.
- **Pitch decks, from the same strategy.** So the deck, the website and the proposal say one thing.

[Join the RocMode beta](#)

The Events Playbook

What this build is for

An event puts buyers in a room with you, and with each other, at a moment a page can't settle.

Some questions are decided by people, face to face. "Do we need outside help at all?" is one an owner asks other owners over dinner, not an assistant. "What do people I trust say about them?" is answered best by hearing those people in person. An event is the only build where Kestrel can be in that room.

Marketers rate events highly: in the Content Marketing Institute's 2025 survey of 980 B2B marketers, in-person events topped the channels said to produce the best results, at 52% ([CMI](#)). That's marketers' own opinion, not a measure of buyers. It's a reason to run events well, not to run more of them.

This build is in the Direct category: a person, live. It covers the events Kestrel hosts, the ones it attends, and the ones where it speaks. Its rule is the same as every build's: **an event exists because a moment needs a room**. A stand at a trade show because competitors have one isn't a reason.

The footprint, and what it reads from the chain

Events lead **2 cells and 2 moments**, both the owner's:

Moment	Who	Their question
Options	Owner	"Do we need outside help at all?"
Checking	Owner	"What do people I trust say about them?"

And support three more cells (12 moments): the COO's trigger for each of the seven problems (where a talk on flat output reaches COOs in the room), the plant manager's framing ("will they just tell us what we already know?", answered by seeing Kestrel's people on a floor), and the COO's checking by sector (where a client speaking at an event is the reference).

The owner is the person this build exists for. Owners decide in owner-led firms, and they trust other owners more than any firm. Of all Kestrel's choosers, they're the hardest to reach any other way: they don't read guides, rarely search, and ask their peers.

What it reads, and never rewrites:

From	It takes
Brand Strategy	The owner as a chooser, with their jobs and fears
Credentials	Every claim a speaker may make, and every client who's agreed to speak
Customers	The clients who'll speak, host or attend: the community and the results reviews
Partners and channels	The associations and accountants who fill rooms with owners
Content	The benchmark, which gives every talk a finding to present

Its slice of the Score, and its own audit

Kestrel today	Known	Found	Quality	Moment score
Owner · options (out in the market)	12	10	29	3
Owner · checking (inside a deal)	Open	Open	61	61

The two moments tell the whole story. Inside a deal, when an owner hears from people they trust, Kestrel scores 61: owners who talk to Kestrel's clients like what they hear. Out in the market, when an owner wonders whether to bring anyone in at all, Kestrel scores 3: it's named in one association directory, and nowhere owners actually talk.

The audit.

Event	What Kestrel does	What it produced last year
The annual manufacturing trade show	A stand, two partners for three days	140 badge scans; 3 fit calls; 0 reviews
Two industry conferences	Attended; no speaking	Business cards; nothing traced
Association events	Member of one association; doesn't attend	Nothing
Hosted events	None	—

Event	What Kestrel does	What it produced last year
Client events	None	—

The trade show cost the most and produced the least. The rooms where owners actually decide, small, peer-led, often hosted by an association or an accountant, Kestrel has never been in.

Step 1: the shape, the category's calendar

The shape is **the events the category already runs**, by type and by who's in the room.

Event type	Who's in the room	In Kestrel's category	Kestrel today
Trade shows	Everyone, briefly; mostly buyers of equipment	One major show a year	A stand
Industry conferences	COOs, improvement leads, consultants	Two a year	Attends
Association events	Owners, COOs, by sector	Four associations, three or four events each	Not attending
Peer round tables	Owners, 10 to 20 at a time	Run by associations, accountants and banks	Never
Plant visits	Plant managers, improvement leads	Hosted by manufacturers, sometimes by programmes	Never
Webinars	Anyone, at a distance	Common; mostly unwatched	None
Client and alumni events	Past clients	Large firms only	None

The calendar is set by the category, not by Kestrel: trade-show season, budget season, the associations' annual dinners. The question for each is the same: who is in that room, at which moment, and is Kestrel the best answer to their question there?

Step 2: coverage, one conversation per moment a room decides better than a page

The coverage unit is a **conversation**: who needs to hear whom, about what, at which moment. For each of the five cells events lead or support, ask whether a room would answer the question better than anything Kestrel could publish.

Moment	Who needs to hear	From whom	About what	Better in a room?
Owner · options	Owners	Another owner who used outside help	Whether it was worth it, and what it cost them in disruption	Yes: peers, in private
Owner · checking	Owners weighing Kestrel	Kestrel's client owners	What working with Kestrel was really like	Yes
COO · trigger (7 problems)	COOs	A COO and a Kestrel engineer, together	What flat output, rising costs and the rest usually come from	Partly: a talk reaches the room; the guide reaches everyone else
Plant manager · framing	Plant managers	Kestrel's engineers, on a floor	Whether outsiders find anything they don't already know	Yes: seeing it on a floor
COO · checking, by sector	COOs in a sector	A client COO from the same sector	Their results, and how the review ran	Yes, if the speaker is from their sector

Four conversations need a room; one needs a room and a page. That gives three event formats: **the owners' round table**, **the open plant visit**, and **a talk with a client** at events the category already runs.

Step 3: gaps, biases and jobs

Moment	Stage	Bias	The job
Owner · options	Known, Found	We follow others	Two owners' round tables a year, co-hosted with an association, with a client owner as the speaker

Moment	Stage	Bias	The job
Owner · checking	Believed	Familiar feels safe	Invite owners in live deals to the next round table, to hear from a client owner in person
COO · trigger	Known	Familiar feels safe	Replace the trade-show stand with a talk: a client COO and Kestrel's head of delivery on the benchmark's finding
Plant manager · framing	Believed	We follow others	Open plant visits, hosted by a client: half a day on a floor where a review ran, for plant managers and improvement leads
COO · checking, by sector	Believed	We follow others	At each sector association's annual event, a client COO from that sector speaking
After every event	Bought	A hard choice stalls	A personal follow-up within 48 hours, from the person who met them, with the one thing they asked about

One pattern: **in every job, a client or a peer does most of the talking.** Kestrel hosts, invites, prepares and follows up. Clients and peers do the talking.

Step 4: reconcile, the events plan

Event	Format	How often	Moments	Signals	Priority
Owners' round table	12 to 16 owners, dinner, one client owner speaking, co-hosted with an association or an accountant	Twice a year	Owner options, owner checking	Coverage, strategy	P1
Talks at sector association events	A client COO and Kestrel's head of delivery, 30 minutes, on the benchmark	Four a year, one per sector	COO trigger, COO checking by sector	Canon, coverage	P1
Open plant visit	Half a day at a client's plant, 10 to 15 plant managers and improvement leads	Twice a year	Plant manager framing; improvement lead checking	Coverage, strategy	P2

Event	Format	How often	Moments	Signals	Priority
Trade show	A talk on the main stage, not a stand	Once a year	COO trigger	Canon	P3
Improvement leads' community	Run by Customers; Events helps with the room	Twice a year	Improvement lead extend	Coverage	Customers'

Rejected, with reasons:

- **The trade-show stand.** 140 badge scans, three fit calls, no reviews. The same money funds both round tables.
- **Webinars.** No moment asks for them. The benchmark is presented once a year, recorded, and that recording is enough.
- **A large annual client event.** Owners want one peer in a small room, not a conference.
- **Sponsoring conference dinners without speaking.** It buys a logo on a menu. That's advertising's nod, not this build.

The anatomy library

Event asset	Parts, in order	Length
The invitation	The question the evening answers, in the guest's words · Who's speaking, and why they're worth hearing · Who else is coming (by role, never by name without permission) · Date, place, length · How to say yes	Under 150 words
The run of show	Welcome, 5 minutes, by the co-host · The speaker, 20 minutes: their problem, what they did, what it cost, what they'd do differently · Questions, 30 minutes · Dinner, unstructured · Close, one line	Two to three hours
The speaker brief	Who's in the room, and what they're wondering · What the speaker might cover · What they shouldn't be asked to say · Timing	One page

Event asset	Parts, in order	Length
The talk (association events)	The finding, with its number · The client's story: the problem, the floor, the result · What to check this week · One slide of method	30 minutes, 10 slides at most
The plant visit	What the host plant changed · Walk the line with the host's plant manager · Questions on the floor · Coffee, one question each: what would you check first in your plant?	Half a day
The follow-up	Thanks, and the one thing they asked about · The smallest next step, if any · An easy no	Under 100 words, within 48 hours

Three rules hold for every event:

- **No selling in the room.** Kestrel's partners host and listen. The next step comes in the follow-up, and only if the guest asked a question that leads there.
- **Speakers say what happened, including what went wrong.** A peer is believed because of the part that went wrong.
- **Every guest is followed up by the person who met them,** within 48 hours, about what they actually asked.

Briefs

EVENT	Owners' round table, March
MOMENTS	Owner · options: "Do we need outside help at all?" Owner · checking: "What do people I trust say about them?"
CO-HOST	The metal and engineering association
SPEAKER	A client owner: managing director of a 230-person fabricator
GUESTS	14 owners: 9 from the association's list, 3 from accountants' introductions, 2 in live deals with Kestrel
KESTREL	Jo Brady and the managing partner: host, listen, no pitch
THE QUESTION	"When does an owner need outside help, and when not?"
FOLLOW-UP	Within 48 hours, from whoever sat with each guest
MEASURE	Guests who ask for a conversation within a month; owners in live
	deals who move forward

The finished example: the owners' round table, written in full

The invitation, the speaker's brief, and one follow-up, for the March round table. The association, the speaker, the guests and Kestrel are invented.

1. The invitation, sent by the association's chair, co-signed by Kestrel's managing partner.

Subject: When does an owner need outside help? Dinner with owners of engineering firms, 12 March

Dear Graham,

Every owner I talk to has asked themselves the same question at some point: are our problems ones we can fix ourselves, or do we need someone from outside?

On 12 March, Helen Dorsey, managing director of a 230-person fabrication business, will tell a small group of owners how she answered it: what she tried first, why she brought someone in, what it cost her in time and disruption, and what she'd do differently.

It's dinner for 14 owners of manufacturing businesses, from 6pm to 9pm, at the association's rooms. No presentations beyond Helen's 20 minutes, and nothing will be sold. Kestrel Advisory, who worked with Helen's business, is hosting with us.

If you'd like to come, reply to this email. Places are limited.

Margaret Kaye, chair, Metal and Engineering Association

2. The speaker brief, for Helen Dorsey.

Who'll be there: 14 owners of engineering and fabrication businesses, 80 to 400 staff. Most have never hired an operations consultant. Two are talking to Kestrel now. **What they're wondering:** whether outside help is worth the disruption; whether a smaller firm can be trusted with their people; what it really cost. **What you might cover:** what you tried first; why you brought someone in; the first two weeks on the floor, including what your people thought; what changed, in your figures; what you'd do differently. **What we won't ask you to do:** recommend Kestrel. Tell it as it was. If you'd use someone else next time, say so. **Timing:** 20 minutes after the welcome, then questions for as long as people want.

3. A follow-up, from Jo Brady to an owner who sat at her table, 36 hours later.

Subject: The changeover question you asked Helen

Graham, good to meet you on Tuesday. You asked Helen how long it took her team to get changeovers down, and whether it held. The short answer from her plant: about four weeks to halve them, and it's held for two years because her own team runs the routine.

If it's useful, I'd be glad to show you how she measured it before deciding anything, so you can check your own lines first. If not, I hope to see you at the next one.

Jo Brady, Kestrel Advisory

What to notice. The invitation is signed by the association, about the guest's question, and says plainly that nothing will be sold. The speaker brief tells a client not to recommend Kestrel, which is why the owners will believe her. The follow-up answers the one thing that owner asked, offers a way to check for himself, and makes saying no easy. One evening; two a year.

Templates

The events calendar (one row per event)

Date	Event	Format	Moments	Co-host	Speaker	Guests (by role)	Kestrel attendees	Follow-ups done in 48 h	Outcomes at 30 days

The event brief

EVENT	[name, date]
MOMENTS	[chooser · step, and their question]
CO-HOST	[association, accountant, client]
SPEAKER	[who, and why the room will believe them]
GUESTS	[how many, by role, from where]
KESTREL	[who attends, and their job in the room]
THE QUESTION	[the one the evening answers]
FOLLOW-UP	[who, within 48 hours]
MEASURE	[what should happen within 30 days]

The guest log (one row per guest)

Guest	Role, company	Who met them	What they asked	Follow-up sent	Next step

In the campaign plan

Events are occasions: dated, and shared by the category. The Campaign Engine builds around them, and this build makes sure each occasion has a room where a moment is decided.

Occasion	Campaign	The event	What other builds bring
The benchmark's release (Q2)	The annual benchmark	Talks at the four sector associations	Content: the sector cuts; PR: the trade coverage; Email: invitations to the list
Budget season (Q3)	Board season; "Before you buy another line"	The trade show's main-stage talk	Advertising: the COO's trigger; the pitch: the kit
Spring and autumn	Owners' round tables	Two dinners	Partners: the associations and accountants fill the room; Customers: the speaker
Twice a year	Plant visits	Half a day at a client's plant	Customers: the host; Content: the guide the visitors take away

Running it

Who owns what. At Kestrel, the marketing lead owns the events calendar and the logistics; the managing partner owns the round tables and the relationships with co-hosts; the head of delivery owns the talks and the plant visits. Every partner who attends owns the follow-ups for the guests they met.

When	What happens
Eight weeks before	Co-host and speaker confirmed; the speaker brief agreed
Four weeks before	Invitations from the co-host; guests confirmed by role

When	What happens
The day before	The guest list read by everyone attending: who's coming, which moment they're in, what they might ask
Within 48 hours after	Every follow-up sent, by the person who met the guest
30 days after	Outcomes logged: conversations, fit calls, deals moved
Yearly	The calendar against the moments; anything that produced nothing for a year is dropped

The speaker's thanks. A client who speaks gives up an evening and some privacy. Thank them in writing from the managing partner, send them the evening's questions, and don't ask again for a year.

Common mistakes

The mistake	The stage it breaks	The bias it ignores	At Kestrel	Instead
A stand because competitors have one	Known	We follow others	140 badge scans, 3 fit calls	Be where the moment is decided, not where the crowd is
Counting badge scans	All, unseen	Good enough beats best	"Great show, 140 leads"	Count conversations that moved a moment
Kestrel doing the talking	Believed	We follow others	Partners presenting the firm	Clients and peers speak; Kestrel hosts
Selling in the room	Preferred	Losses loom larger	A pitch over dessert	No selling; the follow-up offers one small step
Following up a week later, with a brochure	Bought	A hard choice stalls	A generic "thanks for visiting our stand" email	Within 48 hours, about what they asked

The mistake	The stage it breaks	The bias it ignores	At Kestrel	Instead
The wrong room	Known	Familiar feels safe	Conferences full of consultants	Rooms full of the people who decide: owners, COOs
Webinars nobody asked for	Answered	Good enough beats best	Not yet, and rightly	Record the one presentation that's worth it
Asking clients to speak too often	Believed	Losses loom larger	—	Once a year per client, and thank them properly

The bias to watch in the team: **the spotlight effect**. Teams think a stand or a logo makes them noticed. Buyers remember the one conversation that answered their question, and who was in the room when it happened.

The business check and the measures

- **Revenue:** owner-led firms are a buying path Kestrel serves, and owners are the hardest chooser to reach. The round tables are worth the partners' evenings.
- **One truth:** speakers say what happened; Kestrel's people say only what Credentials can sign.
- **The proof is real:** every client speaker has a confirmed result. No one speaks before their 90-day check.
- **Capacity:** each round table takes two partners for an evening and about 15 hours of preparation and follow-up. The trade-show stand took 60.

