PinpointProof

A PINPOINTPROOF CASE STUDY ON ITS OWN OPERATIONS

Deciding where AI belongs

How we mapped our own marketing, scored every step, and built a bespoke platform in days, not quarters.

7 processes mapped | Every step scored | Built in 7 working days | The right steps kept human

EXECUTIVE SUMMARY

We ran our method on ourselves first

PinpointProof is a new startup building process intelligence and AI opportunity mapping software for regulated financial firms. Like every early-stage firm, it faced a resourcing question with a familiar shape: marketing needs doing every day, across many channels, and the founder's time is worth most where only the founder can spend it, leading client engagements, running discovery sprints, and building the product.

So we did what we ask our clients to do. We mapped our own marketing processes in PinpointProof, scored every step for AI and automation potential, decided deliberately what to automate, what to AI-assist under approval, and what to keep human, and then built exactly that as a bespoke platform.

The result is Beacon: a marketing operating platform tailored to one firm's processes, built in 7 working days with AI tools, for a fraction of the cost of the software it replaces.

The outcome, in numbers we can stand behind

7 processes
High-level marketing processes mapped; the performance process alone holds 11 nodes in detail

7 working days
From first design mockup to a running platform, feature by approved feature

28 features
Across intelligence, content, publishing, engagement, performance and attribution

£125 per month
AI tooling and platform combined, against roughly £3,000 per month for a comparable SaaS marketing suite

40 to **4** hours
A full-time marketer's week of work now runs in around 4 hours of founder time

Kept human, on purpose
Approving every draft, the first hour after posting, and real conversations

The method transfers directly. Swap the word marketing for client onboarding, fund launches, breach reporting or board MI, and the questions are identical: where does AI belong in this process, what must stay human, and how do you evidence the decision for a board or a regulator?

Why this document exists. Firms are told to adopt AI and given no method for deciding where. This is the method, demonstrated on the makers, with the screenshots to prove it.

CONTENTS

What is in this case study

Reading time is about 25 minutes. Sections 3 to 6 mirror the stages of a PinpointProof engagement, so the document doubles as a walkthrough of the method itself.

1 Everyone says automate. Choosing where is the hard part.

Every regulated firm now has an AI mandate, and most have a budget to match. What almost none have is a method for deciding where the AI should actually go. The result tends to be one of two failure modes. Paralysis: working groups, no baseline, nothing shipped. Or its opposite: automating whatever a vendor demo made look easy, which is rarely the step where automation was safe or valuable.

We had the same problem, in miniature. A founder-led startup, eight marketing channels, competitor and regulatory news to monitor, content to draft, publish, measure and learn from, and a founder whose highest-value hours belong to clients and product.

Run manually, this is a full-time job. Monitoring competitors and regulatory news, drafting for eight channels in a consistent voice, publishing on schedule, staying in the comments, working the inbox, capturing performance numbers and actually learning from them adds up to roughly 40 hours a week of a competent marketer's time.

Figure 1. The manual week: roughly 40 hours across the seven marketing processes, every week.

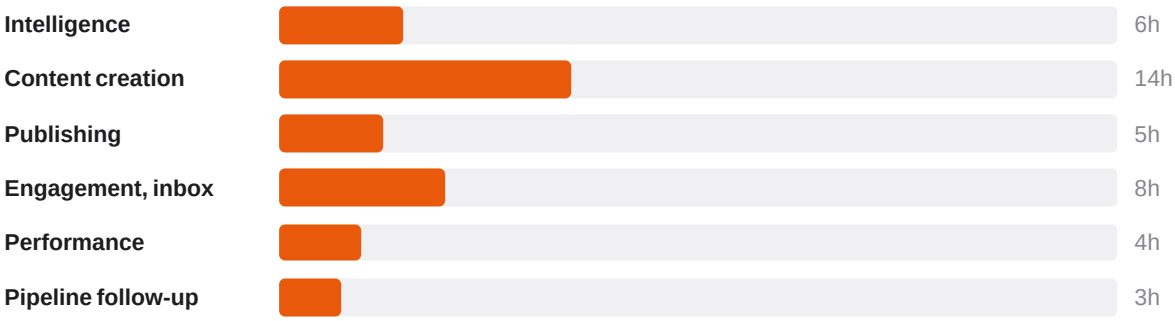

Indicative split of a 40-hour marketing week for a firm of this shape, before any automation.

Why the conventional answers fell short

A marketing SaaS suite at roughly £3,000 per month would have imposed someone else's processes on us and left the judgement work undone. Handing everything to AI would have put a language model in charge of the two things a young firm in a regulated market protects most carefully: its voice and its credibility. Somewhere in our marketing were steps that should never be automated. We needed to find out which ones.

That is precisely the question PinpointProof exists to answer. So we pointed it at ourselves.

2 The method in one page

The approach is the same whether the process is marketing, client onboarding or breach reporting.

Figure 2. The PinpointProof method.

01	**Map**	How the work actually happens, by sitting with the work itself. No event logs, no desktop surveillance, no transformation programme.
02	**Score**	Every step, for judgement required, repeatability, impact when it goes wrong, and whether there is data to verify against.
03	**Decide**	Deliberately: automate, *AI*-assist with human approval, capture with assistance, or keep manual on purpose. Every decision recorded.
04	**Prove**	With a designed build or test: small enough to be safe, real enough to be evidence a board can rely on.
05	**Learn**	Measure what shipped, feed the evidence back, and let the decisions improve rather than fossilise.

The rest of this document is that method applied to ourselves, with the screenshots to show it: the map (section 3), the scores and decisions (section 4), the operating model they produced (section 5), the governed build (section 6), the resulting platform (section 7), and the real numbers (section 9).

3 Stage one: mapping the marketing process

We started where every engagement starts: with a map of how the work actually happens. The real sequence of steps, decisions, handoffs and waits, drawn with the person who does the work rather than from a job description.

The marketing function decomposed into seven high-level processes: intelligence and research, content creation, publishing and engagement, performance and learning, and the pipeline processes that turn attention into conversations. Each expanded into a detailed map. The performance process alone, the one that answers whether any of this worked, held 11 distinct nodes once it was drawn in full.

Figure 3. The full marketing process map in PinpointProof.

