GTM PHILOSOPHY · MARKETING ENGINEER TAKE-HOME · AUGUST 2026 · FOR JULIAN CLARKE

Three plays, one system, one number.

Julian, the number is submission volume: more of your 40,000 offices sending risk, and the brokers who already do sending more of it. Three broker-facing plays, built to feed each other, each with a working tool behind it. Click around instead of taking my word for it.

TOP OF FUNNEL | THE ORDER IS THE ARGUMENT | DECISION

03

AI visibility

Clips and cited answers reach a broker who has never heard of Shepherd.

02

Targeting

The same named offices, made familiar before the email.

01

Signal Engine

Brokers, signals, drafts, briefs: the moment a signal fires.

THE NUMBER

Submission volume

Every touch becomes a signal on the broker record.

Wider: offices that never sent a risk.

Deeper: one or two accounts a year becomes seven-plus, where quote rate jumps ~20 points and hit rate 10 to 15.

ORCHESTRATION

How the three are orchestrated

They are one system, not three campaigns, and the order is the argument.

01

Play 3 earns the attention

A clip or a cited answer reaches a broker who has never heard of Shepherd. Cheap, slow, compounding.

02

Play 2 makes it familiar

The same named offices see paid placement, and anyone who visits the site gets retargeted. Days, not months.

03

Play 1 converts the moment

A signal fires at that firm and the underwriter's email lands into recognition instead of cold. This is where a submission actually happens.

04

Everything feeds back

Every touch from any play is logged as first-party intent on the broker record, which raises their priority and re-orders tomorrow's work.

01

CONVERTS THE MOMENT

Signal-led broker outbound

Score the broker universe, watch for the moments a broker is about to place risk, and have agents write the first email for the underwriter.

OBJECTIVE

New submitting offices, and first submissions from brokers who have never sent one.

HYPOTHESIS

A broker replies to an underwriter, not a marketer, when the email lands the week something changed at their firm and names it. If reply rate clears 20%, the thesis holds.

SIGNALS AND CHANNEL

Third-party: project awards, permits, practice hires, carrier pullbacks. First-party: logged engagement. Lifecycle and outbound email, sent from the underwriter's own mailbox.

LIFECYCLE, AND WHO ACTS

Acquisition, then first-to-repeat: open offices that never submitted, then walk them toward the 7-plus pace where quote rate jumps ~20 points. The agent drafts, the underwriter approves and sends. Marketing owns the signal book.

WHAT I BUILT

Every construction and energy practice title at your 18 appointed brokerages, pulled from Apollo and scored on four arguable axes: fit, timing, access, intent. A sweep reads the trade press for moments that mean risk is about to move. When a broker at a flagged firm floats to the top, the agent writes the first touch in your underwriter's voice and waits for approval. The one-pager is here too: type a name, get the brief.

WHY IT MOVES SUBMISSIONS

The universe finds offices that never sent anything; the signals tell the underwriter which week to call. One of its own cold emails became this interview.

FEASIBILITY

Running today: 338 real contacts, 40 real signals in 45 days, clickable in the tool. Full-strength permit and project data is a licensed upgrade, priced in the budget.

> The signal selection itself is an argument, not a feed: seven types ranked into act, aim, and context tiers, each with its sourcing, criteria, and confidence written down, and loss events deliberately excluded from outreach. Read the signal thesis →

OPEN THE FUNNEL MATH: 21 SUBMISSIONS A MONTH AT 60 EMAILS A DAY, ASSUMPTIONS EDITABLE +

Apollo about $200 a month at 2,000 brokers. Claude under $50. Mail goes through whatever you already use.

Kill rule. Under 20% reply rate on signal-led email after 300 sends, I change the copy, not the target.

02

MAKES IT FAMILIAR

Broker ABM and paid media

The scored universe is already an audience. Point paid media at the same named offices the underwriters are about to email, so the email lands into recognition instead of cold.

OBJECTIVE

Be already familiar at the offices Play 1 is about to email, and capture the brokers who visit the site without ever replying.

HYPOTHESIS

Almost no carrier here runs a modern paid stack, so a small spend against named offices is cheap and uncrowded. If targeted offices don't out-reply the control, it stops.

SIGNALS AND CHANNEL

First-party: site visitors and content viewers, retargetable. Third-party: Play 1's signals, used to time a flight. Paid media and ABM. LinkedIn by job title inside named brokerages, plus the trade newsletters brokers already read.

LIFECYCLE, AND WHO ACTS

Acquisition support: makes the office familiar before Play 1's email lands. Claims no stage of its own. Marketing end to end; underwriters only approve the target list.

WHAT I BUILT

The broker table exports as a LinkedIn ad audience: company names and job titles, filtered to the offices Play 1 is emailing that month. Paid and outbound point at the same names on purpose. Spend and outcomes land in the Spend tab, so cost per submitting office is a real number.

WHY IT MOVES SUBMISSIONS

Indirectly, and I'll say so: brokers do not click an ad and send a risk. Paid makes the underwriter's email the second time an office has seen the name that week. Smallest bet on the board, pointed at a list you can name.

FEASIBILITY

Feasible this week: the audience exports from data already here, no credits, no emails. Needs a LinkedIn ads account and $500 to $1,500 a month.