Measure	What it tells you	Stage
Guests who ask for a conversation within 30 days	Whether the room moved the owner's options moment	Found
Owners in live deals who move forward after a round table	Whether hearing a peer settles checking	Believed
Follow-ups sent within 48 hours	Whether the team does the part that matters	Bought
Plant-visit guests whose plants become fit calls	Whether seeing it on a floor answers the plant manager	Believed

Measure	What it tells you	Stage
Cost per conversation that moved a moment	Whether an event earns its place against the others	All

The Loop for this build

After each event, and each half-year across all of them, the Loop (step 11) asks: **did the conversations in the room move the moments they were for?**

What the Loop finds	What it means	What changes
Guests enjoy it, nobody asks for a conversation	The question wasn't theirs, or the follow-up missed	Ask guests what they'd want next time; check the follow-ups
Owners in live deals don't move after a round table	The speaker didn't match their situation	Match speakers by size and sector
One format produces most outcomes	The others are weaker rooms	Move effort to it; drop the weakest
Co-hosts keen to repeat	The exchange is working	Make it a fixture on their calendar

What Kestrel's first year would look for. Two round tables with 12 or more owners each, and at least two owners asking for a conversation from each. Owner checking moments moving forward in live deals after a guest heard a client speak. And the trade-show budget, moved to rooms that produce conversations, with the same or less spent.

Shapes by business type

Business type	Typical mode	The events that usually matter most
B2B professional services	Committee or owner-led	Peer round tables, talks with clients at association events, site visits
B2B software	Committee	User conferences, customer advisory boards, workshops for one buying group, live demos at the category's main event

Business type	Typical mode	The events that usually matter most
Considered consumer (finance, health, education)	Considered personal	Open days, seminars, one-to-one consultations
Ecommerce	Quick choice or habitual	Pop-ups, launches, sampling; rarely decisive on their own
Local services	Owner-led or considered personal	Local shows, community events, workshops in the shop or on site
Marketplaces	Varies on each side	Seller meet-ups and training days; rarely buyer events

Feeling a bit overwhelmed?

You should be, a little. One firm: a calendar set by the category, two round tables a year with a co-host and a client speaker, four talks, two plant visits, a brief for each, a guest list read the night before, a follow-up within 48 hours for every guest, and outcomes logged a month later. Most firms book a stand, count badge scans, and send everyone the same brochure a week later.

RocMode does part of the heavy lifting:

- **Campaigns timed to your category's key dates.** RocMode's campaign concepts are built around your category's calendar, so the occasions events serve are in the plan from the start.
- **The people in the room, and their questions.** Your choosers and their moments, mapped from what's known about your brand, so every event starts from who needs to hear what.

[Join the RocMode beta](#)

The Email Playbook

What this build is for

Email reaches people who already know you, with the one thing they need next, in words they can forward.

It's the oldest channel in this playbook and the most misused. Most firms run it as a broadcast: a newsletter on a schedule, a promotion when sales are slow, a sequence for every download. That reaches everyone and answers no one. This playbook runs it the other way: a message exists only because someone, at a known moment, has a question it answers.

Email belongs to the Relationship category: your system, talking to people you already know. That gives it two jobs. **Inside a relationship**, it keeps the people around a deal or a review informed at the moments they'd otherwise wonder: is it on track, are the savings real, what happens next. **Around the relationship**, it keeps Kestrel's thinking in front of people who aren't buying yet, through a newsletter worth reading, cut from Content's core pieces.

Email writes every message other builds send. Automation decides when a message goes; the pitch decides what goes in the proposal; Customers decides when to ask for a referral. The words come from here, written once, in the brand's voice, checked against Credentials.

The footprint, and what it reads from the chain

In Kestrel's map, email leads **3 cells and 3 moments**, all during and after the review, where a busy executive needs to know without asking:

Moment	Who	Their question
Delivery	COO	"Are we on track, and what are they finding?"
Delivery	Finance director	"Are the savings real, and are they being counted properly?"
Extend and refer	Finance director	"Are the savings showing in the accounts?"

It supports **14 more cells**, among the widest support of any build: the business case moments for the COO, owner and finance director (7 of them), the investor's trigger, week zero and the plant manager's weekly tasks, the follow-on for the COO, owner, plant

manager and investor, and the improvement lead's delivery and next review. In almost all of them, another build decides the moment and email carries the words.

What it reads, and never rewrites:

From	It takes
Identity	The voice, and the words to use and avoid
Credentials	Every claim and figure a message states
Where you stand	Which moment each message serves, and the chooser's question
Content	The core pieces the newsletter is cut from
The pitch and Automation	When each message goes, to whom, and what happens on a reply

Its slice of the Score, and its own audit

Email's part of [the Selectability Score](#) is **Answered** and **Believed** at the moments after the signature: does the COO know how it's going, and does the finance director believe the savings?

Kestrel today, the 3 moments email leads	Answered	Believed	Preferred	Bought	Moment score
COO · delivery	75	80	80	75	78
Finance director · delivery	45	50	70	65	57
Finance director · extend	40	45	60	50	49

The pattern is clear: Kestrel keeps the COO informed and leaves the finance director to guess. The COO hears from the engineers on site; nobody writes to the finance director until the invoice. Yet the finance director decides the next review, and their "are the savings real?" is the question every future proposal has to answer.

The audit.

What Kestrel sends today	To whom	What it misses
Engineers' weekly notes, each in their own format	The COO, sometimes	The finance director; a consistent record of findings and savings

What Kestrel sends today	To whom	What it misses
A final report	The COO	A plain summary the finance director can check against the accounts
Nothing after the review	—	The 90-day check; the programme decision; the referral
A quarterly newsletter of firm news	1,400 contacts, all of them	Any question a reader has. Opened by 14%, answered by none

The newsletter is the same mistake as the old blog: it exists because a calendar said so.

Step 1: the shape, the lifecycle programme

Email's shape is **the lifecycle programme**: the streams of messages a relationship goes through, from first contact to the next piece of work. Every category has the same skeleton; what changes is how long each stream lasts and who reads it.

Stream	When it runs	Who it's for at Kestrel	In the category (of 5 firms)
Welcome	After a first enquiry or download	Whoever asked	4
Newsletter	Regularly, to people not yet buying	COOs, finance directors, improvement leads, investors	5
Deal	From fit call to signature	The buying group	2
Onboarding	Signature to week one	Plant manager, improvement lead, procurement	1
Delivery	Each week of the work	COO, finance director, plant manager	1
Results and renewal	After the work	Finance director, COO, owner	0
Re-engagement	When a past client might need more	COO, procurement, investor	1

Canon is just the welcome and the newsletter: everyone does them, so their absence would be noticed. The opening is in the streams almost nobody runs: delivery, results and renewal. That's where Kestrel's moments are, and where its competitors leave the finance director guessing.

Step 2: coverage, one message per question at each step

The coverage unit is **a message per question per step**. For each of the 17 cells email leads or supports, take the chooser's question and ask: is this answered best by a message arriving on its own, at the right moment? If yes, it gets a message.

Stream	Moments	Messages
Welcome	Investor trigger; any enquiry	2: an acknowledgement with a next step; for investors, the portfolio note
Newsletter	Trigger and framing, for four choosers (supporting Content)	1 a month, cut from the core pieces: 12 a year
Deal	COO, owner and finance director building the case (9 moments)	3: the business case kit with their figures; the board paper; a note before the board meets
Onboarding	Plant manager and improvement lead · buying; improvement lead · scope	3, written here, sent by Automation's week-zero flow
Delivery	COO, finance director and plant manager · delivery; improvement lead · delivery	3 a week for six weeks: Monday tasks, the Friday update, the finance director's savings tally
Results and renewal	Finance director, COO, owner, plant manager · extend	4: the results summary; the 90-day check; the programme note; the referral ask
Re-engagement	Investor and improvement lead · extend	2: the portfolio note; "running the next review yourselves"

That's 14 messages written once, three weekly templates, and a monthly newsletter. Every message has a moment; none is there because a sequence needed a fourth email.

Step 3: gaps, biases and jobs

Moment	Stage	Bias	The job
Finance director · delivery: "Are the savings real?"	Believed	We follow others	A weekly savings tally from week three: each finding, its value, how it was counted, and who on the client's side agreed it
Finance director · extend: "Are they showing in the accounts?"	Believed	Losses loom larger	The 90-day check: the finding, the figure, a single question, "does this match what you see?"
COO · delivery: "Are we on track?"	Answered	A hard choice stalls	One Friday update in one format: what was found, what changed, what's next, what's needed from them
Plant manager · delivery	Answered	A hard choice stalls	Monday tasks short enough to read on the floor; SMS if they prefer
COO, owner and finance director · building the case	Believed	Losses loom larger	The kit and the board paper sent with a covering note in each reader's terms, not one note for all
COO · extend: "Do we go into the programme?"	Preferred	A hard choice stalls	The results summary with the programme's cost and its likely return, side by side
Owner · extend: "Would I tell another owner?"	Preferred	We follow others	The referral ask, only after confirmed savings, made easy: one forwardable paragraph
Every newsletter reader	Known	Familiar feels safe	A newsletter that answers one buyer question a month, from Content's pieces; firm news moved to a short footnote

The biggest gap is one person. **The finance director is written to at none of the moments where they decide what Kestrel is worth.**

Step 4: reconcile, the email plan

Stream	Messages	Moments	Sent by	Signals	Priority
Delivery	Friday update; finance director's savings tally; Monday tasks	4	Automation (review rhythm)	Coverage, strategy	P1
Results and renewal	Results summary; 90-day check; programme note; referral ask	4	Automation (savings confirmation, programme decision)	Coverage, strategy	P1
Deal	Kit cover note per reader; board paper note; pre-board note	9	The deal owner, from the pitch	Coverage	P1
Onboarding	Week-zero messages to plant manager, improvement lead, procurement	3	Automation (week zero)	Coverage	P2
Newsletter	One a month, cut from Content	4 choosers' trigger and framing	Marketing lead	Canon, coverage	P2
Welcome	Acknowledgement; investor portfolio note	2	Automation (enquiry response)	Canon	P2
Re-engagement	Portfolio note; "running the next review yourselves"	2	Automation (re-engagement)	Coverage	P3

Every moment email leads or supports has a message; every message has a sender. The two P1 streams are the ones no competitor runs.

Rejected, with reasons:

- **A five-email sequence after every download.** Nobody downloading a cost article is waiting for five emails. One acknowledgement, then the newsletter if they ask for it.
- **Promotional emails.** Kestrel has nothing to discount, and no moment asks for an offer.
- **The firm-news newsletter.** Replaced. News survives as two lines at the foot of the monthly issue.
- **Personalisation by first name only.** It answers nothing. Personalise by role and moment, or not at all.

The anatomy library

Message type	Parts, in order	Length
The update (Friday, delivery)	Subject: the week and the headline finding · What we found · What changed on the floor · What's next · What we need from you, if anything	Under 200 words
The tally (finance director)	The running total, with its range · Each finding: value, how counted, who agreed it · What isn't counted yet, and why	A table and three lines
The check (90 days)	The figure we said · One question: does it match? · How to reply in one line	Under 100 words
The cover note (deal)	The reader's question, answered in the first line · What's attached and why it's theirs · The one thing to do · Who to call	Under 120 words
The ask (referral)	What was achieved, in their figures · The ask, specific: one person, one call · A paragraph they can forward · An easy no	Under 150 words
The newsletter issue	One buyer question, in the subject · The answer, in the first paragraph · One figure, sourced · A link to the full piece · Two lines of firm news	300 to 500 words

Five rules hold for every message:

- **It comes from a person, by name**, who will read the reply.
- **The subject line answers or names the question**. Never "Update" or "News".
- **One thing to do, or nothing**. An update with three asks gets none of them done.
- **Every figure is signed or theirs**.
- **It can be forwarded as it is**. The COO forwards the Friday update to the board; write it for the board.

Briefs: the programme brief, then a brief for every message

STREAM	Delivery
RUNS	Each week of a review, weeks 1 to 6
MESSAGES	Monday tasks (plant manager); Friday update (COO, copied to the
	finance director from week 3); savings tally (finance
director,	from week 3)
MOMENTS	COO, finance director, plant manager · delivery
FROM	The lead engineer, by name
SENT BY	Automation: review rhythm, from the delivery tracker
VOICE	Plant language; findings in the plant manager's words
OWNER	Head of delivery
MEASURE	Replies that confirm or correct a finding; the COO forwarding

MESSAGE	The savings tally
JOB	Believed · finance director · delivery · we follow others
THEIR QUESTION	"Are the savings real, and are they being counted properly?"
ANSWER FIRST	The running total, with a low and a high, and how it's counted
PROOF	Each finding agreed by a named person at the client
THE ONE THING	Tell us if any line doesn't match your view
GUARD	No finding counted until the client's own person has agreed it

The finished example: Hartwell's week four, and the 90-day check, written in full

Three messages from Kestrel's review at Hartwell Plastics, the account from the Pitch Playbook. Hartwell, its people and its figures are invented.

1. The Friday update, week four. To the COO, Lena Marsh; copied to the finance director, Tom Avery.

Subject: Week 4 at Hartwell: warm-up scrap after changeovers is the biggest loss so far

Lena, four weeks in, here's where we are.

What we found. After each mould change on lines 2 and 4, the first 40 minutes of parts are mostly scrap while the tool comes up to temperature. That's nine changes a week, and the largest single source of the 5.8% scrap rate. **What changed on the floor.** Dave's crew on line 2 trialled pre-heating the next mould during the last hour of the run. Warm-up scrap fell by about two-thirds on the four changes we measured.

What's next. The same trial on line 4 next week, and a look at regrind ratios, which drift on night shift. **What we need from you.** Nothing this week.

Priya Raman, lead engineer, Kestrel Advisory

2. The savings tally, week four. To the finance director, Tom Avery.

Subject: Savings found so far at Hartwell: A\$185,000 a year (range A\$150,000 to A\$210,000)

Tom, this is the running tally, and how each line is counted. A finding only appears here once someone at Hartwell has agreed it.

Finding	Annual value	How it's counted	Agreed by
Warm-up scrap after changeovers, lines 2 and 4	A\$96,000	Scrap kg saved × material cost, from your own scrap logs	Dave Okafor, plant manager
Regrind ratio drift on night shift	A\$58,000	Virgin material saved at your current resin price	Anna Lee, quality lead
Compressed-air leaks	A\$31,000	Metered kWh saved × your tariff	Sione Taufu, maintenance lead

Not counted yet: the line 4 trial, until it's measured next week. If any line doesn't match your view, reply and I'll walk you through it.

Priya Raman

3. The 90-day check. To Tom Avery, 90 days after the review closed with A\$214,000 agreed.

Subject: Hartwell, 90 days on: are the A\$214,000 of savings showing?

Tom, when the review closed in December, we agreed savings of A\$214,000 a year, mostly scrap and resin. Three months on, does that match what you see in the accounts?

A one-line reply is plenty: "yes", "about A\$x", or "not yet". If it's not showing, I'd like to understand why, and we'll come back on site to find out.

Priya Raman

What to notice. Each subject line is the answer. The update is written for the board as much as the COO, because it will be forwarded. The tally counts nothing the client hasn't agreed, which is why the finance director believes it, and why, at 90 days, a "yes" becomes a signed case for the next proposal. And none of the three asks for anything the reader can't do in a minute. Three messages, for one week and one check, at one client.

Templates

The message inventory (one row per message; the build's single list)

ID	Stream	Moment (chooser · step)	Their question	Subject line	From	Sent by (person or flow)	Claims used	Live?