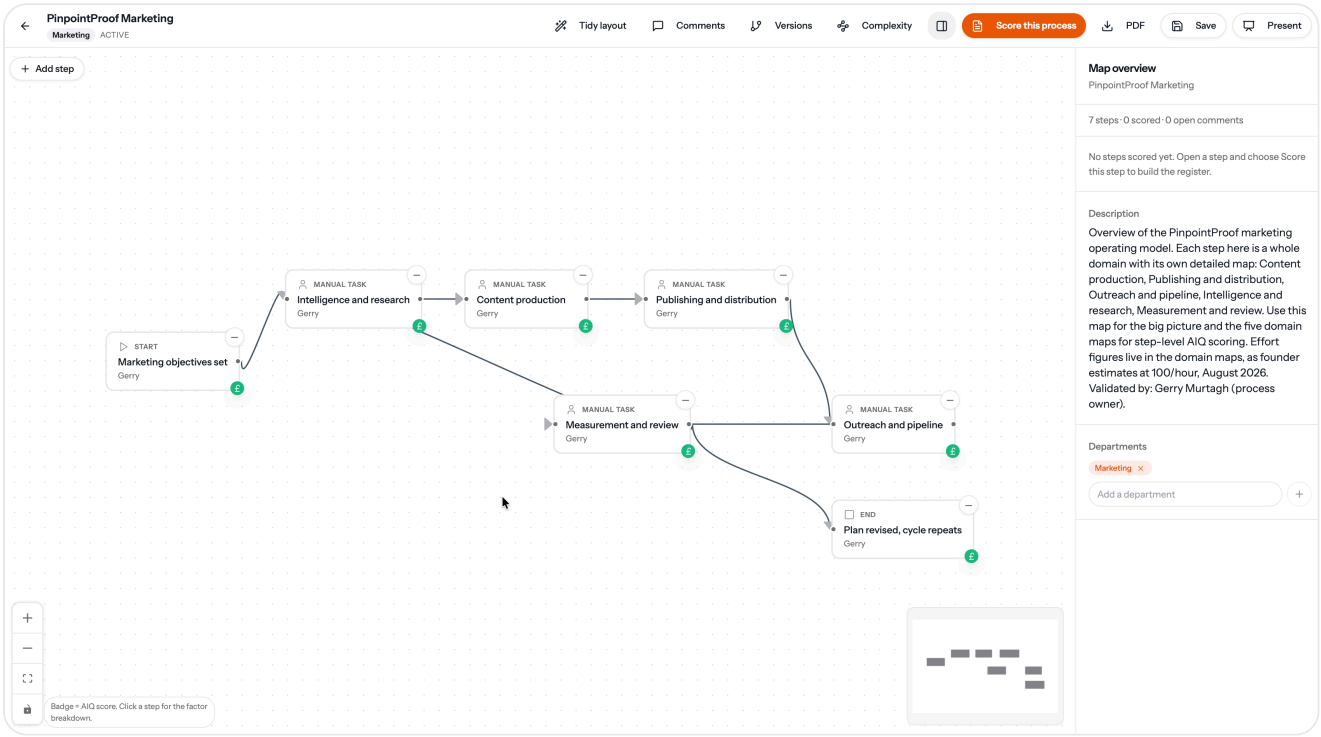

If we had skipped the map, we would have built a posting tool. The map is why we built a learning loop instead.

The value of mapping showed up immediately in the nodes nobody would have listed from memory. Stamping the live URL of a published post so its performance can be measured later. The first hour after posting, when replies drive reach more than the post itself does. Triaging connection requests so warm interest is answered while it is still warm. None of these appear on a marketing suite's feature list, and all of them turned out to matter.

Figure 4. The intelligence and research branch in detail: watchlist sources, sweeps, findings and the decision points between them.

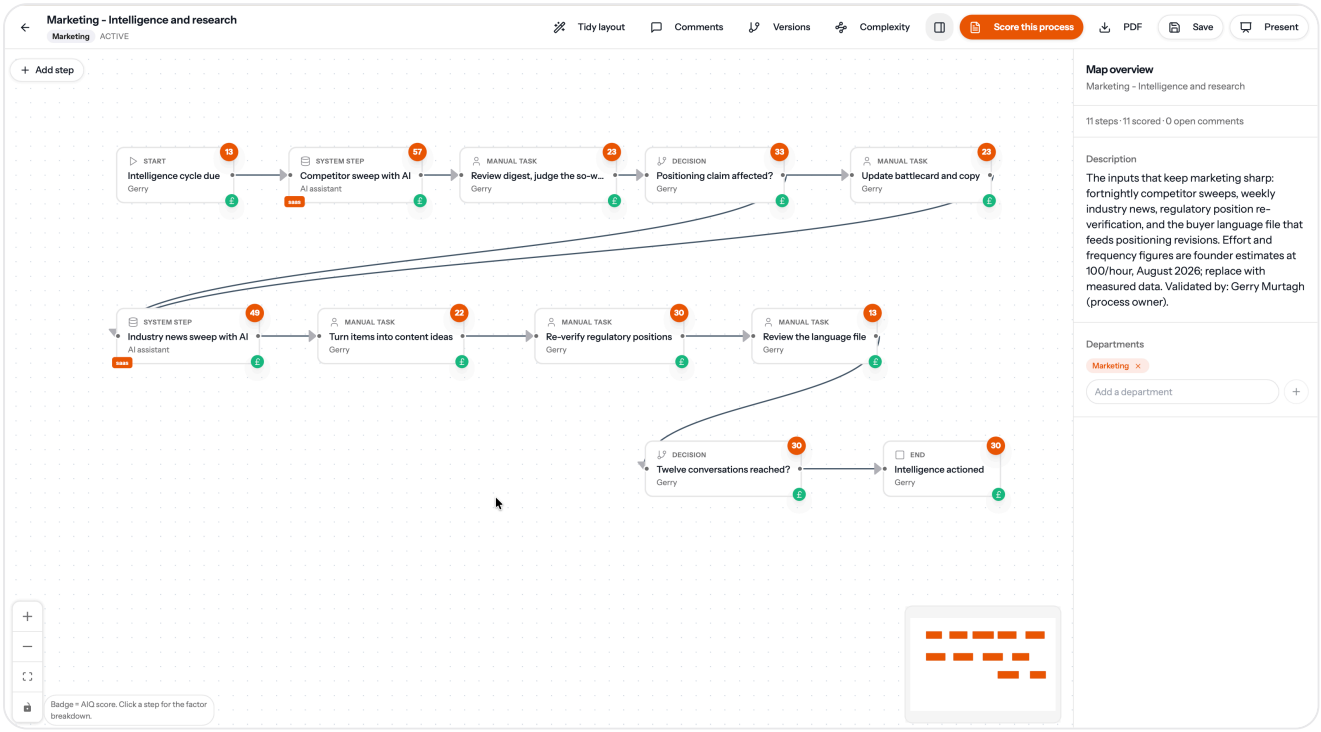

What a good map captures that memory misses

The connective steps

Stamping dates and URLs, logging sources, routing outcomes: the bookkeeping that makes measurement possible later.