$500 to $1,500 a month to start, which is small enough to kill without a meeting. No agency.

Kill rule. 90 days. If reply rate at targeted offices does not beat the control, the budget moves into Play 1's data.

03

EARNS THE ATTENTION

Broker-facing content and AI visibility

One recorded hour a month with Justin and Mo becomes a month of clips, and the same material becomes the pages an LLM cites when a broker asks it who writes this risk.

OBJECTIVE

Be the market a broker recognises before the email lands, and the market an LLM names when a broker asks where to place a risk.

HYPOTHESIS

Brokers research markets the way everyone now does: by asking a model. Legacy carriers carry twenty years of sentiment, Shepherd carries none, and nobody else is competing for the answers.

SIGNALS AND CHANNEL

First-party. Engagement with a clip or a page is a signal that raises that broker's priority in Play 1. Video and social, plus answer-engine optimisation.

LIFECYCLE, AND WHO ACTS

Deepening: engagement feeds the Intent axis and climbs brokers toward the 7-plus pace. Also the front door for brokers who discover markets by asking a model. Marketing produces; Justin and Mo give one recorded hour a month; the audit runs itself. Underwriters spend zero time here.

WHAT I BUILT

An audit that asks every answer engine the questions brokers actually ask, records who gets named and cited, and drafts the comparison page that closes each gap. It runs today on five engines; six more are written and start with a key. The clip side turns one recorded founder hour into a month of posts.

WHY IT MOVES SUBMISSIONS

Slowest of the three and I'd fund it last, but it compounds: a clip costs nothing to distribute again, and a page that earns a citation keeps earning it. I ran this playbook once and got 1.5 million views on under $1,000. The mechanics translate; the content changes.

FEASIBILITY

The audit is live on five engines today, so the baseline is real. Clips need one founder hour a month and an editor.

> Prompts vary infinitely, so the program is ranked, not sprayed: comparison prompts first, objection prompts second, shortlist prompts last, argued by who is asking, how close they are to a submission, and how winnable the answer is. Read the prompt ranking →

The five clips I'd cut first, the posting cadence, and the tracking live on the Social page →

One recorded founder hour becomes a month of posts. About $30 a month in model calls across engines. An editor at $1,500 a month, or me. No paid distribution until a clip earns it.

Kill rule. Six pages live, 90 days, under 15 points of share-of-voice movement. Stop.

ROLLOUT

Ninety days, one number

The budget below is the full amount if every optional line turns on. The realistic first month is under $400, because the AEO tool and paid video wait for proof.

BUDGET, MONTHLY

Apollo (seats + credits for ~2,000 brokers) Search is free; credits only on reveal.	01	$200
Claude and other model calls Scoring, drafts, briefs, and the audit across engines. Nothing runs on a timer.	01 03	$80
Email sending and tracking Whatever the underwriters already use, plus tracking.	01	$100
LinkedIn ABM against the exported audience Starts small with a held-out control so the lift is measurable. Kill at 90 days if it does not beat the control.	02	$750
AEO measurement (AthenaHQ Starter), optional Only if the first 30 days show movement. Otherwise $0. It buys Google AI Overviews and AI Mode, which no model key can reach.	03	$295
Video editor (contractor) or my time Half a day a month. Captions and scheduling under $100.	03	$1,500
Paid distribution on the two best clips Starts day 61, targeted at the same broker audience Play 2 builds.	03	$500
Everything on		$3,425
Realistic month one		$380

No agency retainers. No new headcount. Tooling and one contractor I can manage.

TARGETS

METRIC	30	60	90
New submitting offices Offices at appointed firms sending a first risk.	baseline	+5	+12
First submissions from never-submitted brokers From the funnel math at 60 emails a day.	baseline	10	25
Brokers on a 7+ submissions-a-year pace Your stat from our call: at 7+ a year, quote rate jumps ~20 points and hit rate 10-15. The deeper half's whole target.	baseline	+5	+15
Reply rate, signal-led email Below 20% and the copy changes, not the target.	20%	22%	25%
Share of voice on comparison prompts Claude, Llama now; ChatGPT, Gemini, Perplexity with keys.	baseline by engine	+8 pts	+15 pts or kill
Brokers engaging with content, per month Counted against the Play 1 universe, and logged as first-party intent on the broker.	first clips live	100	250
Hours of Julian's week this costs Review the draft queue, approve pages, say yes to clips.	3	2	2

Targets are mine, built from the funnel math above. They become real when I have your baseline.

ON THE CALL

What I want to ask you

I guessed at a lot of numbers to build this. These four answers would change the plan the most.

01

What's the submission baseline? Monthly count, active offices, and brokers who have sent at least one. Where does it live today?

02

What does the broker data look like right now: appointments by office, a CRM, spreadsheets? How is underwriter outreach tracked, if at all?

03

Is there any paid budget today, and what would you need to see in 90 days to keep funding it?

04

Which risk classes do you want more of this half, and which are you fine losing on price?

How this was built: Next.js on Vercel, Supabase, Claude for the agents, a VPS for the scheduled sweeps, all in a week of evenings. The stack details live on the stack page, because the argument here is the plays, not the plumbing.