The stream brief

STREAM	[name]	RUNS	[when, how long]
MESSAGES	[each one, to whom]		
MOMENTS	[chooser · step]		
FROM	[a named person]	SENT BY	[person or flow]
VOICE	[from Identity: words to use and avoid]		
OWNER	[who reads the replies]		
MEASURE	[the stage it serves]		

The message brief

MESSAGE [name]
 JOB [stage · chooser · step · bias]
 THEIR QUESTION [in their words]
 SUBJECT [the answer, or the question]
 ANSWER FIRST [the first line]
 PROOF [signed claims, or their own figures]
 THE ONE THING [what they do, or nothing]
 GUARD [passes only if it can be forwarded as it is]

The newsletter plan (one row per issue)

Month	The buyer question	Chooser	Core piece it's cut from	The one figure, with source	Footnote news

In the campaign plan

Every campaign has an email moment, usually more than one. The rule: **a campaign's emails are written from its core piece and sent to people for whom its question is live, never to the whole list.**

Kestrel campaign	The issue or message	To whom
The annual benchmark	The newsletter issue on the headline finding; a personal note from a partner to past clients with their sector's figure	Newsletter readers; past clients by sector
"Before you buy another line"	The issue "What flat output really costs"; a cover note with the kit for anyone who downloads it	COOs and finance directors on the list
Board season	A dated note to finance directors in live deals: the kit, ready for this budget round	Live deals only
Owners' round table	The invitation, from the speaker, by name; the write-up afterwards	Owners on the list; attendees
Results from the year's reviews	A results issue: what the year's reviews found, confirmed by clients	Everyone on the list, once a year

Running it

Who owns what. At Kestrel, the marketing lead owns the message inventory, the voice and the newsletter. Every message in a deal or a review is sent in the name of the person who'll read the reply: the deal owner, or the lead engineer. The head of delivery signs the delivery templates; the managing partner signs anything that states the way out.

When	What happens
Weekly, during each review	The Friday update and the tally, checked by the lead engineer before they go
Monthly	One newsletter issue, from the content plan; replies read and answered
Quarterly	Every template re-read against the claims register and the voice; unsubscribes and replies reviewed
When a claim changes	Every message that states it, the same week
Yearly	The list cleaned: anyone who hasn't opened in a year is asked once whether to stay

Permission. Messages in a deal or a review go to the people in it, about it. The newsletter goes only to people who asked for it. Keep the two lists apart, so nobody joins the newsletter because they once signed a contract.

Common mistakes

The mistake	The stage it breaks	The bias it ignores	At Kestrel	Instead
Writing to the champion only	Believed, for the finance director	Losses loom larger	The finance director heard nothing until the invoice	A message to each person whose moment it is
Broadcasting to the whole list	Known, worn down	Good enough beats best	Firm news to 1,400 people, 14% opening	One buyer question a month, to people who asked
Subject lines that hide the point	Answered	A hard choice stalls	"Update", "Newsletter Q3"	The answer in the subject

The mistake	The stage it breaks	The bias it ignores	At Kestrel	Instead
Counting savings nobody agreed	Believed	We follow others	Totals the finance director couldn't check	Only findings a named client person has agreed
Sequences for their own sake	Preferred	Good enough beats best	Five emails after a download	Send what someone is waiting for, then stop
Sent from "the team"	Believed	Familiar feels safe	noreply@ on the final report	A named person who reads the reply
Asking for the referral too early	Preferred	Losses loom larger	A referral ask in the final report, before any savings showed	After the 90-day check, and only after a yes
Messages written in the flow tool	Believed	Familiar feels safe	Terms in an automated email that didn't match the contract	Every message from the inventory, every claim from Credentials
Judging by opens	All, unseen	Good enough beats best	A 30% open rate celebrated	Replies, confirmations and forwards to the board

The bias to watch in the team: **the curse of knowledge** again, in its email form. The engineer knows what "OEE up 6 points" means; the finance director needs the money.

The business check and the measures

- **Revenue:** the delivery and results streams protect the next review at every client and make the proof for every new one. If only one stream can be built, build the savings tally and the 90-day check.
- **One truth:** every figure in every message matches Credentials and the client's own records.
- **The offer:** the 90-day check asks a question the firm must be ready to hear "no" to. Decide in advance what Kestrel does when savings aren't showing.
- **Capacity:** the Friday update takes the lead engineer 20 minutes with the template, an hour without it.

Measure	What it tells you	Stage
Finance directors who reply to the tally during the review	Whether the savings are being believed as they're found	Believed
90-day checks answered, and the share answered "yes"	Whether results are real, and becoming proof	Believed
Friday updates forwarded inside the client	Whether the update is written for the board	Answered
Programme or next-plant decisions within 60 days of results	Whether the results stream leads somewhere	Preferred
Referrals given after a "yes"	Whether proof turns into the next client	Preferred
Newsletter replies per issue	Whether the issue answered a question	Known

The Loop for this build

Each quarter the Loop (step 11) reads replies and decisions beside the Score's delivery and extend moments, and asks: **did the people we wrote to believe the results, and act on them?**

What the Loop finds	What it means	What changes
Tallies sent, finance directors silent	The tally isn't in their terms, or isn't reaching them	Rewrite in the money and method they use; check the address
90-day checks answered "not yet"	Savings aren't sticking	A Customers job, not an email one: the 90-day visit
Results sent, no programme decisions	The results summary doesn't lead to a choice	Put the programme's cost and return beside the results
Newsletter opens high, replies none	Issues answer nothing anyone asked	Pick next quarter's questions from sales calls
Updates forwarded to boards	The delivery stream is working	Use the format in proposals: "here's what you'll get every Friday"

What Kestrel's first year would look for. Finance directors replying to the tally during the review, which nobody did last year. Then the 90-day checks: if most say "yes", Kestrel has more confirmed cases in a year than it has in its history, and the finance director's "consultants overstate savings" has an answer in every proposal.

Shapes by business type

Business type	Typical mode	The streams that usually matter most
B2B professional services	Committee or owner-led	Delivery updates, results and renewal, a newsletter that answers buyer questions
B2B software	Committee	Onboarding to first value, usage and adoption notes per role, renewal and expansion
Considered consumer (finance, health, education)	Considered personal	Application progress, plain-language documents, reminders before deadlines, annual reviews
Ecommerce	Quick choice or habitual	Order and delivery, replenishment, reviews after use, a few well-chosen offers
Local services	Owner-led or considered personal	Booking confirmations, the day-before reminder, a follow-up after the job, seasonal reminders
Marketplaces	Varies on each side	Seller onboarding and payouts; buyer order updates; trust and safety notices

Feeling a bit overwhelmed?

You should be, a little. One firm: seven streams, 14 messages written once, three weekly templates for every review, a monthly newsletter cut from the content plan, every figure agreed by a named person at the client, and two lists that must never mix. Most firms send a quarterly newsletter of their own news and wonder why the finance director never believes the savings.

RocMode does part of the heavy lifting:

- **Emails and sequences, drafted on brand.** RocMode's Email Studio drafts your messages and sequences in your voice, from your strategy.
- **The moments behind every message.** Your choosers, their steps and their questions, mapped from what's known about your brand, so each message starts from a question someone actually has.

[Join the RocMode beta](#)

The Automation Playbook

What this build is for

Automation is for the moments that are lost because nobody replied in time. Not the moments nobody cares about: the ones everyone cares about, but nobody owns at five on a Friday.

The classic evidence is about speed. An audit of more than 2,000 US companies, published in the Harvard Business Review, found that firms which replied to an online enquiry within an hour were nearly seven times as likely to have a meaningful conversation with the buyer as firms that replied even an hour later ([HBR](#)). Speed matters most where a buyer is waiting, and in a committee sale, someone almost always is: procurement waiting for the insurance certificate, the plant manager waiting to know what to have ready, the improvement lead waiting for the data request.

So the job isn't "automate marketing". It's narrower and more useful: **notice the moment, send what that person needs, and hand anything that matters to a named human.** Automate the noticing, not the talking.

The Relationship category's nodes live here: **CRM** (the record every flow reads from and writes to) and **SMS and chat** (two more ways a flow can reach someone), each with a short section below.

The footprint, and what it reads from the chain

In Kestrel's map, automation leads **4 cells and 5 moments**, all in the gap between "yes" and the first day on the floor, plus the weeks that follow:

Moment	Who	Their question
Buying (2: planned and urgent)	Procurement	"Site inductions, payment terms and data handling?"
Buying	Plant manager	"What do we need ready before they arrive?"
Buying	Improvement lead	"What data do they need from us before week one?"
Delivery	Plant manager	"What do my people need to do this week?"

It supports nine more cells, where another build leads but a timely reply decides whether the moment holds: procurement's supplier-list and safety checks, the COO's and investor's start dates, the COO's and finance director's progress and savings updates, the COO's programme decision, the finance director's look at the accounts, and procurement's re-engagement.

What it reads, and never rewrites:

From	It takes
Where you stand	Which moments each flow serves, and each chooser's question in their words
Sales Strategy	The motion steps and routing: who owns the account at each step, and when a human must take over
Credentials	The documents flows send (insurance, safety, terms), each one signed and current
The pitch	Week zero, the start dates, the deal map's people and roles
Email	The messages flows send, written once and reused
Identity	The voice every automated message speaks in

Its slice of the Score, and its own audit

Automation's part of [the Selectability Score](#) is mostly **Bought** (did the step run without a stall?) and **Answered** (did the person get what they needed, when they needed it?).

Kestrel today, the 5 moments automation leads	Answered	Believed	Preferred	Bought	Moment score
Procurement · buying	35	45	50	40	43
Plant manager · buying	55	60	70	60	62
Improvement lead · buying	50	55	65	55	57
Plant manager · delivery	80	80	85	75	81

These are small moments in points: together they earn 1.7 of Kestrel's 18 and lose 1.3. Their weight is in what they protect. A deal that stalls in procurement, or a review that starts a week late, damages the moments after it: the COO's confidence, the finance director's savings date, the reference the client might give.

The audit. List every hand-off in the motion, who does it today, and how long it takes.

Hand-off	Who does it today	How long	What it costs
Reply to a website enquiry	Whichever partner sees the inbox	Median 26 hours	Enquiries that went cold before the first call
Procurement's documents (insurance, safety, terms)	An office manager, by hand, from a shared drive	Median 4 working days	Contracts signed a week after the verbal yes
The data request before week one	The lead engineer, from memory	Sent the week before the start, or not at all	3 of 9 reviews last year started late, waiting for data
The weekly update during the review	Each engineer, in their own format	Some weeks, not others	The COO hearing progress second-hand
The 90-day savings check	Nobody	Never	Savings confirmed only when a client volunteered it, as two did

The CRM holds contacts, not moments: nothing in it knows that a deal has reached week zero, so nothing fires.

Step 1: the shape, the systems map

For automation, the shape is **the systems map**: every system that holds a fact about a buyer, every event it can notice, and every way it can reach someone. You can't automate a moment no system can see.

System	What it knows	Events it can notice	Kestrel today
CRM	People, accounts, deal stage, who owns the account	A deal moves stage; a new contact joins an account; no contact for N days	Has contacts; deal stages not kept
Website and forms	Who asked what, from which page	A form sent; a kit downloaded; a return visit from a known account	A contact form to a shared inbox

System	What it knows	Events it can notice	Kestrel today
Email	What was sent, opened, replied to	A reply; no reply after N days	Partners' own inboxes
Calendar	Meetings booked	A fit call booked; a start date set	Partners' own calendars
Document store	The signed claims, insurance, safety, terms, templates	A document expires; a new version is signed	A shared drive, unversioned
E-signature	Contracts out and signed	A contract signed	Contracts signed on paper and scanned
Delivery tracker	The review's week, tasks, findings	A week ends; a finding is logged; a task is late	Each engineer's spreadsheet
Finance system	Invoices, payments	An invoice paid; a payment late	Accounts package, not connected
SMS and chat (nod)	Short messages to people on the floor	A message sent or answered	Not used

The canon question here is simple: what do buyers in the category already expect to happen automatically? For B2B services, three things: an instant acknowledgement with a real next step, documents on request without a wait, and a calendar link instead of email tennis. Kestrel has none of the three.

The CRM, as a nod. The CRM isn't a build; it's the memory every flow reads from and writes to. Its only rule in this playbook: **every moment that matters must be a field or a stage someone can see.** If "week zero" isn't a stage, nothing can fire on it.

Step 2: coverage, one trigger per moment lost for lack of a timely reply

The coverage unit is a **trigger**: an event a system can notice, tied to a moment where someone is waiting. For each of the 14 cells automation leads or supports, write four things: **who is waiting, for what, how long they'll wait before it costs you, and what the system can see.**

Moment	Who's waiting, for what	How long before it costs	What the system can see
Website enquiry (feeds every early moment)	Whoever asked: a reply with a real next step	An hour	Form sent
Procurement · comparing and checking	Procurement: insurance, safety record, supplier-list forms	Two days	Request on the procurement page; deal at proposal
COO and investor · buying	A start date	A day after the verbal yes	Deal moves to verbal yes
Procurement · buying (planned and urgent)	Site inductions, payment terms, data handling	Two days; the same day if urgent	Contract sent
Plant manager · buying	What to have ready	Two weeks before the start	Start date set
Improvement lead · buying	The data request	Two weeks before the start	Start date set
Plant manager · delivery	This week's tasks for their people	Monday morning, every week	Week begins in the delivery tracker
COO · delivery	Progress and findings	Friday, every week	Week ends; findings logged
Finance director · delivery and extend	Savings counted, then confirmed in the accounts	End of the review; then 90 days on	Review closed; 90 days pass
COO · extend	The programme decision, with the results	30 days after the review	Results confirmed
Procurement · re-engaging	A route without a new tender	When the next review is discussed	A new opportunity on an existing account

That's 11 triggers. Each one names a person and a moment; none is "nurture".

Step 3: gaps, biases and jobs

Hold the audit against the 11 triggers. Each one nothing fires on today is a gap; each gap becomes a job.

Moment	Stage	Bias	The job
Any enquiry	Bought	A hard choice stalls	Reply within the hour, with a calendar link for the fit call and one page that answers the question they asked
Procurement · checking	Believed	We follow others	Send the current, signed insurance and safety pack the moment it's asked for, from the document store, never from someone's inbox
Procurement · buying	Bought	Losses loom larger	When the contract goes out, send site-induction, payment and data-handling answers with it, so procurement doesn't have to ask
COO · buying	Bought	Losses loom larger	On the verbal yes, confirm the start date within a day, and alert the partner if it isn't set
Plant manager · buying	Answered	A hard choice stalls	Two weeks before the start, the week-zero checklist, in the plant manager's language
Improvement lead · buying	Answered	Losses loom larger	Two weeks before the start, the data request, with what each item is for and who will see it
Plant manager · delivery	Answered	A hard choice stalls	Every Monday, this week's tasks for their people, short enough to read on the floor; by SMS if they prefer
COO · delivery	Believed	We follow others	Every Friday, the week's findings, from the delivery tracker, in one format
Finance director · extend	Believed	We follow others	At 90 days, ask the finance director to confirm the savings; turn a yes into a signed case

Moment	Stage	Bias	The job
Procurement · re-engaging	Bought	A hard choice stalls	When a new opportunity opens on an existing account, send the framework terms first

Two rules come out of the jobs. **Every flow hands to a named person** when the buyer replies with anything but a yes. And **nothing is sent that a person wouldn't sign**: the documents come from Credentials, the words from the Email build.

Step 4: reconcile, the automation plan

Group the triggers into **flows**: a chain of steps that starts on one event and ends with a person, a document or a decision. This is the automation build's equivalent of a sitemap.