The judgement points

Where a person genuinely decides: which finding deserves a point of view, which contact is warm, which draft is ready.

The waits and handoffs

Where work sits between steps, and who notices when it has sat too long.

The steps with external impact

The places where an error reaches a client, a prospect or a regulator rather than staying internal.

4 Stage two: scoring every step

With the map drawn, every node was scored. The scoring lens is deliberately simple, because it has to survive contact with a board: how much judgement does this step need, how repeatable is it, what is the impact if it goes wrong, and is there data to verify it against?

High repetition with low judgement scores for automation. High judgement with reputational impact stays human. The wide middle becomes *AI*-assisted: a model does the labour, a person makes the decision, and the approval is recorded.

Figure 5. The scoring picture: judgement required against repeatability, with the four verdicts.

Figure 6. *A* scored node in PinpointProof, with its *AI* opportunity assessment.

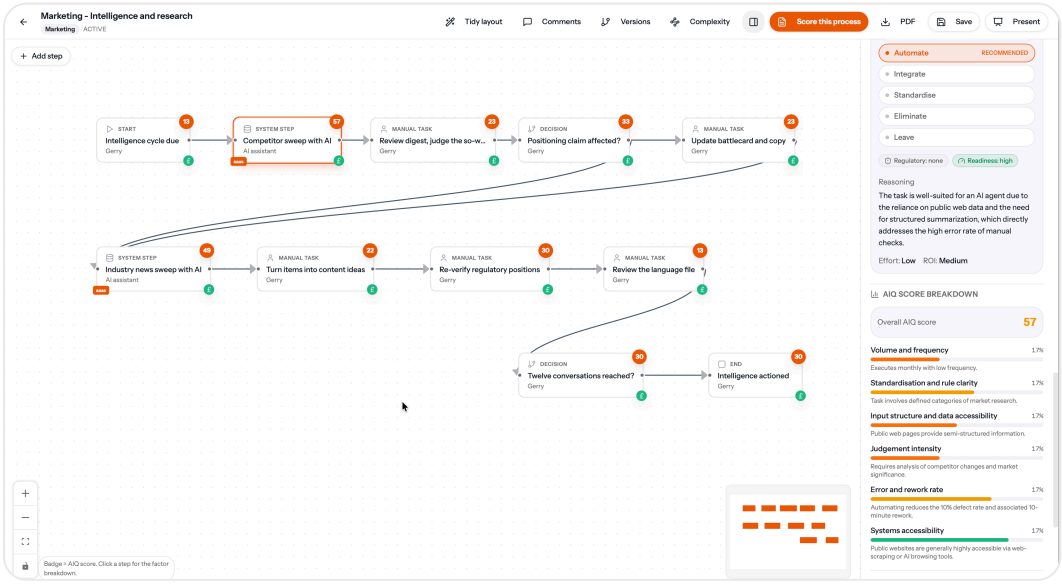

Figure 7. The decision table: an extract from the full register. The complete node-by-node table is in the appendix.

Process node	Verdict	Why
Competitor sweep	Automate	Repeatable research against a defined watchlist; findings arrive as structured records with sources
Industry news sweep	Automate	Defined topics, defined sources, weekly cadence
Deciding which findings deserve content	Keep manual	Commercial judgement about positioning; AI supplies the summary and the so-what, a person makes the call
Drafting for eight channels	AI draft, approve	High labour with consistent rules; nothing publishes without sign-off
Search and AI-answer packaging	AI draft, approve	Mechanical once the article exists: metadata, structured data, FAQs
Publishing	Assisted	A calendar with readiness states; nothing sends itself, by policy
Stamp published date and live URL	Automate	Pure bookkeeping, and the foundation for measurement
First-hour comment engagement	Keep manual	Replies in the first hour drive reach more than the post does; high judgement, poor automation candidate, by design
DM and connection request triage	Assisted	The queue fills itself and replies are drafted in our voice; a human sends every one
Routing warm contacts to the CRM	Assisted	The decision is human; the clerical work is one click
Capturing performance numbers	Assisted	Platforms expose little; an assisted browser run reads the analytics and records the source of every figure
Classifying published posts	Automate	Hook, tone, structure and theme, captured at the moment of publishing to feed the learning loop
Turning performance into writing rules	AI proposes, approve	The playbook changes only when a person accepts the evidence
Applying approved rules to new drafts	Automate	Once approved, a rule should never be forgotten
Weekly digest of state and actions	Automate	The system nudges its operator rather than waiting to be opened
Link tracking and attribution	Automate	Every click logged; every post traceable to the conversations it started

Three first-pass scores that changed, and why that is healthy

We first scored publishing as a full automation candidate, and pulled it back to assisted after weighing what a mis-post costs a firm that sells credibility. DM triage began as an afterthought and became one of the most valuable assisted workflows in the platform. And performance capture turned out to need assisted capture with every figure's source recorded, because the platforms expose very little. First scores usually contain a comfortable assumption; the method is iterative so it gets found.

5 Stage three: the target operating model

Scores in hand, the design fell out almost by itself: a closed loop rather than a toolbox of point features.

Signals come in from the sweeps. Signals worth acting on become topics; topics become drafts across eight channels; drafts pass guardrails and approval; approved work is published on a calendar, with the first hour protected for human engagement. Published work is measured. Measurement becomes proposed writing rules. Approved rules reshape the next round of drafts. The loop closes.

Figure 8. The operating loop, with each stage owned by the founder, the AI in the platform, or the AI assistant working alongside in the browser.

