Fractional Demand

THE GTM ENGINEERING PLAYBOOK

11 Plays Working Right Now.

The GTM RevOps stack we use to turn signal into pipeline.

WHAT'S INSIDE

→ fractionaldemand.com

FULL-STACK REVOPS

How we think about GTM RevOps.

Less about clicking buttons inside a CRM. More about orchestrating a stack that fits your business.

RevOps used to be about knowing where every button lived inside HubSpot. That still matters. But the bar moved. With AI, the value isn't navigating someone else's platform anymore. It's about doing things like building custom dashboards that don't depend on what the CRM lets you build, consolidating disparate signals into something a rep can actually act on, and knowing when to buy a tool versus when to build one.

1. Infrastructure

- Pipeline setup
- ICP scoring and account prioritization
- Target account list creation
- Lead enrichment
- Outbound signal triggers
- AI research automation
- Custom audience scraping for hard-to-reach buyers

2. Operations

- CRM sync and field mapping
- Lead routing and assignment rules
- Contact and company data cleanup
- CRM hygiene and ongoing enrichment
- Custom reporting dashboards (beyond CRM defaults)
- Knowing what to measure and how to measure it

3. GTM Activation

- Automated LinkedIn outreach sequences
- Automated email cadences and frameworks
- Paid ads audience building
- Custom landing pages at scale
- Slack notifications and signal routing
- Meeting prep automation

WHERE REVOPS IS GOING

Orchestration beats button-pushing. The hard part isn't running Clay or HubSpot in isolation. It's deciding when to buy, when to build, and how to wire it all into the workflow your sales and marketing teams already live in. The slickest tool in the world doesn't matter if reps don't use it.

Our approach is to think like marketers first. Solutions get designed around the end result (the trade show you're attending, the channel you're trying to measure, the signal you're trying to act on) rather than the limits of whatever platform happens to be in the stack.

PLAY #1 · TRADE SHOW INTELLIGENCE

Trade Show Intelligence Dashboard.

Walk into a conference with every A-tier account pre-scored before the doors open.

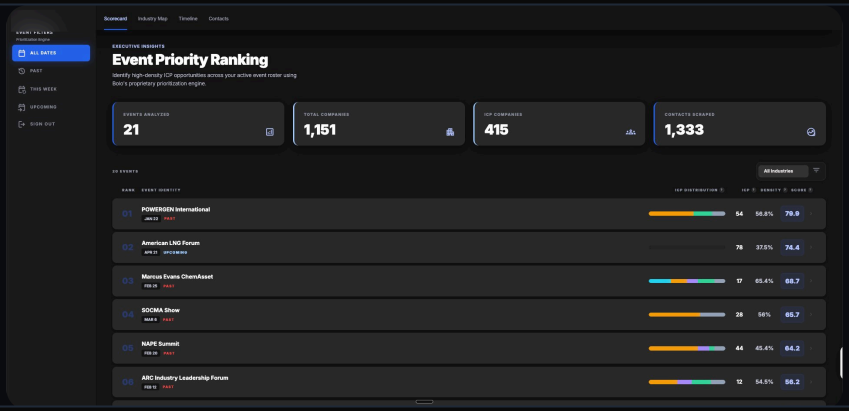

The Problem

Trade shows are the densest ICP moment in the B2B calendar. Three thousand of your exact buyers in one building for three days. Most teams walk in with a printed exhibitor list, a sales deck, and the hope of bumping into the right person between sessions or at their booth. You get a prospectus from the show, which helps, but it doesn't tell you who's actually going to be there. The prep work every team knows they should do rarely happens, because it's manual, the list changes until the show opens, and nobody has hours to pull it all together.

The Result

Your AEs walk the floor with a purpose-built plan, not a lanyard and a prayer. The booth still matters. The prep is what decides whether the booth converts.

STEAL THIS FOR ANY EVENT

Works for customer summits, partner events, even conferences you're not attending. If your ICP is going to be in one building, you can scrape the event site for speakers, panelists, sponsors, anyone mentioned as attending, score the list, and arm your team before anyone sets foot inside. The pipeline looks identical. Just point it at whatever event page you've got.