Flow	Starts when	Triggers it carries	Ends with	Signals	Priority
Enquiry response	A form is sent	1	A fit call booked, or the partner told	Canon, coverage	P1
Procurement pack	Procurement asks, or the deal reaches proposal	2	Documents received; procurement's questions to a person	Canon, coverage	P1
Verbal yes to start date	The deal moves to verbal yes	2	A start date in the calendar, the contract out with the procurement answers	Coverage, strategy	P1
Week zero	A start date is set	2	Checklist and data request confirmed by the plant manager and improvement lead	Coverage, strategy	P1
Savings confirmation	90 days after the review closes	1	The finance director's confirmation, turned into a signed case	Coverage, strategy	P1
Review rhythm	Each week of the review	2	Monday tasks read; Friday findings sent	Coverage	P2
Programme decision	30 days after results	1	A meeting with the COO, or a reason why not	Coverage	P2
Re-engagement	A new opportunity on an existing account	1	Framework terms sent; the partner owns it	Coverage	P3

8 flows, 11 triggers. Every lead moment has a flow; every flow ends with a person or a confirmed fact. The savings confirmation is P1 for a reason the footprint doesn't show: it's where the proof for every other build comes from.

Rejected, with reasons:

- **Lead scoring.** No moment asks for it. With 31 enquiries a year, the partners can read every one.
- **A drip "nurture" sequence.** It answers no question a buyer is waiting on. Newsletters belong to Email, sent because there's something worth reading.
- **A chatbot answering technical questions.** A wrong answer about a plant is worse than a slow one. Chat, as a nod, routes to a person and books calls.
- **Automatically sent proposals.** The proposal is the pitch's job, written for the people in the room.

The anatomy library

Every flow is built from four parts, and each part has a standard shape.

Part	What it holds, in order	Rule
The trigger	The system · the event · the conditions (account, stage, role) · what it must not fire on	One event, one moment. If it needs three conditions to be right, the CRM stage is wrong
The message	Who it's to · their question, answered in the first line · the one thing they need · who to reply to, by name	Written once in the Email build; never written inside the flow tool
The handover	When a person takes over · who · how they're told · by when	Any reply that isn't a plain yes goes to a person within a working day
The guardrail	What the flow must never do · the stop condition · who can switch it off	Every flow has an off switch and an owner

And three shapes of flow cover almost everything Kestrel needs:

Flow shape	Kestrel's flows	Steps
Acknowledge and route	Enquiry response, re-engagement	Reply at once · route to the owner · remind the owner if nothing happens

Flow shape	Kestrel's flows	Steps
Prepare	Procurement pack, verbal yes to start date, week zero	Send what's needed · confirm it arrived · chase once · hand to a person
Rhythm and check	Review rhythm, savings confirmation, programme decision	Send on the date · ask one question · record the answer · act on it

SMS and chat, as a nod. Both are just other ways for a message to reach someone. Use SMS where people don't sit at a desk: Kestrel's Monday tasks reach line leads on the floor by text, with their permission. Use chat on the website only to book a call or route a question to a named person, never to answer about a plant.

Briefs: the flow brief, then a brief for every step

A flow is briefed like a page: one brief for the whole, then one for each step, because each step serves one person at one moment.

The flow brief for Kestrel's savings confirmation:

FLOW	Savings confirmation
STARTS WHEN	90 days after a review is marked closed in the delivery tracker
MOMENTS	Finance director · extend: "Are the savings showing in the accounts?"
PEOPLE	Feeds: every build's proof (Credentials: confirmed cases) The client's finance director; Kestrel's lead engineer; the head of delivery
ENDS WITH	The finance director's confirmation in writing, or a call booked to understand why the savings aren't showing
OWNER	Head of delivery
OFF SWITCH	Head of delivery; stops automatically if the client is in dispute
MEASURE	Share of closed reviews with savings confirmed within 120 days

Then each step gets its own brief:

STEP	2: the confirmation request
JOB	Believed · finance director · extend · we follow others
TO	The client's finance director, from the lead engineer by name
THEIR QUESTION	"Are the savings showing in the accounts?"
THE ONE THING	Confirm or correct the figure in one reply
MESSAGE	From the Email build: "90-day check", filled from the delivery tracker
HANDOVER	Any reply but a yes goes to the lead engineer within a working day
GUARDRAIL	Never sent if the review's final findings weren't accepted

The flow check. Walk the flow as each person it touches, on a real date. Would they know why they got each message, what to do, and who to ask? Then walk it as the person who owns it: on the day it misfires, how would you know?

The finished example: Kestrel's week-zero flow, written in full

Week zero is where Kestrel lost three starts last year. The flow covers three of automation's five lead moments: the plant manager's and improvement lead's "what do you need before you arrive?", and procurement's site-induction questions. Here it is for Hartwell Plastics, whose review starts on 3 November. Hartwell and its people are invented.

The trigger. The deal moves to "start date set" in the CRM, with a date at least 10 working days away. It does not fire if the contract is unsigned, or if the start is in less than five days (the partner calls instead).

Day 0, the same hour. To the plant manager, Dave Okafor, from Kestrel's lead engineer, Priya Raman:

Subject: Hartwell review, starting 3 November: what we'll need from your team

Dave, we start on your floor on Monday 3 November. Here's everything we'll need, so nothing waits on us or on you.

From your team, week one: two hours a day of your time; one hour a shift from each line lead on lines 2 and 4; access to the floor on every shift. **Before we arrive:** a site induction booked for 8am on 3 November; two hi-vis vests and site passes; a desk near the lines. **What we won't need:** anyone taken off the line, or any change to your schedule in week one.

Reply with a yes, or with anything that doesn't work. If it's easier to talk, my number's below. Priya Raman, lead engineer, Kestrel Advisory

Day 0, the same hour. To the improvement lead, Sam Ito: the data request, with a line on each item saying what it's for and who will see it. Scrap logs for lines 2 and 4 since January; the last three months of downtime records; energy readings by line if they exist. "If any of this doesn't exist, that's a finding, not a problem. Tell us and we'll measure it on site."

Day 0. To procurement: the site-induction form, the insurance certificate and safety record (both current, from the document store), and the data-handling note. No one asks for them, because they've already arrived.

Day 3, if no reply from the plant manager. One reminder, shorter. **Day 5, still nothing:** Priya is told, and calls. The flow never chases twice.

Day 5, if the data hasn't arrived. Priya calls Sam. Not a reminder email: a call, because missing data is the usual reason a start slips.

Day 8. The COO, Lena Marsh, gets one line: "Week zero is on track: your plant manager and improvement lead have everything they need, and the induction is booked." Or, if not: what's missing, and who's on it.

The guardrail. If anyone at Hartwell replies with a question, the flow stops for that person and Priya takes over. If the start date moves, every step re-dates itself; nobody gets the old date.

What to notice. Each message answers the reader's own question in its first line, says what they won't have to do, and comes from a named person. The flow makes two decisions a busy team gets wrong: it chases once, then hands to a human; and it tells the COO before the COO has to ask. It's up to five messages and two phone calls, for one start. Kestrel starts nine reviews a year.

Templates

The trigger inventory (one row per trigger; the build's single list)

ID	Moment (chooser · step)	Who's waiting, for what	How long before it costs	System and event	Flow	Live?

The flow brief

FLOW	[name]
STARTS WHEN	[system · event · conditions]
NEVER FIRES	[when it must not run]
MOMENTS	[chooser · step, and what each is waiting for]
PEOPLE	[everyone it touches, and the owner]
ENDS WITH	[a person, a document or a confirmed fact]
OFF SWITCH	[who, and the automatic stop conditions]
MEASURE	[the stage it protects, and how you'll see it]

The step brief

STEP	[number and name]	WHEN	[day or event]
JOB	[stage · chooser · step · bias]		
TO / FROM	[a named person, to a named person]		
THEIR QUESTION	[answered in the first line]		
THE ONE THING	[what they need to do, or need to know]		
HANDOVER	[what sends it to a human, and to whom]		
GUARDRAIL	[what this step must never do]		

The flow log (one row per flow per month)

Flow	Fired	Completed without a person	Handed to a person	Misfired	Time from trigger to end	Stage it protects, this month

In the campaign plan

Campaigns create waiting. A benchmark launch brings in downloads; a board-season push brings in fit calls; a round table brings in owners who want a follow-up the same week. Automation's rule: **before a campaign starts, every response it invites has a flow, and every flow has an owner who knows the date.**

Kestrel campaign	What it invites	The flow that must be ready
The annual benchmark	Report downloads and questions about the figures	Enquiry response, with a reply that answers the reader's sector and problem; routed to a partner for anyone asking about a review
"Before you buy another line"	Fit-call requests from COOs	Enquiry response, with the calendar link and the equipment comparison

Kestrel campaign	What it invites	The flow that must be ready
Board season	Business case kit downloads	A one-line offer from a partner to go through the kit with their numbers, within a day
Owners' round table	Owners who want to talk	A personal follow-up from the speaker within 48 hours, set as a task, never sent automatically

The last row is deliberate. Some moments are too personal to automate: the flow's job there is to make sure the person remembers.

Running it

Who owns what. At Kestrel, the office manager owns the systems map and the flow log. Each flow has an owner who would notice if it stopped: the managing partner owns enquiry response and re-engagement; the head of delivery owns week zero, the review rhythm and savings confirmation; the finance lead owns the procurement pack. A flow without an owner is switched off.

When	What happens
Daily	Handovers read: every reply a flow passed to a person has an answer within a working day
Monthly	The flow log: what fired, what completed, what was handed over, what misfired
Quarterly	Each flow walked end to end as the buyer, on a test account; messages checked against the claims register and the voice
When a document expires	Insurance, safety, terms: the document store flags it 30 days before, and the procurement pack stops until it's replaced
When the motion changes	The CRM stages first, then every flow that reads them

The failure to plan for. Automation fails quietly. A broken flow sends nothing, and nobody notices for months. So every flow needs a heartbeat: a monthly count that would look wrong if it stopped. Kestrel's week-zero flow should fire about once a month; two months

at zero means either no reviews started, or the flow is broken. Someone should know which.

Common mistakes

The mistake	The stage it breaks	The bias it ignores	What it would do at Kestrel	Instead
Automating the talking, not the noticing	Believed	We follow others	A bot answering "will this work in our plant?"	Flows notice and route; people answer
Nurture for its own sake	Known, worn down	Good enough beats best	A weekly drip to every enquiry, until they unsubscribe	Send only what someone is waiting for
No handover	Bought	Losses loom larger	Procurement's question answered by the next scheduled email	Any reply but a yes goes to a named person within a day
Chasing until they give up	Preferred	Losses loom larger	Five reminders about the data request	Chase once; then a person calls
Sending stale documents	Believed	Familiar feels safe	Last year's insurance certificate, sent automatically	Documents from the store, with expiry dates that stop the flow
Messages written inside the tool	Believed	Familiar feels safe	A flow's email that contradicts the proposal's terms	Every message from the Email build, every claim from Credentials
Stages nobody keeps	All, unseen	A hard choice stalls	Week zero never fires because nobody moves the deal stage	Few stages, each one a moment, kept by the deal owner
No off switch, no owner	All	The planning fallacy	A flow still emailing a client who left in dispute	Every flow has an owner and stop conditions
Judging by volume	All, unseen	Good enough beats best	"2,000 automated emails sent this quarter"	Judge by the stage each flow protects

The bias to watch in the team: **automation bias**, the habit of trusting what a system did because a system did it. A flow that ran isn't a buyer who was served. Read the handovers.

The business check and the measures

Before any flow is switched on:

- **Revenue:** the P1 flows protect the deals already won and the proof for the next ones. If the team can build only two, build week zero and savings confirmation.
- **One truth:** every message's facts come from Credentials, every document from the store. A flow that states a price, a date or a term the proposal doesn't is switched off until it matches.
- **Consent:** contacts are messaged only for the deal or review they're part of. Newsletters need their own permission, kept by the Email build.
- **Capacity:** each handover needs a person within a day. If nobody can answer procurement's questions within a day, don't automate the request; fix the capacity first.
- **The systems can see it:** every trigger in the plan has a system and an event. Where none exists, the first job is a CRM stage or a field, not a flow.

The measures. Each flow is judged by the stage it protects, never by how much it sends.

Measure	What it tells you	Stage
Median time from enquiry to first reply	Whether buyers wait (Kestrel: 26 hours today)	Bought
Days from verbal yes to signed contract	Whether procurement waits (today: a week or more)	Bought
Reviews that start on the agreed date	Whether week zero works (today: 6 of 9)	Bought
Plant managers who confirm the checklist before the start	Whether the floor is ready	Answered
Closed reviews with savings confirmed within 120 days	Whether proof is being made (today: 2 of 9)	Believed
Handovers answered within a working day	Whether the people behind the flows are there	Preferred
Flows that misfired this month	Whether the system can be trusted	All

The Loop for this build

Each quarter the Loop (step 11) reads the flow log beside the Score's deal moments, and asks: **did the people who were waiting stop waiting, and did the moments after them hold?**

What the Loop finds	What it means	What changes
A flow completes, but the stage it protects doesn't move	The message is on time but doesn't answer the question	Rewrite the message with the Email build; check the reader's question
Most runs end in a handover	The flow is asking people to do something unclear, or the timing is wrong	Simplify the one thing it asks; move the date
Handovers aren't answered within a day	Capacity, not automation	Fix who owns the handovers before adding flows
A flow never fires	No such moments, or a stage nobody keeps	Check the CRM; if the moment is real, fix the stage
Savings confirmations rise	New proof for every build	Credentials signs each confirmed case; the website, the pitch and PR use it

What Kestrel's first year would look for. Reply times down from 26 hours to under one in the first month, because that flow is simple. Starts on the agreed date rising from six in nine to every one. And, by the year's end, savings confirmed for most closed reviews: the change no other build can make, because it turns every finished review into a case the finance directors at the next board will believe.

Shapes by business type

The moments where someone waits change with the business, and so do the flows that matter most:

Business type	Typical mode	The flows that usually matter most
B2B professional services	Committee or owner-led	Enquiry response, procurement pack, week zero, results confirmation

Business type	Typical mode	The flows that usually matter most
B2B software	Committee	Trial activation, usage-based nudges to a named rep, security review pack, renewal warnings
Considered consumer (finance, health, education)	Considered personal	Application progress, document requests, cooling-off reminders, a person on call for the hard questions
Ecommerce	Quick choice or habitual	Order and delivery updates, back-in-stock, abandoned basket (once), replenishment reminders
Local services	Owner-led or considered personal	Instant quote acknowledgement, booking confirmations, the day-before reminder, review requests after the job
Marketplaces	Varies on each side	Seller onboarding steps, first-sale prompts, buyer-seller response-time nudges, dispute routing

Feeling a bit overwhelmed?

You should be, a little. One firm: nine systems, 11 triggers, eight flows, each with a brief, a brief for every step, an owner, an off switch and a heartbeat, and a log that has to be read every month, because automation fails quietly. Most firms buy a tool, switch on a newsletter sequence, and wonder why procurement still waits a week for an insurance certificate.

RocMode does part of the heavy lifting:

- **From a signed brief to tasks, by channel.** When you sign a campaign brief in RocMode, it breaks the work into tasks for each channel: email, social, ads, landing page and pitch deck.
- **The moments behind every trigger.** Your choosers, their steps and their questions, mapped from what's known about your brand, so every flow you build starts from a moment someone is waiting on.

[Join the RocMode beta](#)

The Customers Playbook

What this build is for

Customers is the build for everything after the signature: making the work land, making staying the easy choice, and making the client the reason the next one calls.