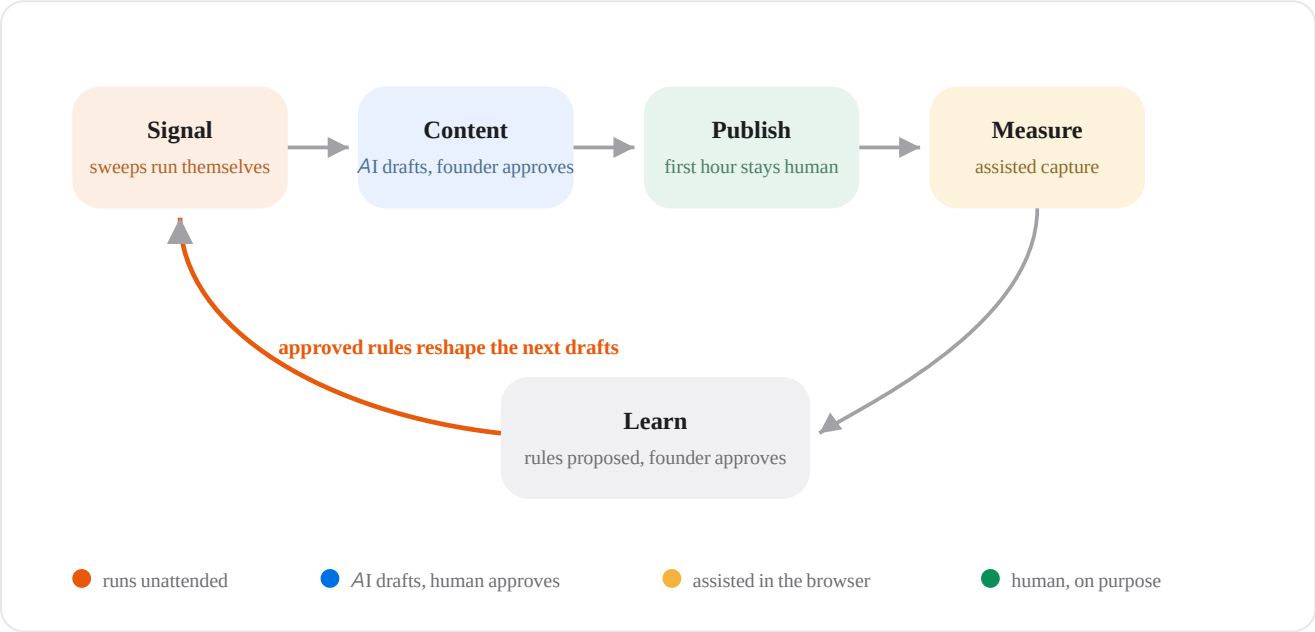

Around the loop sits a division of labour, made explicit and colour-coded in the platform itself. Some things run unattended. Some things AI drafts for approval. Some things an AI assistant does alongside the founder in the browser. Four things belong to the founder alone: judgement calls in approvals, the first hour after posting, real conversations with warm people, and the client work that pays for all of it.

6 Stage four: building it, under control

This chapter answers the question every operations leader asks next: can this actually be built without a programme, a vendor and a year?

It took 7 working days. The interesting part is the control that made the speed safe.

Mockups before code, every time

For every feature, the AI produced design mockups first: full screen layouts with real interaction states, in the house design language. Nothing was built until the mockups were approved. This sounds like ceremony and works out as economy: rejecting a mockup costs minutes, while rejecting a built feature costs days. Several features changed substantially at mockup stage, including one redesigned from a pop-up window into a full working screen after the first version was declined on sight.

Verify, checkpoint, record

Every feature ended the same way: automated checks pass, a named checkpoint is created so the platform can be restored to any known-good state, and a written record is filed of what was built and why. The build has an audit trail for the same reason our clients' processes need one: because a working system should be able to show how it came to work.

Figure 9. The build rhythm, feature by feature.

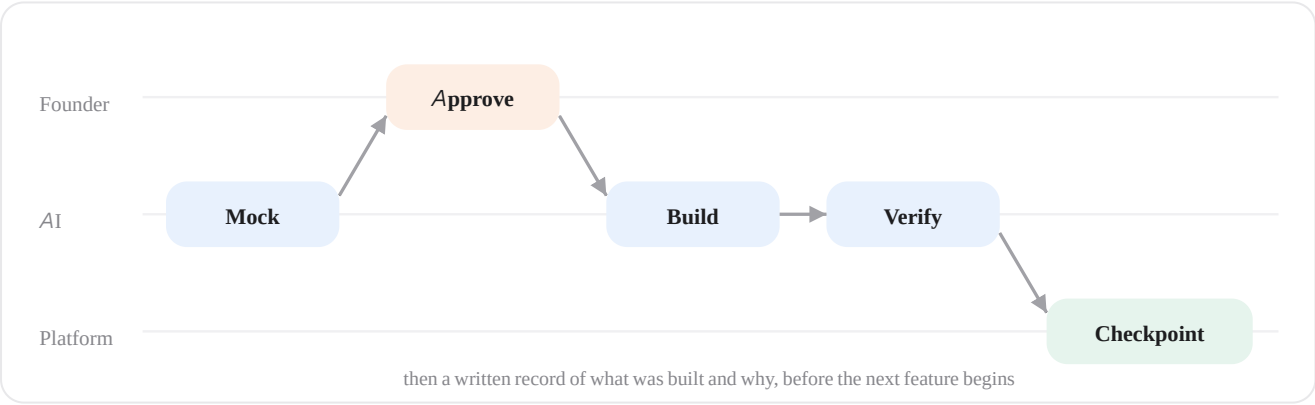

Figure 10. The approved design mockup of the publishing calendar (left) and the shipped screen (right). *A*pproval happened on the mockup; the build followed it.

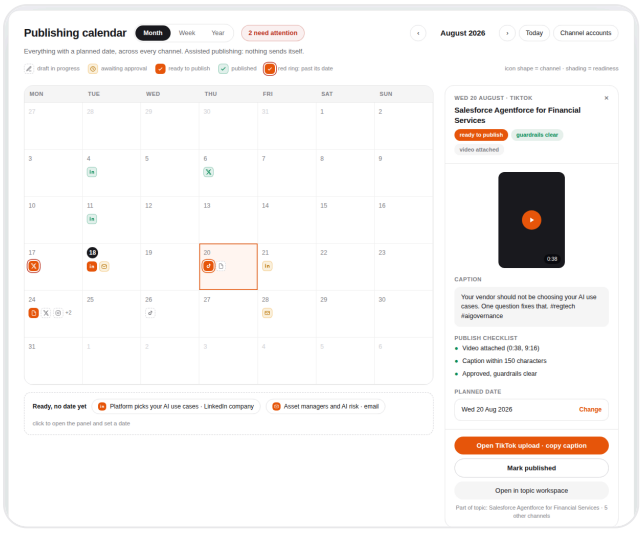

The mockup, as reviewed and approved

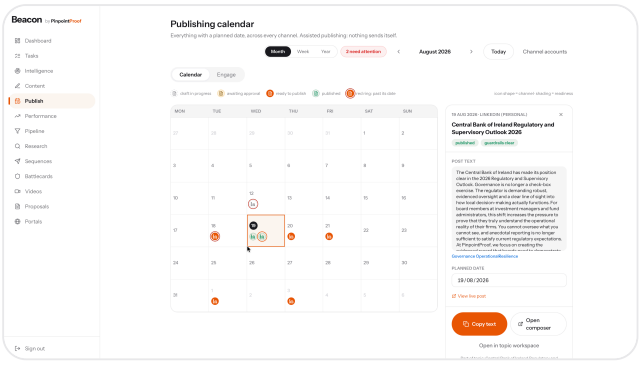

The shipped screen in Beacon

The map was the specification

Feature requests were process nodes, often almost verbatim. "Stamp the published date and store the live post URL for the metrics module" is a sentence from the map and a shipped feature, unchanged. *A* good process map turns out to be the cheapest possible specification for an *A*I builder, and that point deserves a moment: the mapping discipline that satisfies a regulator is the same discipline that makes *A*I delivery fast.