The Play

1. **Scrape the conference site.** Exhibitors, sponsors, panelists, speakers. Structured as JSON with company, name, role, session topic, and booth number.
2. **Enrich the company and the contacts in Clay.** At the company level: domain → firmographics → tech stack via BuiltWith → funding, news, and job-change signals from the last 90 days, so you know which accounts are worth your time. At the contact level: for every name you scraped (speakers, panelists, exhibitor reps), pull title, tenure, LinkedIn, recent activity, so you know who to lead with.
3. **Score against ICP.** Run a scoring prompt from the Clay table using Claude, ChatGPT, or whatever LLM you prefer. Feed it your ICP and have it score each account on fit. Every account gets a 1–5 score with the reason attached.
4. **Draft outreach for A-tier.** For every account scoring 4 or 5, AI drafts a LinkedIn message your AEs or BDRs can send, tailored to what the person talked about on a recent panel, what they're scheduled to speak on at the show, or something specific from their profile. Layer in invites to dinner, lunch, happy hour, or whatever you're hosting. "Saw you're on the panel Tuesday. Walking the floor Monday afternoon, worth 15 minutes?"
5. **Ship to the team in Slack.** Each AE wakes up Monday of the show to a Slack thread with their personalized list. Booth numbers, names, session times, pre-drafted DMs ready to send.

The Stack

`FireCrawl` `Clay` `Claude Code` `Slack` `HubSpot`

PLAY #2 · DYNAMIC EXPANSION EMAIL

Dynamic Add-On Engine → Expansion Email.

How to turn your add-on catalog and account usage into personalized expansion revenue.

HOW THE ENGINE RENDERS THIS EMAIL

Catalog → eligibility → match → render → send. Every step runs per-account, per-segment, on every send.

1 **Pull the catalog**
Add-ons, modules, integrations, SKUs. Scrape product pages, hit the API, or load a feed.

2 **Apply eligibility filter**
AI flags which items are eligible for this account: plan tier, prerequisites, what's already activated, price band.

3 **Match to account history**
AI matches to account history and look-alike customers, what teams like this one activate next.

4 **Render and send**
Top items render into HTML cards. HubSpot pulls the per-account block on send.

The Problem

Most expansion emails are generic because the data required to personalize them lives in five different places and nobody has wired it together. A SaaS customer is on a $4K-per-month plan with four of your twelve add-ons activated. The "consider upgrading" email goes to every account in that plan tier with the same copy. It drives almost nothing, because it asks the customer to do all the work of figuring out which add-ons are relevant. The ones they'd actually activate are sitting in your product catalog and your usage data.

The Result

The generic email was a nothing-burger. The dynamic version turns your add-on catalog into a merchandised, personalized expansion offer per account. Built once, runs forever, re-renders per customer on every send.

OUR USE CASE: CLICKLEASE

The example above is real. We built it for clicklease, an equipment leasing platform: customer approved for $25K, used $10K, gets a dynamic block of complementary equipment priced inside their remaining credit window. Same pattern works wherever you have a catalog plus account data: SaaS add-ons, marketplaces, field service parts.

The Play

1. **Pull your add-on catalog.** Pull from your product API if you have one. If not, scrape your own pricing and add-ons pages and structure as JSON: name, description, prerequisite plan, dependencies, price.
2. **Apply eligibility logic with AI.** Not every add-on is eligible for every account. Feed the catalog plus the account's plan tier and current activations into Claude or your LLM of choice. Each row gets an `eligible: true/false` flag plus a reason (prerequisite missing, already activated, plan-tier mismatch).
3. **Match to account history and look-alikes.** Look at what the account has already activated, plus what similar customers in their segment activate next. AI ranks the eligible add-ons by likelihood-to-activate.
4. **Generate the HTML block.** Top 3–5 matched add-ons render into a responsive HTML card grid. Name, one-line value prop, price, one-tap "activate" CTA.
5. **Drop into HubSpot.** The HTML block lives as a custom module. When the email fires, HubSpot pulls the block rendered for that specific account. Plan-aware. Eligibility-filtered. Personalized per customer, every send.