Most marketing stops at the sale; most delivery teams don't think of themselves as marketing. Between them sits the stage most firms leave to chance: Kept. Yet it's where Kestrel's best evidence comes from, where its repeat work comes from, and, through 22 referrals out of 31 fit calls, where most of its new work comes from too.

The build covers four things:

- **Onboarding:** the handover from the pitch to the people who'll live with the work, so the promise survives the first week.
- **Loyalty:** results reviews and the programme decision, so staying is chosen, not drifted into.
- **Community:** a place where clients' teams help each other, and write the evidence the next buyer reads.
- **Referral:** making it easy for a satisfied client to tell another one.

The Direct category's two nodes live here, because the person who renews is the person who was served: **account management** (one owner for each client relationship) and **customer service** (what happens when something goes wrong mid-review).

One finding sets the tone. Byron Sharp and Anne Sharp studied a large loyalty programme and found substantial "excess loyalty" in only two of its six member brands, with gains that showed up among non-members too ([Sharp and Sharp](#)). Points don't keep clients. Keeping the promise does. This build is about the promise.

The footprint, and what it reads from the chain

Customers leads **4 cells and 4 moments**, the ones after the work where the relationship is decided:

Moment	Who	Their question
Delivery	Improvement lead	"Will we own the methods when they leave?"
Extend and refer	COO	"Did it pay back? Do we go into the programme?"
Extend and refer	Owner	"Was it worth it? Would I tell another owner?"

Moment	Who	Their question
Extend and refer	Plant manager	"Results are slipping. What do we do now?"

It supports four more: the plant manager's weekly needs during the review, the finance director's "are the savings in the accounts?", the improvement lead's "can we run the next review ourselves?", and the assistant's "is Kestrel any good?", which depends on reviews from clients this build keeps happy.

It's a small footprint with a large shadow. Every confirmed result this build produces becomes proof for the website, the pitch, PR and reviews and Being named. Every referral becomes a fit call that skipped the 93 moments where Kestrel scores zero.

What it reads, and never rewrites:

From	It takes
Credentials	What was promised: the scope, the six weeks, the way out
The pitch	Who's in the account, their questions and fears, what was said in the room
Sales Strategy	The motion's last three steps: the review, results and renewal, referral
Email and Automation	The delivery updates, the 90-day check and the programme decision, which this build acts on

Its slice of the Score, and its own audit

Kept isn't a separate check in [the Selectability Score](#): it's the same checks at the moments after the work starts. So this build's slice is the quality of Kestrel's delivery and follow-on moments.

Kestrel today, the 4 moments Customers leads	Answered	Believed	Preferred	Bought	Moment score
Improvement lead · delivery	60	65	70	70	66
COO · extend	65	70	71	60	68
Owner · extend	60	75	71	55	68
Plant manager · extend	55	60	65	55	60

These are Kestrel's best moments, because Kestrel is good at the work: 5 of the 7 clients who finished a review last year went on to more. Preferred here is that rate, 71%. The weaker scores are Bought ("what happens next" is left vague) and the plant manager's slipping results, which nobody comes back for.

The audit.

After the signature	What happens today	What it misses
Handover from the pitch	The partner who sold it briefs the lead engineer verbally	What each person in the account was promised, and feared
The improvement lead's methods	Left behind in the final report	Training their team to run the next review
Results review	The final presentation, then nothing	A review at 90 days, with the finance director, on confirmed figures
The programme decision	Raised by the client, if at all	A clear offer with its cost and return
Slipping results	Found out when the client calls, or doesn't	A check at six months on the plant floor
Clients' teams	Never meet each other	A place to compare notes with plants like theirs
Referrals	Arrive unasked and unthanked	An easy way to refer, and thanks when it happens

No one owns the relationship once the review ends. The partner moves to the next sale; the engineer to the next plant.

Step 1: the shape

The shape is **the customer lifecycle the category runs on**: the touches clients expect after they've bought, and what competitors offer.

Part	What clients expect	In the category (of 5 firms)	Kestrel today
Kick-off	A meeting that sets out who does what	5	Yes, informal

Part	What clients expect	In the category (of 5 firms)	Kestrel today
Handover between sales and delivery	The people on site know what was promised	3	No
Results review	A formal look at what changed	4	The final presentation only
An account owner	One person to call after the work	3	No
Training the client's team	Methods left behind, usable	2	No
A follow-up check	Someone comes back to see if it held	1	No
A client community	Peers to compare notes with	0	No
A referral route	A way to introduce others	1	No

Canon is the kick-off, the results review and an account owner: buyers who've used consultants before expect them. The opening is everything after: training, the follow-up check, the community and referral. Almost nobody does them, and they're where Kestrel's "we leave the methods with you" becomes something a client can see.

Step 2: coverage, one touch per moment after the signature

The coverage unit is **a touch**: one meeting, call, visit or invitation that answers a client's question at a moment after the signature. For each of the eight cells this build leads or supports, ask what the person needs, and whether only a person can give it.

Moment	Who	What they need	The touch	Only a person?
Delivery	Improvement lead	The methods, in their team's hands	Training sessions in weeks four to six, and a handbook	Yes
Delivery	Plant manager	Clear tasks each week	Monday tasks (Automation, Email)	No

Moment	Who	What they need	The touch	Only a person?
Extend	COO	Whether it paid back, and what's next	A results review at 90 days, with the finance director	Yes
Extend	Finance director	Savings in the accounts	The 90-day check (Email), then the results review	Partly
Extend	Owner	A reason to tell another owner	The results review; an invitation to the owners' round table; a forwardable paragraph	Partly
Extend	Plant manager	Help when results slip	A six-month visit to the floor	Yes
Extend	Improvement lead	Confidence to run the next review	A community of improvement leads from client plants; a check-in call before their first solo review	Partly
Checking	Assistant	Something to cite when asked "any good?"	Reviews, asked for by PR after confirmed results	No

Eight touches: four need a person, and three more need one in part. That's the point of the build: the moments that keep a client are the ones you can't automate.

Step 3: gaps, biases and jobs

Moment	Stage	Bias	The job
Handover, before week one	Answered	A hard choice stalls	A written handover from the pitch: each person's question, fear and what they were told, read by the lead engineer before the kick-off

Moment	Stage	Bias	The job
Improvement lead · delivery	Preferred	Losses loom larger	Their team runs the last two weeks' changes, with Kestrel beside them; a handbook they keep
COO · extend	Preferred	A hard choice stalls	A results review at 90 days: confirmed figures, what held and what didn't, and one clear offer for the programme or the next plant
Finance director · extend	Believed	We follow others	The finance director in the results review, presenting their own confirmed figure
Owner · extend	Preferred	We follow others	An invitation to the owners' round table; the forwardable paragraph after a "yes"
Plant manager · extend	Believed	Losses loom larger	A six-month floor visit: what's held, what's slipped, a short plan to recover it
Improvement lead · extend	Preferred	Familiar feels safe	A community of client improvement leads: two meetings a year, and a channel between them
Every client relationship	Bought	A hard choice stalls	One account owner, named at the kick-off, who stays when the engineers leave
When something goes wrong mid-review	Believed	Losses loom larger	Customer service, as a nod: a named person, a response within a day, and the issue in the next Friday update

The two rules: **the person who sold it hands over in writing**, so nothing promised is lost; and **someone owns every client after the work**, so the next conversation isn't left to chance.

Step 4: reconcile, the customer plan

Programme	What happens	Moments	Signals	Priority
Handover	Written handover from the pitch; kick-off with each person's question answered	Every delivery moment	Canon, coverage	P1
Account ownership	A named owner per client, from kick-off to the next piece of work	Every extend moment	Canon, coverage	P1
Results review	At 90 days, with the COO and finance director, on confirmed figures; one clear offer	COO, finance director · extend	Canon, coverage, strategy	P1
Methods handover	The improvement lead's team runs the last changes; a handbook they keep	Improvement lead · delivery	Coverage, strategy	P1
Six-month visit	Half a day on the floor; a short recovery plan if results slip	Plant manager · extend	Coverage	P2
Improvement leads' community	Two meetings a year, hosted at a client plant; a channel between them	Improvement lead · extend	Coverage, strategy	P2
Referral	The forwardable paragraph after a "yes"; the owners' round table; thanks for every referral	Owner · extend	Coverage	P2
Customer service	A named person; a response within a day; the issue reported in the Friday update	Delivery	Canon	P2

Eight programmes. Every moment has one. The four P1 programmes cost nothing but time and attention, and they protect Kestrel's best number: 5 of 7 clients going on to more.

Rejected, with reasons:

- **A loyalty scheme or retainer discount.** Kestrel's clients stay because the work paid back. A discount suggests it might not.
- **A client portal.** No moment asks for it. The Friday update and the account owner answer the questions a portal would.
- **A large annual client event.** Owners want to hear from one owner like them, not a conference. The round table does that.
- **Satisfaction surveys after every review.** The 90-day check asks the only question that matters, in figures.

The anatomy library

Touch	Parts, in order	Length
The handover note (pitch to delivery)	Each person: role, question, fear, what they were told · What was promised, word for word · What was said about the alternatives · Who decides the next piece of work	One page
The kick-off	Each person's question, answered · Who does what, week by week · How findings will be credited · Who to call when something goes wrong	One hour
The results review	The confirmed figure, presented by the client's finance director · What held, and what didn't · What the team now runs itself · One offer: the programme or the next plant, with cost and return	One hour, 90 days on
The six-month visit	Walk the lines with the plant manager · What's held, what's slipped · A three-item recovery plan · Whether help is needed	Half a day
The community meeting	One client's improvement lead presents what they did · Questions from peers · One method shared · The next host	Two hours, at a client plant
The referral moment	Thanks, for their figure · The forwardable paragraph · An easy no	A short note, after a "yes"

Four rules hold for every touch:

- **The client's people do the talking where they can.** The finance director presents the savings; the improvement lead presents the method.
- **Each touch ends with one decision**, never a vague "let's stay in touch".
- **Bad news travels first.** Slipping results are raised by Kestrel at the six-month visit, before the client has to.
- **Nothing is sold that wasn't asked about.** The results review makes one offer; the account owner never pitches between touches.

Briefs

The account brief is written once per client, at the handover, and kept by the account owner:

ACCOUNT	Hartwell Plastics, 260 staff, plastics and rubber
ACCOUNT OWNER	Jo Brady, partner (sold it; stays with it)
LEAD ENGINEER	Priya Raman
PROMISED	Fixed fee A\$60,000; six weeks; lines 2 and 4; way-out clause; the improvement team leads the changes
THE PEOPLE board	Lena Marsh, COO (champion): wants a result she can take to the board Tom Avery, finance director: "consultants overstate savings" Dave Okafor, plant manager: worried about his team's time Sam Ito, improvement lead: worried about a verdict on his work The owner: deciding on an A\$1.4m press
NEXT WORK TOUCHES	Line 1 and the tool room; the owner decides Kick-off 3 Nov; results review at 90 days; six-month visit; Sam into the improvement leads' community
MEASURE 90 days	Savings confirmed; the next piece of work; a reference after

And a touch brief for each meeting:

TOUCH	The results review
JOB	Preferred · COO · extend · a hard choice stalls
THEIR QUESTION	"Did it pay back? Do we go into the programme?"
WHO PRESENTS	Tom Avery: the confirmed figure; Sam Ito: what his team now runs
THE ONE OFFER GUARD	Line 1 and the tool room, fixed fee, with its likely return No offer if the 90-day check said "not yet"; fix first

The finished example: Hartwell's handover and results review, written in full

The same account from the Pitch, Automation and Email playbooks, picked up at the two moments this build owns. Hartwell, its people and its figures are invented.

1. The handover note, from Jo Brady (the partner who sold it) to Priya Raman (the lead engineer), the week before the kick-off.

Hartwell Plastics: what each person was told, and what they're worried about

Lena Marsh, COO. Champion. Needs a result she can take to the board by March, when the press decision comes back. Told: weekly updates every Friday, in a format she can forward. **Tom Avery, finance director.** Sceptical: "consultants always overstate savings". Told: nothing counted unless someone at Hartwell agrees it; a running tally from week three; a 90-day check. Don't round anything up. **Dave Okafor, plant manager.** Worried about his team's time. Told: two hours a day from him, one hour a shift from the line leads on 2 and 4, nobody taken off the line. Hold us to that. **Sam Ito, improvement lead.** Worried the review is a verdict on his team. Told: his team leads the changes and keeps the methods. He spoke to Brennock's improvement lead before we were chosen. Credit his team by name in every update. **The owner.** Didn't meet us. Decides on the A\$1.4m press. Told, through Lena: the review will show how much of the cost is in the existing lines first. **Promised, word for word:** fixed fee A\$60,000; six weeks; lines 2 and 4; "if the review doesn't find more than it costs, you don't go ahead with the changes, and you don't pay for them". **Next work, if it goes well:** line 1 and the tool room. I'll stay as account owner.

2. The results review, 90 days after the review closed, at Hartwell. Present: Lena, Tom, Dave, Sam, the owner; Jo and Priya for Kestrel.

Agenda (one hour)

1. The figure: Tom Avery (10 minutes). Tom presents what Hartwell's own accounts show: A\$198,000 a year of the A\$214,000 agreed at the end of the review, mostly scrap and resin. The gap is energy, where the tariff changed in January. **2. What held, and what didn't: Dave Okafor (15 minutes).** Warm-up scrap on line 2 has held. On line 4 it crept back on night shift; Dave's recovery plan is already running. **3. What the team now runs: Sam Ito (15 minutes).** Sam's team has run the changeover routine on line 1 without Kestrel. He shows their first results. **4. The one offer: Jo Brady (10 minutes).** Line 1 and the tool room, fixed fee A\$45,000, six weeks, the same way-out clause. On the evidence of lines 2 and 4, a likely return of A\$90,000 to A\$140,000 a year. And for the owner: the review found most of the capacity the press was meant to add. Whether the press is still needed is now a question the figures can answer. **5. The decision (10 minutes).** Go ahead, wait, or no. Any of the three is fine; we'd like to know which.

What happened next (invented, to complete the example). The owner deferred the press and asked for line 1. Tom agreed to be a reference for finance directors. Sam joined the improvement leads' community and hosted its next meeting. Dave's six-month visit is in the diary.

What to notice. The handover note stops the most common failure in services: the people on site not knowing what the person who sold it promised. The results review is presented by the client. Kestrel speaks only for the offer, and the offer comes after the client's own figure, including the part that didn't hold. That's why the owner believed the rest. Two documents, one account. Kestrel hands over nine a year.

Templates

The handover note

```
ACCOUNT          [name, size, sector]
FOR EACH PERSON  [name, role] · [champion / decider / blocker]
                  What they asked · What they fear · What they were told
PROMISED         [scope, fee, timing, the way out, word for word]
ALTERNATIVES     [what they were weighing, and what was said about it]
NEXT WORK        [what it could be; who decides]
ACCOUNT OWNER    [name]
```

The account plan (one row per client)

Client	Account owner	Review closed	90-day check	Results review	Six-month visit	Community	Reference?	Next work

The results review agenda

- | | | |
|-------------------------------|-----------------------------|--------|
| 1. The figure | [client's finance director] | 10 min |
| 2. What held, and what didn't | [plant manager] | 15 min |
| 3. What the team now runs | [improvement lead] | 15 min |
| 4. One offer | [account owner] | 10 min |
| 5. The decision | [the client] | 10 min |

The service log (one row per issue raised during a review)

Date	Client	Raised by	The issue	Owner	Responded within a day?	Resolved	In the Friday update?