When it went wrong, the governance absorbed it

Early in the build, a platform-level data issue silently dropped newly added fields for several hours: records saved, and parts of them vanished. In an ungoverned build this is the kind of fault that erodes trust in the whole system, because nobody can say what else might be wrong. Here it was diagnosed against the audit trail, worked around with a design change, and written into the build rules so it cannot recur. *A* governed build turns surprises into recorded lessons.

Guardrails live in the product

Beacon's approval flow runs every draft through a guardrail engine before it can be approved: no invented statistics, no claimed customers or results while we have none to claim, house tone rules enforced mechanically across every visible field. *A* draft that breaks the rules requires a written override note before approval will stick. Nothing publishes itself, on any channel, ever.

Figure 11. The guardrail engine at work: a deliberately invented claim ("One million people use PinpointProof") is flagged twice, and approval is withheld pending an override note.

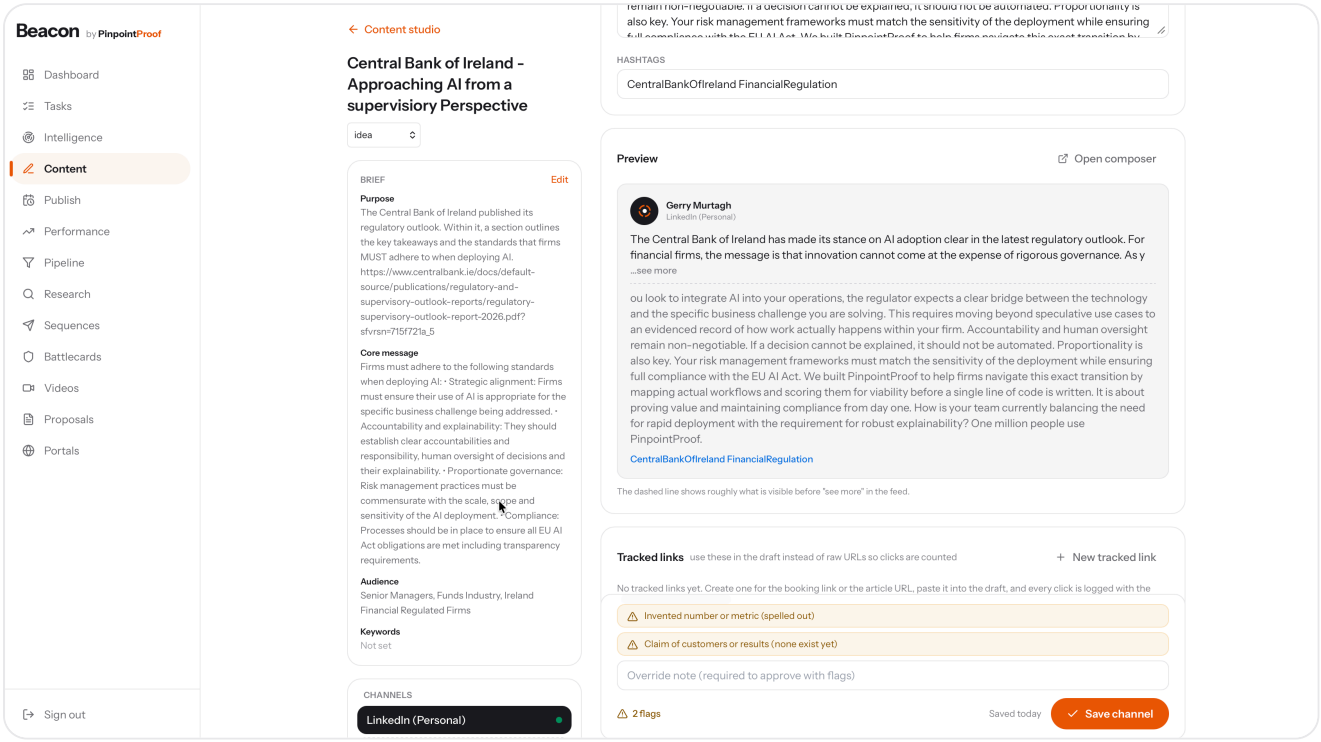

The stack, for completeness

PinpointProof produced the map and the scores. Claude, *A*nthropic's *AI*, working in an agentic coding environment, did the building, the testing and the record-keeping. Base44 provided the application platform. Combined running cost: £125 per month. The method matters more than the particular tools; what the method did was make fast tools safe to use.

7 The result: a tour of Beacon

Every screen that follows exists because a node on the map earned it.

The marketing pipeline dashboard

The whole engine on one screen: a funnel from intelligence findings through content and publishing to measured performance, with conversion rates between the stages, and an action queue where every waiting task is tagged with its owner: the founder, the *AI* in the platform, or the *AI* assistant in the browser. The dashboard routes work to its owner.

Figure 12. The marketing pipeline dashboard.

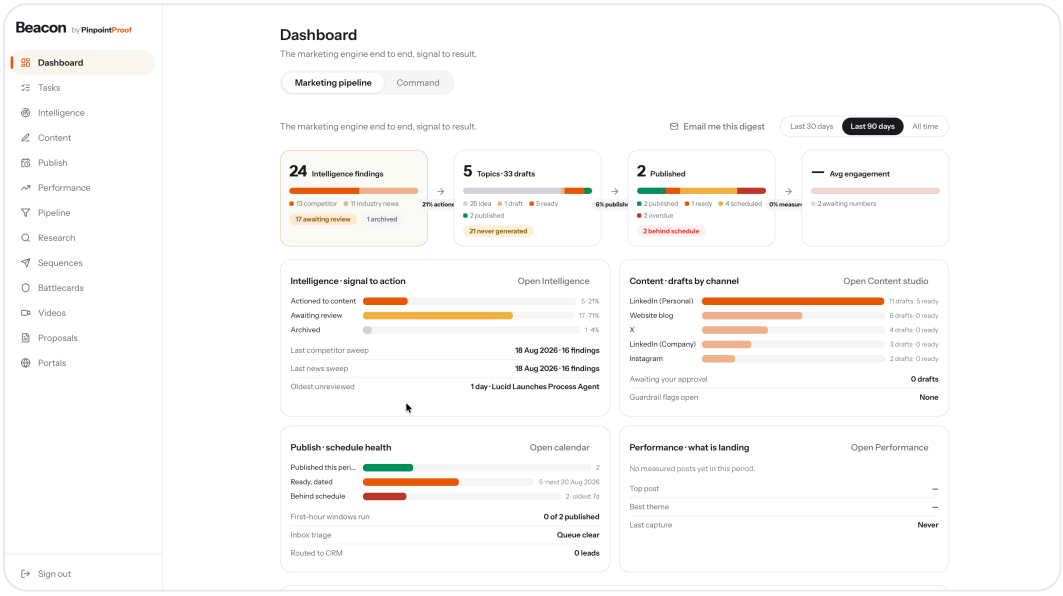

Intelligence

Scheduled sweeps of competitors and industry news write structured findings with sources and a why-it-matters line. Findings are actioned into content or archived with a note, so the queue stays clean.