The Stack

`FireCrawl` `Clay` `Claude / GPT` `HubSpot` `Custom HTML Module`

PLAY #3 · INBOUND TO CEO DM

Inbound Form → CEO LinkedIn DM + Auto Competitor Intel.

The 90-second response that makes every other follow-up look asleep.

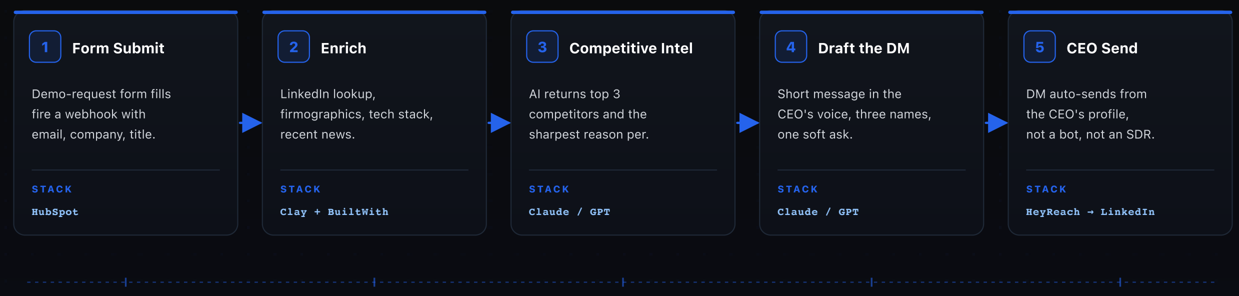

~90 SECONDS END-TO-END

Industry average: 42 hours · Best-in-class SDR: 5 minutes · This play: 90 seconds.

The Problem

Demo-request forms are the warmest moment in your funnel. The average B2B company takes 42 hours to respond to one.[1] Best-in-class teams hit 5 minutes. By 42 hours, your prospect has already booked three other calls. The fix isn't picking a faster channel. It's running more than one at the same time: a real email and an automated LinkedIn DM that goes out within 90 seconds. The kicker: the DM lands from the CEO's profile, not a shared inbox or an SDR.

The Result

The CEO looks personally attentive without actually being personally attentive. Response time collapses from hours to seconds. The message opens a real conversation, not a demo calendar link, on a channel the prospect actually checks.

> **IF THE CEO WON'T SIGN UP**
>
> **The CEO sender is the kicker, not the play.** Plenty of founders aren't comfortable handing over their LinkedIn, and that's fine. The DM can land from an AE, a VP of Sales, or a BDR. Anyone with a credible profile works. You lose some of the surprise factor. You keep the speed, the multi-channel coverage, and the personalized reasoning. That's most of the lift.

The Play

1. **Form submission fires the webhook.** HubSpot form → Clay webhook → record created with email, company, title.
2. **Clay enriches the record.** LinkedIn URL lookup, firmographics, tech stack via BuiltWith, recent news, funding signals. By the end of this step you know everything publicly available about the prospect and their company.
3. **AI identifies their top 3 competitors.** Feed Claude (or your LLM of choice) the company's website, product description, segment, and stage, then ask: "Who are the three most likely competitive alternatives this buyer is evaluating, and what's the sharpest reason per?" Returns three named competitors with one-line reasoning. No manual research, runs in seconds.
4. **AI drafts the DM.** Prompt includes the prospect's name, role, enriched context, the three competitors, and a tone-and-voice sample from the sender (recent emails, posts, anything that captures how they actually write) so the DM sounds like them. The opener signals genuine attention. Something like "just started digging into [Company], saw you're evaluating this category..." so it lands warm, not boilerplate. Then the three competitors, the specific reason, and a soft ask. "Anyone else on your shortlist we should know about?" beats "Want to book a call?" because it invites a reply instead of asking for a commitment.
5. **HeyReach auto-sends within 90 seconds.** Message goes out from the CEO's profile when they're game. If not, the AE, VP, or BDR sender works the same way. Replies flow back with the full enriched context attached.

The Stack

`HubSpot` `Clay` `HeyReach` `Claude / GPT` `LinkedIn`

[1] Drift, Lead Response Survey, 2018. Average response time across 433 B2B companies sampled.