In the campaign plan

Clients are the most believable part of almost every campaign. The rule: **a campaign asks clients to take part only after a confirmed result, and only in a way that helps them too.**

Kestrel campaign	What clients bring	What they get
The annual benchmark	Their data, with permission; a finance director's quote	Their plant's position against the benchmark, privately, before publication
Owners' round table	An owner who speaks	A room of peers, and the chance to host
Improvement leads' community	A host plant; a method to share	Peers who solve the same problems
"Before you buy another line"	A client who didn't buy the line, as a reference	Recognition for the decision, in their words
Results from the year's reviews	Confirmed figures	A results issue that credits their teams by name

No client appears in a campaign they haven't seen and approved, and no client is asked twice in a quarter.

Running it, with account management and customer service

Account management, as a nod. Every client has one account owner, named at the kick-off: usually the partner who sold the work, who stays with the client after the engineers leave. The account owner doesn't sell between touches. They make sure each touch happens, read every reply, and are the person the client calls. At Kestrel, no partner holds more than eight accounts.

Customer service, as a nod. When something goes wrong during a review (a missed shift, a finding the plant disputes, an engineer who's off sick), the client has one named person to tell, a response within a working day, and the issue reported in that week's Friday update, with what was done about it. Problems handled in the open build more trust than a review with none.

When	What happens	Who
Before every kick-off	The handover note, read by the lead engineer	Account owner, lead engineer
Weeks four to six of every review	The improvement lead's team runs the changes	Lead engineer
90 days after	The 90-day check (Email), then the results review	Account owner
Six months after	The floor visit	Lead engineer
Twice a year	The improvement leads' community meets	Head of delivery
Quarterly	The account plan: every client's next touch, and any client with none	Managing partner

The question for every account, every quarter: when did someone from Kestrel last help this client, and when will they next?

Common mistakes

The mistake	The stage it breaks	The bias it ignores	At Kestrel	Instead
Handing over by word of mouth	Answered	A hard choice stalls	The engineer learning on day three what the finance director was promised	A written handover, read before the kick-off
Ending at the final presentation	Preferred	Losses loom larger	No contact after week six unless the client called	A results review at 90 days; a visit at six months
Kestrel presenting the results	Believed	We follow others	Engineers presenting their own savings	The client's finance director presents the figure
Leaving the methods in a report	Preferred	Losses loom larger	A 60-page final report nobody used	The improvement lead's team runs the changes before Kestrel leaves
Hiding the slip	Believed	Losses loom larger	Results slipped on a night shift; the client found out first	Raise it at the six-month visit, with a plan
Selling at every touch	Preferred	Feeling first, reasons after	A proposal for more work in the final presentation	One offer, at the results review, after the client's figure
No owner after the work	Bought	A hard choice stalls	The partner moved on; the engineer moved on	An account owner, named at the kick-off
Loyalty schemes and discounts	Believed	Familiar feels safe	A "repeat client discount" offered once, and regretted	Keep the promise; nothing else keeps a client
Leaving referrals to luck	Preferred	We follow others	22 referrals, none thanked	A forwardable paragraph after a "yes"; thanks within a week

The bias to watch in the team: **the end of the project as the end**. The work feels finished when the engineers leave; for the client, the result is only starting to show. Every touch in this build lives in the gap between those two feelings.

The business check and the measures

- **Revenue:** 5 of 7 clients went on to more work last year. This build protects that number first, and makes the 22 referrals deliberate second.
- **One truth:** the handover note quotes the proposal and the contract word for word. What the client was promised is what delivery delivers.
- **The offer:** the results review's one offer uses the same fixed price and way out as the first review. Changing terms for repeat work is a partner's decision.
- **Capacity:** each client needs about four hours of the account owner's time after the review, and half a day of an engineer's at six months. That's roughly 80 hours a year across nine reviews.

Measure	What it tells you	Stage
Clients who go on to more work within 12 months	Whether Kept holds (today: 5 of 7)	Preferred
Results reviews held at 90 days, with the finance director	Whether the relationship has a next step	Preferred
Improvement leads who run a review without Kestrel within a year	Whether "we leave the methods with you" is true	Believed
Results still holding at six months	Whether the savings last	Believed
Referrals per client, and referrals thanked within a week	Whether clients bring the next one	Preferred
Issues raised mid-review answered within a day	Whether service is seen to work	Believed

The Loop for this build

Each quarter the Loop (step 11) reads the account plan beside the Score's extend moments, and asks: **did every client get their next touch, and did the result hold?**

What the Loop finds	What it means	What changes
Results reviews held, few go on to more	The offer doesn't follow from the results	Make the offer from the client's own figures, for the next obvious line or site
Results slip at six months	Changes didn't stick	Longer handover; the improvement lead's team runs more of the review
Improvement leads don't run reviews alone	The methods weren't really handed over	More of the last two weeks led by the client's team
Clients with no touch planned	No account owner, or an overloaded one	Reassign; cap accounts per partner
Referrals rise	The proof is travelling by word of mouth	Thank every one; ask referred clients how they heard, in their words

What Kestrel's first year would look for. Every client with a named account owner and a results review in the diary within the first quarter. The share going on to more work holding at five in seven or rising, now that someone asks. Improvement leads running reviews alone, which is the strongest proof Kestrel can offer an improvement lead at the comparing step. And referrals thanked, every one, so word of mouth stops being luck.

Shapes by business type

Business type	Typical mode	What keeps a customer, usually
B2B professional services	Committee or owner-led	A written handover, a results review on confirmed figures, an account owner, methods left behind, peers to learn from
B2B software	Committee	Onboarding to first value, adoption by role, a customer success owner, renewal conversations on usage, a user community
Considered consumer (finance, health, education)	Considered personal	Clear first steps, a person to call, annual reviews, no surprises at renewal

Business type	Typical mode	What keeps a customer, usually
Ecommerce	Quick choice or habitual	The product working as promised, easy returns, replenishment, a reason to come back that isn't only a discount
Local services	Owner-led or considered personal	The job done as quoted, a follow-up call, a reminder when it's due again, thanks for referrals
Marketplaces	Varies on each side	Sellers: early sales and fair fees; buyers: trust, protection and good matches

Feeling a bit overwhelmed?

You should be, a little. One firm: eight programmes after the signature, a handover note for every client, a results review presented by the client's own finance director, a six-month visit, a community of improvement leads, an account owner who stays, and a referral route that's never paid. Most firms put all their effort into winning the client, then hand them to delivery and hope.

RocMode starts where Kept starts:

- **The moments after the signature, on the same map.** RocMode maps your choosers and their moments from what's known about your brand, through delivery and the follow-on, so keeping a client is planned around the same people and questions as winning them.

[Join the RocMode beta](#)

Step 10: The Campaign Engine

What this step is for

A campaign bundles several builds around one moment or occasion, for a few weeks, so a buyer meets the same story from several directions at once.

The fourteen builds each do their own work all year: the website answers, the pitch persuades, social proves life. Left alone, they drift. The website says one thing in March, the ads another in August, and the pitch hasn't heard of either. A campaign is where they're made to move together, on a date that matters to the buyer.

The Campaign Engine is the step that plans those campaigns. It doesn't invent them. **Every campaign is generated from jobs the builds have already found:** the gaps in their footprints, the moments the GTM strategy chose to win first, and the occasions in the category's calendar. A campaign that can't name the moment it serves and the builds that carry it isn't a campaign. It's a promotion.

Three words, used precisely:

- **A moment** is a chooser at a journey step: "COO · trigger · flat output".
- **An occasion** is a date the category shares: budget season, the benchmark's release, an association's annual dinner.
- **A campaign** bundles builds around a moment, timed to an occasion where one exists.

This step answers one question: **what twelve months of campaigns would move the moments we chose, with the builds we have?**

What it reads from the chain

The Campaign Engine reads every step before it and rewrites none of them. It adds no audience and no claim.

From	It takes
Credentials (step 2)	Every claim a campaign may make, and the proof behind it
Brand strategy (step 3)	The choosers, the reason and the big idea
Brand identity (step 4)	The look, the voice and the distinctive assets every campaign wears
The map and the Score (step 5)	The moments, and where each is lost

From	It takes
GTM strategy (step 6)	Which moments to win first: for Kestrel, flat output and rising costs, at the trigger, framing and options steps, for the COO, the finance director and the assistant; every comparing and checking moment by sector; the finance director's price and way-out questions; and "before you buy another line"
Website, Marketing and Sales strategies (steps 7 to 9)	The pages, the build footprints and budget, and the sales assets each campaign can land on
The fourteen builds	Their jobs, their "in the campaign plan" tables, and their rules for what must be ready before a campaign starts

The last row matters most. Each build playbook ends its plan with a short table of the campaigns it serves and a rule for taking part. The Campaign Engine collects those tables and turns them into one calendar. Nothing appears in a campaign that a build hasn't already planned.

Two layers: always-on and campaigns

Most of what moves the Score never stops. The website answers every day. The pitch handles every deal. The moment ads run all year at a steady weight. Social proves life every week. That's **the always-on layer**, and it carries most of the 149 moments on Kestrel's map.

The campaign layer sits on top, for a few weeks at a time. It does three things the always-on layer can't:

- **It concentrates.** Several builds hit one moment in the same weeks, so a buyer meets the same story from a page, an ad, a letter and a post.
- **It uses a date.** Budget season, the benchmark's release, an association's dinner: occasions when a moment is more likely to be live.
- **It leaves something behind.** A new page, a piece of coverage, a reference, a review: things the always-on layer can use for the rest of the year.

The rule between them: **a campaign never pauses the always-on layer to pay for itself.** A campaign that switches off the moment ads to fund a burst is spending the memory to buy attention. At Kestrel, each build keeps about a fifth of its effort and spend back for campaigns, and the rest runs regardless.

Step 1: the occasions calendar

Start with the dates the category already keeps. They're set by buyers, boards and associations, not by Kestrel, and each one makes certain moments more likely to be live.

Occasion	When	The moments it makes live	Set by
Capital budgets set ("budget season")	Q3, for most of Kestrel's clients	COO trigger and options ("buy another line, or fix this one?"); finance director building the case	Clients' boards
Board season	Q3 and Q4	Finance director building the case: "will the numbers stand up?"	Clients' boards
The benchmark's release	Q2	Framing for the COO, the finance director and the assistant; finance director trigger	Kestrel (an occasion it made)
Association events	Through the year, one or two per association	Owner options and checking; COO checking by sector	Four sector associations
The trade show	Once a year	COO trigger	The industry
A manufacturing acquisition	Whenever one is announced	Investor and owner trigger: "margin within 18 months"	Investors and owners
Year-end shutdown	Late December to mid January	Few: plants close and inboxes go quiet. Campaigns pause; restart plans are made	The calendar
Results from the year's reviews	December	Checking, for everyone weighing Kestrel	Kestrel

Two of these Kestrel made itself: the benchmark and the results issue. Firms can create occasions, and the best ones are useful to buyers whether or not they ever hire you. A sale or an anniversary isn't one.

Step 2: from jobs to campaigns

Read every build's jobs and its "in the campaign plan" table. Group the jobs that serve the same moment. Where a group of jobs shares a moment and an occasion, it's a campaign. Where it shares a moment and no occasion, it's either always-on work or a campaign Kestrel times itself.

For Kestrel, the fourteen builds' tables reduce to **seven campaigns and two recurring routines**:

Campaign	The moment it serves	Stage it must move	Occasion	Lead build	Builds that carry it	GTM tier
The annual benchmark	Framing, for the COO, the finance director and the assistant; finance director trigger	Answered, Known	Its release, Q2	Content	Website, being named, PR, email, social, partners, customers, events, advertising, outbound, creators	Win first
"Before you buy another line"	COO · trigger and options · flat output; owners weighing equipment	Known, Preferred	Budget season, Q3	Advertising	Website, content, the pitch, PR, email, automation, partners, customers, social, outbound	Win first
Board season	Finance director · building the case	Believed, Bought	Q3 and Q4	The pitch	Website, content, PR, email, automation, advertising	Win first
Owners' round table	Owner · options and checking	Believed	The dinner's date	Events	Website, the pitch, PR, email, automation, partners, customers, outbound, creators, social	Win next
Investor 100 days	Investor · trigger and options	Known, Bought	Each acquisition	Partners and channels	Website, the pitch, PR, outbound	Win next
The improvement leads' series	Improvement lead · framing and checking	Believed	Kestrel's choice	Content	Email, customers, events, social	Win next

Campaign	The moment it serves	Stage it must move	Occasion	Lead build	Builds that carry it	GTM tier
Results from the year's reviews	Checking, for every chooser	Believed	December	Email	Customers, PR, website, being named	Win first
<i>Routine: a sector case each quarter</i>	Comparing, by sector	Believed	Quarterly	Being named	Website, content, the pitch	Win first
<i>Routine: review requests after results</i>	"Is Kestrel Advisory any good?"	Believed	After each 90-day check	PR and reviews	Customers, email	Win first

Three things to notice. **Every campaign has one lead build**, which owns the brief, the date and the result. **Every campaign serves a moment the GTM strategy chose**; the "hold" moments get none. And **most builds appear in most campaigns**, because the website, the pitch and email are in every buyer's path, and a campaign that forgets them leaks.

Step 3: the twelve months

A campaign can only run once the builds that carry it exist. So Kestrel's first year follows the build order in Part three: the website, being named and the pitch kit first; content, PR, email and customers next; paid, social, outbound and partners in the third quarter; events and automation last.

Month	Builds coming online	Campaigns	Routines
January	Website wave 1; being named (facts fixed everywhere); the pitch kit	None: there's nowhere yet for a campaign to land	—
February	—	None	—
March	—	None	First sector case
April	Content P1 guides; PR and reviews; email for delivery; customers	Benchmark: clients see their own plant's position, privately	Review requests begin
May	—	The annual benchmark: release (Content leads)	—
June	—	Benchmark: trade coverage, sector cuts, association summaries	Sector case

Month	Builds coming online	Campaigns	Routines
July	Website wave 2; advertising; social; creators; outbound; partners for investors	Moment ads start, always-on	—
August	—	"Before you buy another line" (Advertising leads), six weeks	—
September	—	"Before you buy" ends; Board season begins (the pitch leads); Investor 100 days opens with the first acquisition after partners are briefed	Sector case
October	Website wave 3; events; automation; the rest of content	Board season continues; improvement leads' series (three newsletter issues)	—
November	—	The first owners' round table (Events leads); board season ends	—
December	—	Results from the year's reviews (Email leads), early in the month	Sector case

From the second year, everything runs from January: two round tables (autumn and spring), the benchmark every Q2, and so on.

The capacity rule. No more than two campaigns launch in the same month, and never two led by the same build. Kestrel is a 60-person firm with one marketing lead. September is the tightest month; it works because "Before you buy" is ending as board season begins, and both lean on the same business case kit.

Step 4: the readiness gate

Each build playbook set one rule for taking part in a campaign. Together they make the gate a campaign passes before its first day. A campaign that fails any line waits; it doesn't launch with a hole in it.