Figure 13. The intelligence findings list.

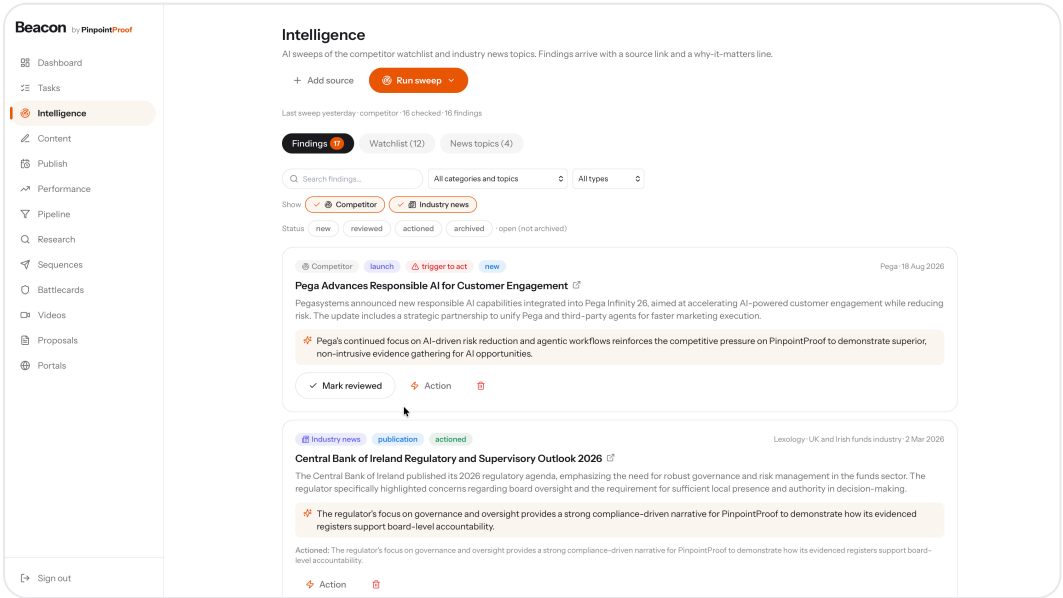

The topic workspace

One topic, eight channel drafts: website article with search and AI-answer packaging, both LinkedIn voices, X, email, and video briefs for the short-form channels. Approved writing rules from the performance playbook are injected into every generation automatically, and each draft shows which rules shaped it.

Figure 14. The topic workspace.

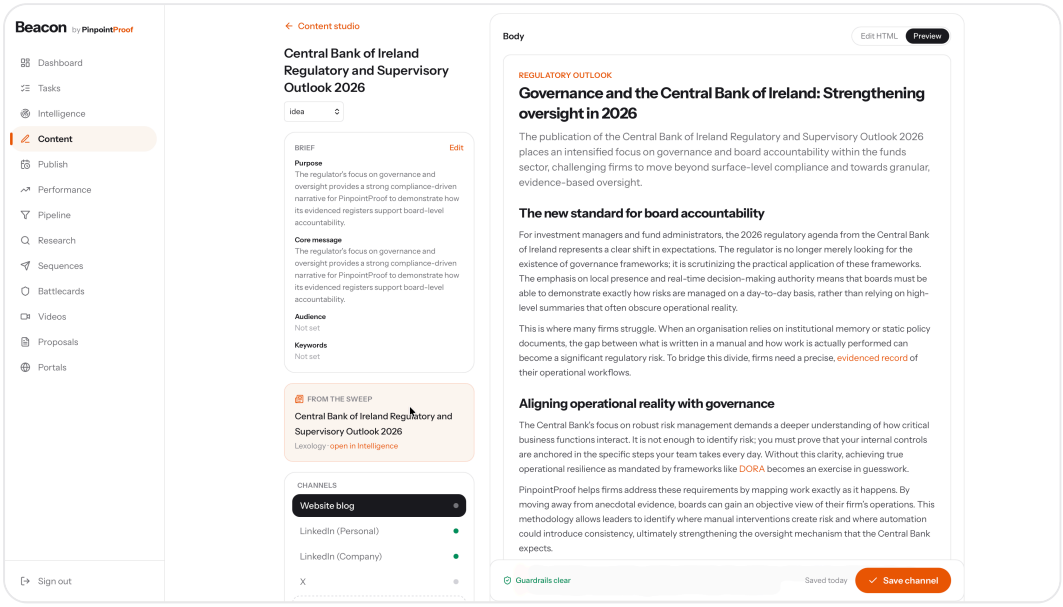

The publishing calendar

A month view with every planned item shaded by readiness, and a red ring on anything past its date. Marking a post published stamps the date, stores the live URL, classifies the post for the learning loop, and offers to start the first-hour window: a sixty-minute countdown with one job, stay in the comments.

Figure 15. The publishing calendar, month view.