PLAY #4 · AD ENGAGEMENT TO OUTBOUND

LinkedIn Ad Engagement → Clay → Outbound.

The ad does the warming. The signal fires. The BDR calls. That's the whole play.

The Problem

LinkedIn ads generate impressions, clicks, comments, reactions. LinkedIn gives you none of it at a person level. Your CEO's thought leader ad might reach the exact 400 VPs of Legal Ops you're targeting and all you'll see in Campaign Manager are engagement metrics. The signal is there. The tool layer that captures it and routes it to a human is what's missing.

The Result

Ad spend finally connects to results. Every dollar you drive into engagement now produces a routed, scored signal to a human who can act, and pipeline you can attribute back to the campaign that sourced it.

> PRO TIP
>
> **LinkedIn's native retargeting needs 300+ engagers before it will serve.** Run Layer 1 TLA spend until your engagement pool clears that floor before you expect the retargeting audience to be usable. The compounding doesn't start on day one. It starts when the pool is deep enough for the algorithm to work with.

The Play

1. **Connect Fibbler to LinkedIn and your CRM.** Fibbler pulls engagement data from the LinkedIn API and matches it to your CRM accounts by domain. No tagging, no scripts on landing pages. Once connected, every account engaging with your campaigns gets logged automatically.
2. **Push engagement to Clay via Fibbler Signals.** Fibbler has a native Clay integration via webhook. Set up a Signal: pick the campaigns to include, the engagement types you care about (impressions, clicks, engagements), and the lookback window. Fibbler pushes the engaging companies to Clay automatically. You get account names at this stage, not contacts yet.
3. **Score and enrich in Clay.** Not every engaging account is equal. Score each account on fit (ICP, segment, tech stack), then enrich with the contacts you actually want to reach (title, role, seniority). Score the contact too if it sharpens the tier.
4. **Route by tier.** Tier 1 fires a Slack alert and an open task in HubSpot for the BDR. Tier 2 gets a personalized DM sent from the AE's profile via HeyReach. Tier 3 goes into nurture.
5. **Close the loop in your CRM.** Fibbler pushes engagement data into HubSpot (or whichever CRM you use), so you can build native reports on influenced pipeline and influenced revenue. Use the Fibbler dashboard or create custom fields in HubSpot to surface engaging accounts and their activity. You can finally answer "did ad spend produce pipeline?" with a number, not a vibe.

The Stack

`LinkedIn Ads` `Fibbler` `Clay` `HeyReach` `HubSpot` `Slack`

PLAY #5 · SIGNAL-BASED OUTBOUND

Signal-Based Outbound.

Listen to what your target accounts and the people inside them are talking about on LinkedIn. Score it. Reach out when it matters.

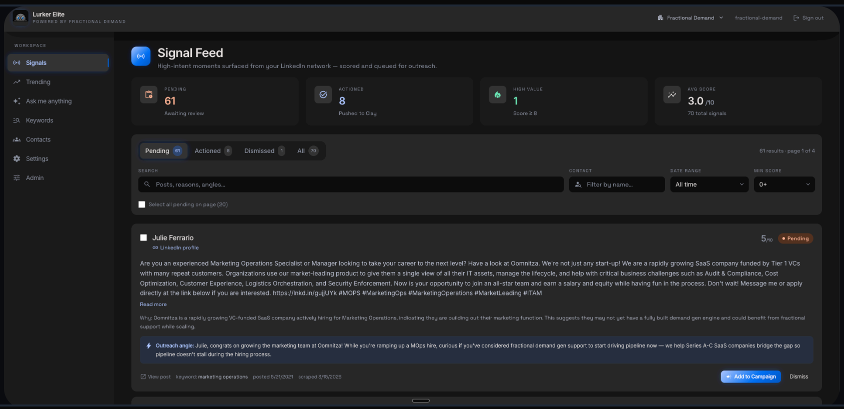

The Problem

LinkedIn is loud. Your buyers are posting, commenting, sharing, and reacting all day. Some of it is about your category. Most of it isn't. Buried in that noise are the conversations that actually matter to your business: the ones where a target account is talking about a problem you solve, a competitor they're frustrated with, a project they just kicked off. Almost nobody surfaces those signals, sorts them, and does anything with the list. The signal is sitting in plain sight.