Build	Its rule	The check
Website	A campaign lands on a page that already answers its moment	The landing page exists, and its first lines answer the moment's question
Being named	Every campaign names the questions it should move	The questions are listed, and the scan will check them a month or two later
Content	A campaign uses a core piece that already exists, or commissions one from the question bank	The core piece is published, or briefed from the bank
The pitch	Every campaign hands the pitch the asset its buyers will ask for next	The deal team has read the brief and confirmed the asset
Automation	Every response the campaign invites has a flow, and every flow has an owner who knows the date	Each response (download, enquiry, reply) has a flow and a named owner
Email	Emails go only to people for whom the question is live, never the whole list	The send list is defined by moment, not by "everyone"
PR and reviews	No campaign makes a claim without naming the outside voice that backs it	Each claim has its outside voice; any claim without one is cut
Customers	Clients take part only after a confirmed result, and in a way that helps them too	Every client named has approved the piece; nobody is asked twice in a quarter
Partners and channels	Partners get a version made for their clients	The partner version exists, and the partner has seen it
Advertising	The always-on layer is never paused to pay for a campaign	Campaign weight comes from the reserved fifth
Social, Creators, Outbound, Events	Real people, disclosed partnerships, named accounts with a reason, rooms that answer a moment	The build's own brief is signed

Build	Its rule	The check
Credentials	Every number is signed	Each figure in the campaign traces to a signed claim

PR's last line is the one to remember: "A campaign can wait for its outside voice. It can't make one up." Kestrel's "Investor 100 days" campaign claims "we start inside 100 days". Until an operating partner will say so, that claim isn't made, and the campaign runs without it.

Two briefs: the campaign brief and the asset brief

Most marketing leaks at the brief. The IPA's BetterBriefs survey asked more than 1,700 marketers and agency staff about briefs: eight in ten marketers said they wrote good ones, and one in ten agency staff agreed ([IPA](#)).

So every campaign has two kinds of brief.

The campaign brief is one page, written by the lead build and signed by whoever owns the Score. It carries:

- **One short idea:** the big idea, the line, and why it's true.
- **The moments and the stages it must move,** by name from the map.
- **The occasion and the dates.**
- **Each build's part,** and what it must have ready (the readiness gate).
- **A separate box of rules:** the claims allowed, the outside voices, what mustn't be said.
- **How it will be judged,** decided before it starts.

The asset brief carries five lines for each thing the campaign makes, whether a page, an ad, an email, a post or a slide:

1. **Moment:** whose, at which step, in their words.
2. **Stage:** the one it must move.
3. **Reason:** the reason, word for word.
4. **Proof:** the signed fact and its outside voice.
5. **Next step:** the one thing the buyer can do after.

The same five lines brief a person or an AI. Each build's own brief (the ad brief, the video brief, the piece brief, the creator brief) adds what that build needs, but starts from these five.

The finished example: "Before you buy another line", in full

The campaign brief, the asset list with five-line briefs, the readiness gate as it was checked, and the six weeks, as it runs in Kestrel's second year, once Hartwell's results have been confirmed at 90 days (Customers' rule: clients take part only after a confirmed result). Kestrel, its clients and their figures are invented.

1. The campaign brief

CAMPAIGN	"Before you buy another line"
LEAD BUILD	Advertising (owner: the marketing lead)
SIGNED BY	The managing partner, who owns the Score
THE IDEA	Find out what this line can do first.
THE LINE	"Before you buy another line, find out what this one can do."
WHY IT'S TRUE	In the 2025 benchmark, 6 of 38 plants had approved a new line before fixing the process on the old one. Brennock Packaging got from 41 to 47 a minute on the line it had. Hartwell Plastics commissioned a review of lines 2 and 4 while weighing a A\$1.4m press; A\$198,000 a year was confirmed at 90 days
MOMENTS	COO · trigger · flat output: "Output's flat and a key customer is complaining. What's going on?" COO · options: "Buy another line, or fix this one?" Owner · trigger · flat output and new line ramp-up Finance director · building the case (hand-off to board season)
STAGES	Known and Preferred: the COO remembers Kestrel when the capital request is written, and prefers a review first
OCCASION	Capital budgets set, Q3. Runs six weeks, mid-August to the end of September
RULES	Claims allowed: the benchmark finding; Brennock's 41 to 47; Hartwell's decision, in Hartwell's own words (approved). Never "you don't need new equipment": sometimes they do. Never name an equipment maker. Outside voices: Brennock's reference; a review saying the same; Hartwell's quote
JUDGED BY	Name searches and fit calls in the ad areas against the holdout; fit calls that mention equipment; the scan in November on the options questions; COO options moment at the next rescore

2. The assets, with their five-line briefs

Asset	Build	Moment	Stage	Reason	Proof	Next step
Landing page: "Kestrel vs new equipment"	Website	COO · options	Preferred	Savings found on the floor in six weeks, at a fixed price	Benchmark: 6 of 38; Brennock 41 to 47	The five-changeover test, then "When output is flat"
Ramp-up and automation ads, extra weight	Advertising	COO · trigger	Known	The line, word for word	Brennock, named	Search "Kestrel Advisory"
Newsletter: "What flat output really costs"	Email	COO and finance director, on the list	Answered	As above	The benchmark's cost of lost capacity	Download the business case kit
Business case kit, with the equipment comparison	The pitch	Finance director · building the case	Believed	A fixed-price review costs a fraction of a line	Hartwell: A\$60,000 review weighed against a A\$1.4m press	A partner goes through the kit with their numbers
"Is it the machine or us?" video, re-cut	Social	Plant manager · trigger · ramp-up	Known	In passing: "that's what six weeks on the floor is for"	A client plant manager, on camera	Check your own changeovers
Capacity check enclosure	Outbound	Owner · trigger, one-to-one firms with a flat-output or ramp-up signal	Known	As above	Benchmark finding	Reply to the partner
Equipment comparison, partner version	Partners	Integrator's and accountants' clients weighing capital spend	Preferred	As above	The same figures	The partner forwards it
Reference call: "the client who didn't buy the line"	PR, Customers	COO · checking	Believed	—	Brennock, in their words	Call Brennock's operations director

3. The readiness gate, as checked on 4 August

☒ ~~Landing page live; its first line answers "buy another line, or fix this one?"~~

- ✓ Questions named for the scan: options, by problem; the new equipment comparison
- ✓ Core pieces published: "When output is flat"; "What flat output really costs"
- ✓ The pitch has the kit with the equipment comparison; the deal team has read the brief
- ✓ Automation: enquiry response with the calendar link and the comparison; owner named
- ✓ Email list: COOs and finance directors only, not the whole list
- ✓ Outside voices: Brennock reference agreed; Hartwell's quote approved
- ✓ Customers: Brennock and Hartwell approved every piece that names them
- ✓ Partner version sent to the integrator and two accountants
- ✓ Ad weight from the reserved fifth; the always-on ads untouched

4. The six weeks

Week	What happens
1	Ads at extra weight; landing page promoted from the flat-output guide; outbound enclosures to 14 firms with a matching signal
2	Newsletter: "What flat output really costs"; partner version forwarded
3	Social: the ramp-up video, re-cut; Ray Hollins's series post on a new line that wouldn't reach rate (his choice of subject)
4	Second newsletter issue: the kit, for anyone who read the first; outbound follow-ups
5	The pitch: every finance director in a live deal gets the kit, dated for this budget round (board season's opening move)
6	Ads return to always-on weight; the brief's measures logged for the Loop

What to notice. The campaign makes one claim and backs it three ways: a number from the benchmark, a client who didn't buy, and a client who weighed a A\$1.4m press against a review. It never says buyers don't need equipment, because sometimes they do, and a COO would stop believing the rest. Every build had something ready before the first ad ran, so the COO who searched, downloaded or called met the same story at every turn. And it ends by handing the finance director to board season, the next campaign, rather than stopping cold.

Templates

1. The campaign brief

CAMPAIGN	[name]
LEAD BUILD	[build, and its owner]
SIGNED BY	[whoever owns the Score]
THE IDEA	[one short idea]
THE LINE	[the words buyers should keep]
WHY IT'S TRUE	[the signed facts behind it]
MOMENTS	[chooser · step · problem, and their question, from the map]
STAGES	[the stage each moment must move]
OCCASION	[the date, and the weeks it runs]
RULES	[claims allowed; outside voices; what mustn't be said]
JUDGED BY	[measures, decided now; the scan's questions; the rescore]

2. The asset list (one row per asset)

Asset	Build	Moment	Stage	Reason	Proof	Next step	Owner	Ready?

3. The readiness gate

- ☐ The landing page exists and answers the moment in its first lines
- ☐ The questions the scan will check are named
- ☐ The core piece is published, or briefed from the question bank
- ☐ The pitch has the asset buyers will ask for next, and the deal team has read the brief
- ☐ Every response the campaign invites has a flow and a named owner
- ☐ Email goes only to people for whom the question is live
- ☐ Every claim has its outside voice; any without one is cut
- ☐ Every client named has approved the piece, and none is asked twice this quarter
- ☐ Partners have a version made for their clients
- ☐ Extra weight comes from the reserved fifth; the always-on layer is untouched
- ☐ Every figure traces to a signed claim

4. The twelve-month calendar (one row per month)

Month	Occasions	Builds coming online	Campaigns launching	Campaigns running	Routines	Gate checked?

Running it

Who owns what. At Kestrel, the managing partner owns the Score and signs every campaign brief. Each campaign's lead build owns its brief, its dates and its result. The marketing lead keeps the calendar and runs the readiness gate. Every other build owner confirms their line of the gate before launch.

When	What happens
Yearly, in December	The next year's calendar: occasions, campaigns, lead builds, launch months; the capacity rule applied
Quarterly	The next quarter's campaign briefs written and signed; the previous quarter's campaigns reviewed in the Loop
Six weeks before launch	Asset list agreed; each build briefed with the five lines
Two weeks before launch	The readiness gate: every line checked by its build's owner; anything unready delays the launch, not the line
During	Weekly: the lead build checks delivery, and that every response is being handled
Two weeks after	Measures logged against the brief's "judged by"; the scan's questions noted for the next run

When the gate fails. Delay the campaign. A week late with a landing page beats on time without one. If an outside voice is missing, drop the claim it would have backed and run without it.

When a new occasion appears. An acquisition, a regulation, a customer crisis in the sector. Ask whether it makes a moment on the map more live. If it does, it may justify a campaign; if it doesn't, it doesn't, however much attention it's getting.

Common mistakes

The mistake	The stage it breaks	The bias it ignores	What it looks like	Instead
A campaign without a moment	All	Good enough beats best	"Let's do something for our anniversary"	Name the chooser, the step and their question, or don't run it
A campaign that lands nowhere	Answered	A hard choice stalls	Ads to the home page	The gate: a page that answers the moment, live before launch
Pausing the always-on layer to fund it	Known	Losses loom larger	The moment ads switched off for a six-week burst	Campaign weight from a reserved fifth
Too many at once	All	A hard choice stalls	Four launches in September, one marketing lead	Two launches a month at most, never two from one build
A claim with no outside voice	Believed	We follow others	"We start inside 100 days", said only by Kestrel	Wait for the voice, or cut the claim
The whole list, every time	Answered	Good enough beats best	Every contact gets every campaign	Only people for whom the question is live
The pitch hears last	Bought	A hard choice stalls	Fit calls the kit can't answer	The deal team reads the brief before launch
Judged afterwards, by whatever went up	All, unseen	Feeling first, reasons after	"Impressions were great"	"Judged by" written into the brief before it starts
Stopping cold	Bought	A hard choice stalls	The campaign ends; the buyers it found hear nothing	Hand them to the next campaign or the always-on layer

The bias to watch in the team: **action bias**. Doing something feels better than waiting, so a quiet month gets filled with a campaign. But a campaign without a moment, or launched before its page and its proof are ready, uses up the buyers it reaches. Sometimes the right campaign this month is none.

Measures, and the hand-off to the Loop

Every campaign is judged on three clocks, and the brief says which measures belong to each before it starts.

Clock	What it asks	For "Before you buy another line"
During (diagnostic)	Did it run, and is every response being handled?	Ads delivered; enquiries answered within a day; kit downloads followed up by a partner
Weeks after (leading)	Did the moment's stage move for the buyers it reached?	Name searches and fit calls in ad areas against holdout areas; fit calls that mention equipment; outbound replies
At the next scan and rescore (the Score)	Did the questions it named move, and did the moment's score?	The options questions, by problem, and the new-equipment comparison, in November's scan; the COO options moment at the next rescore

Two rules from the builds apply to every campaign:

- **Measured on the page as well as in the channel** (Website): did the moment's stage move for the buyers the campaign brought? An ad that brings clicks to a page that doesn't answer has bought traffic, not selection.
- **Something left behind** (Being named): a campaign that leaves nothing for an assistant to read has reached only the people who saw it that week.

At the end of each quarter, every campaign's brief, with its measures filled in, goes to the Loop (step 11). The Loop decides what changes: a campaign repeated, reshaped or dropped, and any new moment it found added to the map.

Shapes by business type

Business type	Typical mode	The campaign rhythm that usually fits
B2B professional services	Committee or owner-led	Four to seven campaigns a year, built on the clients' calendar (budgets, boards) and one occasion the firm creates, such as a benchmark
B2B software	Committee	Campaigns around the buying group's calendar, product releases that change a moment, and the category's main event; always-on carries most of the year
Considered consumer (finance, health, education)	Considered personal	Campaigns timed to the life events and seasons that start the search: tax time, enrolment, renewal dates
Ecommerce	Quick choice or habitual	Many short campaigns on seasons and gifting dates, with fixed assets; the always-on layer still matters most
Local services	Owner-led or considered personal	A few campaigns a year on local occasions and seasons; one done well beats four done thinly
Marketplaces	Varies on each side	Separate rhythms for each side: supply campaigns before the demand season, demand campaigns during it

Feeling a bit overwhelmed?

You should be, a little. One firm: an occasions calendar, fourteen builds' campaign tables read and grouped, seven campaigns and two routines, a twelve-month calendar that follows the build order, a capacity rule, a readiness gate with a dozen lines, a one-page

brief and a five-line brief for every asset, and measures on three clocks. Most firms pick a few dates, make some ads, and judge them by whatever went up.

RocMode does part of the heavy lifting:

Campaign concepts from your strategy. RocMode's marketing strategy proposes campaign concepts, each with a big idea, a line and why it's true, the occasion it's timed to, the choosers and segments it serves, its channels, budget, objective and measures.

From concept to campaign card. Its Campaign Centre turns each concept into a draft campaign card with a name, objective, budget, schedule, measures and audience, for you to review.

From signed brief to tasks. When you open and sign a campaign's brief, RocMode breaks the work into tasks by channel: email, social, ads, landing page and pitch deck.

[Join the RocMode beta](#)

Step 11: The Loop

What this step is for

A score is only worth taking if it changes what you do next. The Loop is where it does.

Every step so far has made a plan: the moments to win, the builds that lead them, the campaigns that bundle them. The Loop asks whether any of it worked, and changes the plan when it didn't. It answers one question: **what changed, and what do we change?**

The book sets out the whole method as six steps: measure, diagnose, write down the truth, prove it, publish it, measure again. Steps 1 to 10 of the Brand Journey are the middle four, done properly. This step is the first and the last: measuring, the same way every time, and recording what moved against what you expected to move.

Most teams run the first step and stop. Or they jump from a disappointing number straight to making more content. The Loop exists to stop both. It ties every fix to a prediction, and every prediction to a result, so the next quarter's plan is built on what actually happened.

What it reads from the chain

The Loop reads everything, and is the only step allowed to send changes back up the chain.