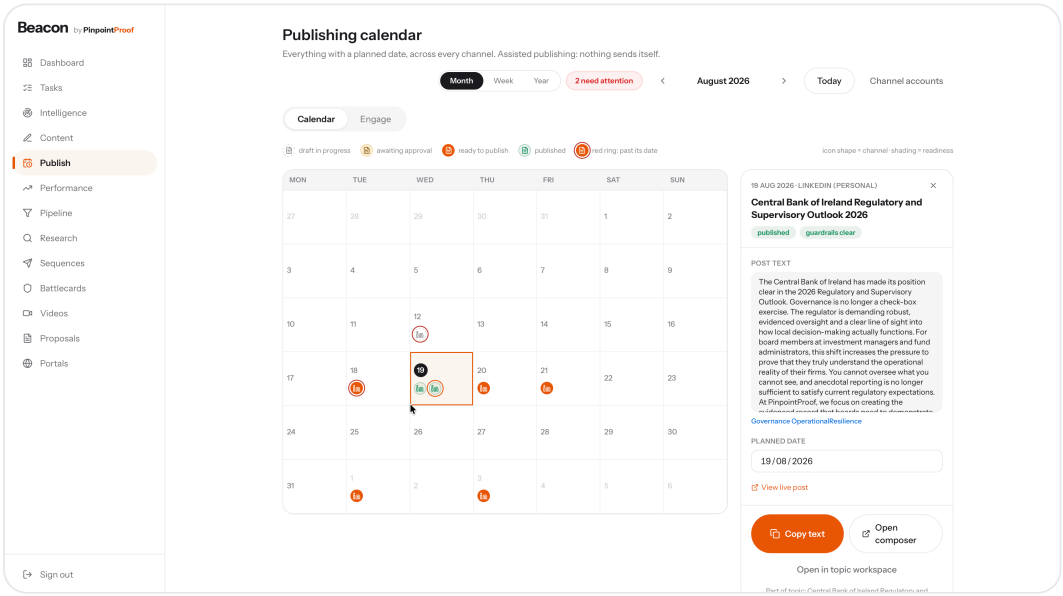

The engage tab

DMs and connection requests arrive in a triage queue with the sender's context. Replies are drafted in our voice for approval, warm contacts route into the CRM as leads with one click, and the originating post is recorded, so any deal can be traced back to the content that started it.

Figure 16. The engage tab, with a live first-hour window.

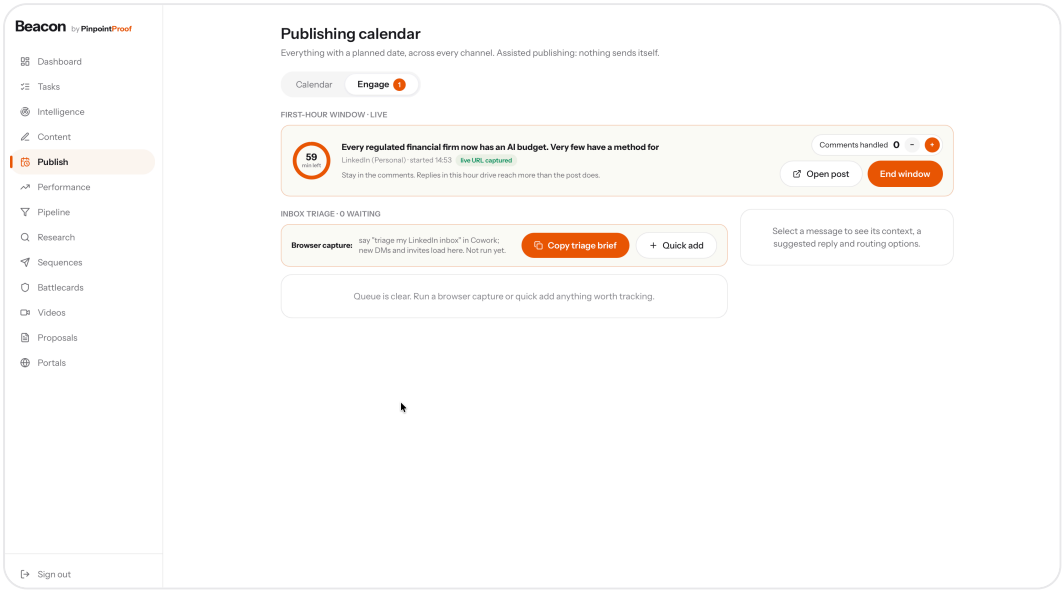

Performance and the playbook

Captured numbers become engagement rates by post, channel and theme. *A*nalysis runs only when there is enough data, proposes writing rules with the evidence attached, and marks thin evidence as an early signal. *A* human approves every rule before it touches a future draft. Tracked links complete the chain: post to click to conversation to lead to demo to signed sprint. The system is designed to learn from revenue.

Figure 17. The performance playbook: proposed insights with evidence, approved rules live in the generation prompts.

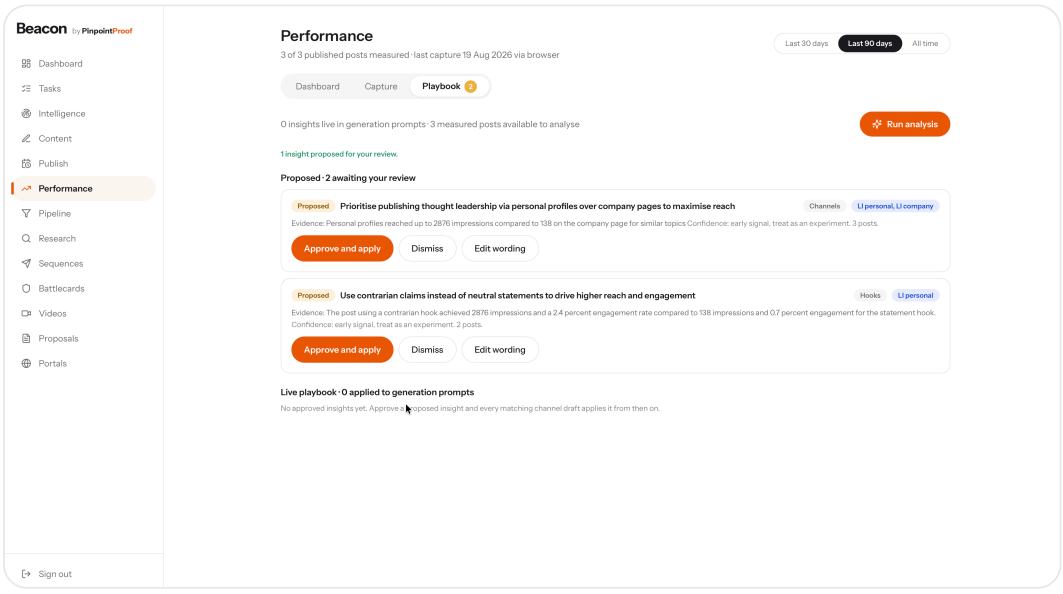

8 What stayed human, and why that is the point

The easy story about AI is what it replaced. The more valuable story is what the scoring protected.

Approving every draft

Voice and credibility are the two assets a young firm in a regulated market cannot re-earn quickly. Every word that reaches a reader passed a person first, with the guardrails as a second pair of eyes.

The first hour after posting

Replies from a person drive reach and relationships in a way automation cannot match. Our own process map called it out: high judgement, poor automation candidate, by design. The platform protects the hour with a countdown; a person spends it.

Real conversations

Warm people get a human, for the obvious reason that people can tell. AI drafts the reply; a person decides, adjusts and sends it.

The judgement calls

Which finding deserves a point of view. Which insight deserves to become a rule. Judgement under uncertainty is precisely what the firm sells, so the platform queues those decisions rather than making them.

The map told us what we could automate, and it told us what we should not. For a regulated firm, where the stakes are client money and regulatory standing, that second list is the one the board needs most.

9 The numbers

Everything in this table is real, current at publication, and deliberately conservative.

Measure	Before	After
Weekly effort to run the marketing engine	Roughly 40 hours, a full-time role	Around 4 hours of founder time
Software cost	A comparable SaaS marketing suite runs at roughly £3,000 per month	£125 per month, AI tooling and platform combined
Build investment	Platform selection and implementation typically runs to quarters	7 working days, mockup-approved feature by feature
Features shipped		28 features across six modules
What is deliberately unclaimed	No revenue or client results are attributed to the system yet. The attribution loop that will prove or disprove them is part of the build.	

Figure 18. The founder's marketing week, before and after.

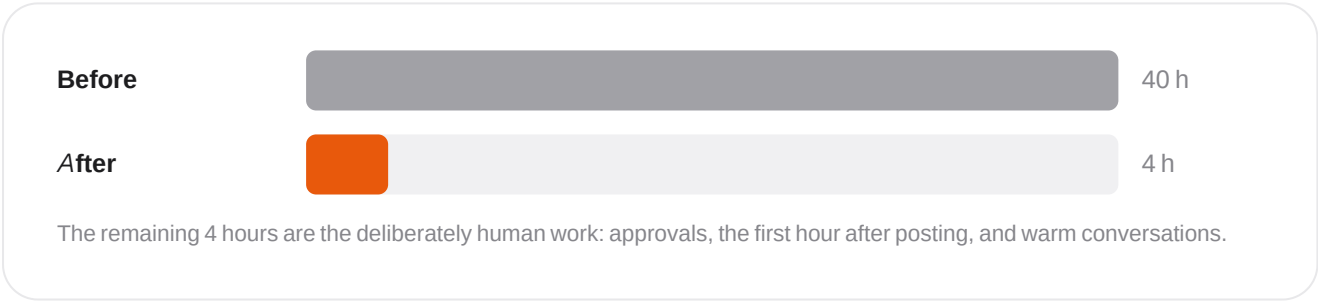

The remaining 4 hours are the deliberately human work: approvals, the first hour after posting, and warm conversations.

The last row of the table is there on purpose. Beacon's own guardrail engine would flag this document if it claimed results it does not have. The measurement loop exists so that when we do make claims, they will be evidenced, which is, after all, the product.

10 What this means for your firm

The method transfers directly beyond marketing. Replace the word and it holds: client onboarding, fund launches, breach reporting, client reporting, board MI. Each is a set of processes someone could map faithfully, score for judgement and repeatability, and split deliberately into automate, assist and keep human, with the decisions recorded well enough to show a board or a regulator how they were reached.

That is what a PinpointProof engagement does, in four weeks, at a fixed fee:

Week 1
Map how the work actually happens, with the people who do it

Week 2
Score every step and agree the decision table

Week 3
Design the test that proves the case, before anyone builds anything

Week 4
The evidenced record: maps, scores, decisions and the business case, in a form a board can stand behind

We ran this method on ourselves before asking any client to trust it. This document is the evidence.

See it on your own process

In the first session of a walkthrough, we map one of your processes live. Bring one process; we will show you where *AI* belongs in it, and where it does not.

Book a walkthrough · pinpointproof.com gerry.murtagh@pinpointproof.com

APPENDIX

The complete decision table

All mapped marketing nodes with their verdicts, in the four-verdict scheme. This is the register the method produces; yours would look like this, for your processes.

Process node	Verdict	Rationale
Watchlist and news source maintenance	Keep manual	A quarterly judgement about who matters; small effort, high leverage
Competitor sweep	Automate	Scheduled research against the watchlist; structured findings with sources
Industry news sweep	Automate	Defined topics and sources, weekly cadence, why-it-matters line per item
Finding triage: action or archive	Keep manual	Positioning judgement; archived items carry a note so the queue stays clean
Topic brief creation	AI draft, approve	Purpose, audience and core message drafted from the finding; person edits
Channel drafting, eight channels	AI draft, approve	Consistent voice rules; approval gate on every draft
Search and AI-answer packaging	AI draft, approve	Metadata, structured data, FAQs, once the article exists
Video briefs and production handoff	Assisted	Script and shot list drafted; production and attachment stay with a person
Guardrail check on every draft	Automate	Mechanical rules, no invented numbers or claims; override requires a written note
Scheduling and the publishing calendar	Assisted	Readiness states and overdue flags; nothing sends itself, by policy
Stamp published date and live URL	Automate	Bookkeeping at the moment of publishing; the foundation for measurement
Post classification for the learning loop	Automate	Hook, tone, structure, theme, captured automatically on publish
First-hour comment engagement	Keep manual	High judgement, poor automation candidate, by design; a countdown protects it

Continued on the next page.

APPENDIX, CONTINUED

Process node	Verdict	Rationale
DM and connection request capture	Assisted	Assisted browser run fills the queue; every item records how it arrived
Reply drafting	AI draft, approve	Drafted in our voice with the source post as context; a person sends every reply
Routing warm contacts to the CRM	Assisted	One click creates firm, contact, lead and first interaction; the decision is human
Lead follow-up timing	Automate the nudge	The reminder is automated; the message stays human
Performance capture	Assisted	Assisted browser run reads platform analytics; every figure records its source
Engagement analysis and insight proposals	AI proposes, approve	Runs only with enough data; thin evidence is labelled an early signal
Applying approved rules to future drafts	Automate	Approved rules inject into every matching generation, visibly
Competitor battlecard refresh	AI proposes, approve	Field-by-field proposals from sweep evidence; person accepts each change
Link tracking and attribution	Automate	Every click logged with source; posts traceable to conversations and deals
Weekly digest and ritual nudges	Automate	Monday email with the week's numbers and the owner-tagged action queue
Selling	Keep manual	Conversations, demos and sprints are human acts; the platform only queues and records

Four verdicts, twenty-four decisions, each with a recorded reason. This is the register the method produces. Yours would look like this, for the processes that matter to your firm.

PinpointProof

From process map to proven AI value. We map how work actually happens, score every step for AI and automation potential, and produce a ranked, evidenced opportunity register your board can rely on.

pinpointproof.com

hello@pinpointproof.com · gerry.murtagh@pinpointproof.com

This case study describes PinpointProof's own operations. No client data or client results are referenced. Product screenshots show the firm's live internal platform; figures are real and current at publication.