From	It takes	What it can send back
The map and the Score (step 5)	The baseline: every moment's score, its range, and where it's lost	A new moment; a moment split or merged; a question reworded (which starts a new baseline)
GTM strategy (step 6)	Which moments to win first, next, or hold	A moment moved between tiers, with the evidence
Credentials (step 2)	Every claim and its proof	A claim that buyers or assistants keep getting wrong, to be fixed at the source
The fourteen builds	Each build's "Loop for this build" table and first-year targets	A job dropped, added or reshaped
The Campaign Engine (step 10)	Every campaign brief, with its measures filled in	A campaign repeated, reshaped or dropped

From	It takes	What it can send back
The deal record	Fit calls, proposals, wins and losses, and why	The mode weights and the moment weights, checked against real sales
New clients	What was happening when they first thought of calling, and where they'd heard of Kestrel	The moments buyers actually name

Three clocks: monthly, quarterly, yearly

Clock	What happens	Who	How long
Monthly	The machine scan: the same 101 questions, the same six assistants, three runs each. Corrections where an assistant misquotes a fact. New clients' answers logged	Being named's owner	Half a day to read
Quarterly	The review: the full rescore (machines, people and places, the deal record); each build's Loop table; each campaign's measures; the levers re-ranked; the next quarter's top three fixes chosen and predicted	The Score's owner, with every build owner	Half a day, prepared over a week
Yearly	The map refreshed: moments added, split or retired; the question set reviewed (a new baseline if it changes); weights checked against the deal record; the next year's campaign calendar	The Score's owner and the managing partner	Two days

The book draws the quarter as four boxes in a circle: **ask** (week 1: re-run the same questions), **fix** (weeks 2 to 11: ship the top three fixes), **ask again** (week 12: the same sample, unchanged), and **record** (predicted against measured; drop what did nothing). Record is the step most teams skip, and the one this step exists for.

Step 1: measure again, unchanged

Run the same three instruments as the baseline in step 5, the same way.

Instrument	Moments	What's asked	Unchanged means
Asked of machines	61	101 questions, six assistants, three runs each: 1,818 answers	The same questions, word for word; the same assistants; the same number of runs
People and places	38	114 searches and 40 places (directories, associations, review platforms, trade titles)	The same searches and places, checked the same way
Inside a deal	50	Read from the deal record: fit calls, proposals, wins and losses, and what buyers said	The same moments, read from the same fields

Three rules keep the comparison fair:

Change nothing in the sample mid-year. A reworded question, a new assistant or a new place starts a new baseline. Make those changes once a year, record them, and rescore the old quarter on the new set if you can.

Read ranges, not just numbers. Kestrel's baseline is 18, with a range of 12 to 27. A quarter that reads 21 hasn't proved anything on its own. A quarter whose range sits wholly above the old score of 18 is starting to.

Read moments, not just the total. The total moves slowly, because most of Kestrel's moments are out in the market and change only when the market notices. Individual moments move faster, and say which fix worked.

Step 2: diagnose, and rank the levers

For every moment that moved, or didn't, find the stage where it's lost and why. Then split the total loss, as step 5 did, and see whether its shape has changed.

At Kestrel's baseline, the Score lost 82 points, and three quarters of them at the gates:

Stage	Points lost at baseline	What it meant
Found	46	Buyers and assistants who never come across Kestrel at their moment
Known	16	Choosers who've never heard of it

Stage	Points lost at baseline	What it meant
Believed	5.8	Claims with no outside voice
Answered	5.1	Questions no page answers
Preferred	5.0	No reason a buyer could repeat
Bought	3.4	Price "on application"; a contact form as the next step

The levers come from that split: each fix names its moments and stages, and its predicted lift is worked out by rerunning the Score with the fix in place. At the baseline:

Fix	Predicted lift (range)
Website: the pages that answer the trigger, framing and options moments	+4.2 (2 to 6)
The finance director's business case kit	+1.5
The proposal: price, scope and the way out, side by side	+0.9
References within a week	+0.9
Procurement's pack	+0.3
All five together	to about 26 (22 to 28)

Each quarter, the levers are re-ranked from the new loss split. A fix that's landed drops off the list; the gap it closed shifts the shares, and the next biggest lever moves up.

Step 3: predict, fix, record

Pick the top three fixes for the quarter, no more. For each, write down before you start:

- **The moments it should move**, by name.
- **The stage it should move them at.**
- **The predicted lift**, as a range, from rerunning the Score with the fix in place.
- **When the effect should show**: this quarter's rescore, or later (being named takes a scan or two; Known can take a year).
- **The build that owns it**, and the date it ships.

At the end of the quarter, put the measured result beside the prediction. Four things can happen:

Measured against predicted	What it usually means	What to do
Within or above the range	The fix worked, and the model of why was right	Keep it; look for the same pattern in other moments
Below the range, but moving	Partly right, or too early	Keep it; check again next quarter before adding to it
Nothing moved	Wrong diagnosis, or it never shipped as planned	Check it shipped. If it did, drop it and re-diagnose
Something else moved instead	A side effect, or a campaign, a launch or the market	Note what else changed that quarter before crediting anything

The record is the asset. After a year, Kestrel knows which kinds of fix move which kinds of moment in its own market: something no benchmark or book can tell it. A fix that did nothing is dropped and recorded, not quietly continued.

Step 4: feed the map

The map in step 5 was the best guess at the time. The builds and campaigns test it every week, and they find things it missed. The Loop is where those findings go back in.

Where the finding comes from	What it looks like	What the map does
New clients' answers	Buyers keep naming a situation the map doesn't have	A new moment, as a hypothesis, until interviews confirm it
Outbound replies and fit calls	Owners keep asking the same question at the trigger	A new question for an existing moment, or a new moment
The scan	Assistants answer a question with a rival the map didn't list	The rival added to the alternatives; the moment's question checked
Campaigns	A campaign needed a page the sitemap didn't have (the Website rule)	The moment it found added to the map, or the page given an end date
The deal record	Moments that never appear in won deals, however well they score	Weights checked; the moment may matter less than assumed

Where the finding comes from	What it looks like	What the map does
GTM	A "hold" moment confirmed by interviews	Moved to "win next", with a build to lead it

Every change is logged with its source and its date, and the GTM strategy re-reads the map. A new moment means a new cell, a lead build and, at the next yearly reset, new questions in the sample.

The finished example: Kestrel's third-quarter review, in full

The review note the managing partner circulated at the end of September in Kestrel's first year, after the half-day review with every build owner. The firm, its clients and every figure are invented.

Kestrel Advisory: the Loop, Q3 review

1. The Score

	Baseline (January)	End Q1	End Q2	End Q3
Score	18 (12 to 27)	19 (13 to 27)	21 (15 to 28)	24 (19 to 30)
Points lost at Found	46	44	41	38
Points lost at Known	16	16	16	16

The range now sits wholly above the January score of 18, for the first time. Known hasn't moved, as expected: it's read again at the year end, and most of the work that moves it (ads, social, outbound) only started in July.

2. This quarter's three fixes, predicted against measured

Fix	Owner	Moments	Predicted	Measured	Verdict
Website wave 2: options and comparing pages, by sector	Website	COO options; comparing, by sector	+2.0 (1 to 3)	+1.6	Within the range. Keep the pattern: answer in the first lines

Fix	Owner	Moments	Predicted	Measured	Verdict
The finance director's business case kit	The pitch	Finance director building the case	+1.5 (0.8 to 2.2)	+0.4	Below. The board paper section shipped on 24 September, too late for this rescore. Re-measure in Q4
The proposal: price, scope and the way out, side by side	The pitch	Finance director comparing; owner buying	+0.9 (0.4 to 1.4)	+1.1	Within. Keep

3. Moments worth noting

- **COO · options · flat output: 0 to 6.** Assistants now name Kestrel in 2 of 18 answers to "buy another line, or fix this one?", citing the "Kestrel vs new equipment" page and the benchmark's coverage.
- **Finance director · trigger: 0 to 2.** The benchmark is being quoted.
- **The mining mix-up: gone** from every answer, as being named predicted, within three scans.
- **Plant manager · framing: still 0.** Ray Hollins's floor-day post went out on 19 September; we don't expect it in the Score before Q1. Two plant managers in live deals have mentioned it already.

4. Campaigns

Campaign	Judged by (from its brief)	Result	Decision
The annual benchmark (May)	Coverage in the three cited trade titles; quoted by assistants on framing questions	Covered by 2 of 3 titles; quoted in 2 of the 7 framing questions	Repeat next May; pitch the third title earlier
"Before you buy another line" (Aug-Sep)	Name searches and fit calls, ad areas against holdout; fit calls mentioning equipment; outbound replies	Name searches up about 30% in ad areas, about 8% in holdout areas; 5 fit calls mentioning equipment against 1; 3 replies from 14 letters. Small numbers: a sketch, not proof	Repeat next year, in the second year's form, once Hartwell's results are confirmed

Campaign	Judged by (from its brief)	Result	Decision
Retargeting holdout	Any difference, on against off	None measurable	Stays off

5. Changes to the map

- **New question, held:** four owners replying to outbound letters asked whether a government grant could pay for a review. Added as a hypothesis to owner options: "Can a grant pay for it?" Held until three interviews confirm it.
- **New alternative:** for the ramp-up question, two assistants named a large firm's "lean in 90 days" programme. Added to the alternatives; the comparison page goes on wave 3's list.
- **Dropped from the levers:** procurement's pack. It's worth 0.3 at most, and nobody in a deal has asked for it. Recorded, not continued.

6. Q4's three fixes

Fix	Owner	Moments	Predicted	Shows by
The board paper, used in every live deal during board season	The pitch	Finance director building the case	+0.8 (0.3 to 1.3)	End Q4
A food and drink sector case on the website and the fact sheet	Being named	Comparing, food and drink	+0.6 (0.2 to 1.0)	End Q4 or Q1
The owners' page, and the first round table	Events	Owner options and checking	+0.5 (0.1 to 1.0)	Q1

What to notice. Every fix had a prediction before it started, and the note puts the measured result beside it, including the one that fell short and why. The total is read with its range, and the review says when the range first cleared the old score. The campaign results are called a sketch where the numbers are small. The map gained a question and an alternative from the builds' own findings, and lost a lever that wasn't worth its effort. And the next quarter's three fixes are chosen from the new loss split, each with a range and a date it should show.

Templates

1. The quarterly review agenda (half a day; papers out a week before)

1. The Score: total and range, against the last four quarters; the loss split by stage.
2. Last quarter's three fixes: predicted against measured, with a verdict each.
3. Moments that moved, or should have and didn't.
4. Each build's Loop table: one line per build, from its owner.
5. Campaigns: each brief's "judged by", with results.
6. Changes to the map: new moments, questions, alternatives; moments moved between tiers.
7. Next quarter's three fixes: moments, stage, predicted range, owner, date it ships, when it should show.

2. The fix log (one row per fix, kept for ever)

Quarter	Fix	Owner	Moments	Stage	Predicted (range)	Shipped on	Measured	Verdict	What we learned

3. The map change log

Date	Change	Source	Evidence	Status (hypothesis / confirmed)	Effect on the sample

4. The new-client question (asked at every first meeting, and logged)

"What was happening when you first thought of talking to someone about this?"

"Where had you come across Kestrel before you got in touch?" (Offer: a search, an AI assistant, a colleague or client, LinkedIn or YouTube, an event, a letter, an ad, somewhere else.)

Running it

Who owns what. One person owns the Score and the list of fixes, across marketing and sales, with the standing to say no to work that moves no moment. At Kestrel it's the managing partner. They don't own the builds; they decide which fix goes first and which build does it. Each build owner brings their Loop table to the review. Being named's owner runs the monthly scan. The marketing lead keeps the fix log and the map change log.

When	What happens
Monthly	The scan; corrections; new clients' answers logged
Week 1 of each quarter	The rescore, on the unchanged sample
Week 2	Papers out: the Score, the loss split, each build's Loop line, each campaign's results
Week 3	The review: half a day, every build owner; the next three fixes chosen and predicted
Weeks 3 to 12	The fixes ship, on the dates their owners gave
December	The yearly reset: the map refreshed, the sample reviewed, weights checked against the deal record, next year's campaign calendar

Three fixes, not ten. A quarter with ten fixes can't tell which one worked. Three, each with a prediction, can.

Common mistakes

The mistake	What it breaks	The bias it ignores	What it looks like	Instead
Measuring and stopping	The whole loop	Good enough beats best	A dashboard everyone looks at and nobody acts on	Every review ends with three fixes, predicted
Changing the questions mid-year	The comparison	Familiar feels safe	"The new questions are better", and the trend vanishes	Change the sample once a year, and record it
Reading the total, not the range	The decision	Feeling first, reasons after	Celebrating 18 to 21	Read the range; read the moments
No prediction	The learning	Losses loom larger	Fixes chosen, results "interesting"	A range written down before the fix starts

The mistake	What it breaks	The bias it ignores	What it looks like	Instead
Ten fixes at once	The attribution	A hard choice stalls	Everything changed; nobody knows what worked	Three a quarter
Crediting the fix for everything	The truth	We follow others	A launch and a fix in the same quarter, one credited	Note what else changed before crediting anything
More content as the answer to every gap	The diagnosis	Good enough beats best	A disappointing score, so more blog posts	Find the stage and the cause first
Keeping a fix that did nothing	The budget	Losses loom larger	"We've put so much into it"	Drop it, and record why

The bias to watch in the team: **the sunk cost fallacy**. The more a team has put into a fix, the harder it is to drop when the rescore says it did nothing. The fix log exists so that dropping a fix is a recorded decision, not an admission.

What the Loop can't tell you

The Score looks forwards, at the choice. The measures a finance director already trusts (attribution, marketing mix models and experiments) look backwards, at sales. The Loop uses both: the Score says which fix to try, and an experiment such as Advertising's regional holdout says whether it paid.

Be clear about the limits:

- **The weights are judgement.** The mode weights and the moment weights need calibrating against real sales, mode by mode. The yearly reset checks them against the deal record, and that's a start, not a proof.
- **Nobody has yet shown that the Score moves with market share.** Until that's tested, read it as an index of your odds, not a forecast of sales.
- **Some numbers aren't independent.** Found and Preferred come from the same sample of answers, so the real margin is wider than the arithmetic says.
- **It can't see a sale.** Only your own data can. Where you have it, it outranks everything above.

And one measure costs nothing: ask every new client how they found you, and include "an AI assistant" as an answer. Over a year, those answers will tell Kestrel more about which moments matter than any model.

Shapes by business type

Business type	Typical mode	How the Loop usually runs
B2B professional services	Committee or owner-led	Quarterly reviews; a slow total and faster moments; the deal record read deal by deal, because there are few of them
B2B software	Committee	Quarterly reviews, with product usage and pipeline data beside the Score; holdouts and, with enough years, a mix model
Considered consumer (finance, health, education)	Considered personal	Quarterly, timed around the seasons that start the search; the new-customer question matters most
Ecommerce	Quick choice or habitual	Monthly on the moments, quarterly on the total; sales data and holdouts carry more weight than for slower categories
Local services	Owner-led or considered personal	Five moments by hand each quarter, using the book's self-score worksheet; an afternoon to a day
Marketplaces	Varies on each side	A loop for each side, read together: a fix on one side often shows on the other

Feeling a bit overwhelmed?

You should be, a little. One firm: 1,818 answers a month from the machines, 114 searches and 40 places, the deal record read moment by moment, a loss split by stage, levers re-ranked every quarter, three fixes predicted and recorded, every build's Loop table and

every campaign's measures, a map that changes as the evidence comes in, and a yearly reset. Most firms look at a dashboard, feel good or bad about it, and make more content.

RocMode does part of the heavy lifting:

The measuring, the same way each time. RocMode asks your buyers' questions, without your name in them, of several AI assistants, several times each, and records who's named, which rivals appear and which sources are cited. Questions that name your brand don't count towards being found.

A range, not just a number. Each moment's score comes with its range, so you can tell a real move from noise.

What kind of moment, and the biggest lever. Each moment is sorted into won, hidden, borrowed or empty, and RocMode names the fix worth the most points, with the points it should add.

Writing down predictions, choosing the three fixes and running the review are still yours.

[Join the RocMode beta](#)