The Result

Your AEs and SDRs know what their target accounts are talking about on LinkedIn before they pick up the phone. Instead of cold outreach with a generic value prop, you reach out with context, referencing the exact post or topic the prospect surfaced last week. The first line of the email is the prospect's own words.

THE TOOL BEHIND THIS PLAY

We built a proprietary tool we call **Lurkr** to run this play for ourselves and our clients. You give it a list of target accounts and the personas you want to listen for. Lurkr monitors what those people are posting on LinkedIn, scores each post for relevance, and routes the high-fit ones to your team. There are off-the-shelf options out there (Trigify is a good one), but we built our own because it's easier to customize the keyword filters, scoring logic, and routing rules to each client's exact ICP.

The Play

1. **Define your target list.** Decide which accounts you want to listen to (your top 500 ICP accounts, say) and which personas at those accounts (VP Eng, Director of RevOps, CEO, whoever fits your buyer). The list of contacts is the input to everything else.
2. **Monitor their LinkedIn activity nightly.** A scraper pulls every recent post from every contact on the list once a day. You're not watching engagement on your own posts. You're watching what your buyers themselves are saying.
3. **Filter for relevance.** Run keyword and regex filters first to gate out the 70 to 80 percent of posts that aren't about your category. What's left goes to an LLM (Claude or your model of choice) for the actual scoring: 1 to 10 on relevance, plus a one-line summary and a suggested outreach angle.
4. **Score and route.** Posts scoring 6 or higher push downstream into Clay, your CRM, and Slack with the full context attached. Below the threshold, the post stays in the database but doesn't fire an alert.
5. **The SDR opens with the prospect's own words.** "Saw your post last week about the deploy pipeline issues your team's been hitting. Curious what tipped it." Beats any cold opener you can write.

The Stack

`Lurkr` `Clay` `Slack` `Outreach / Salesloft`

Fractional Demand

PLAY #6 · ANONYMOUS VISITOR

Anonymous Visitor → Account Alerting.

Someone at an ICP account just hit your pricing page. Your team should know within 30 seconds.

1 Install anonymization tool

Person-level or company-level ID on visitors who never fill out a form.

STACK

RB2B · Vector · Warmly

→

2 Layer for coverage

Add a company-level tool to catch international and smaller firms.

STACK

Vector · Warmly

→

3 Enrich in Clay

Firmographics, stack, intent, recent job-change signals on every visit.

STACK

Clay

→

4 Score by page + fit

Pricing visit + ICP VP = 5. Blog hit + intern = 0. Weighted rule.

STACK

Clay rules

→

5 Alert in Slack

Score 4-5: BDR task fires. Score 2-3: slow drip. Score 0-1: logs silently.

STACK

Slack + HubSpot

The Problem

A huge chunk of your website traffic is your ICP. They visit pricing, they compare integrations, they download a case study, they leave. No form fill, no chat, nothing. Your sales team has no idea they were ever there. You're buying traffic, converting a fraction of it on forms, and letting the rest evaporate anonymously.

The Result

Your site stops being a black box. Every anonymous ICP visit turns into a routed, scored, context-rich task on a real human's desk.

THE PATTERN MATTERS MORE THAN THE TOOLS

The pattern is: identify → enrich → score → route. Plenty of tools can run each stage. Identify: RB2B, Vector, Warmly, or any anonymization tool. Enrich: Clay or your favorite enrichment stack. The sequence is the play. The tools are the implementation. Start with whatever your current stack already pays for and layer in from there.

The Play

1. **Install an anonymization tool.** RB2B is solid for US-based person-level ID. Vector and Warmly cover broader patterns including company-level matching for international and smaller firms. Pick one based on your traffic profile, or layer two for fuller coverage.
2. **Layer in company-level coverage if you need it.** If your traffic skews international or includes smaller firms, layer a company-level tool on top of RB2B. Combining person-level and company-level identification gets you to 70 to 80 percent of B2B traffic identified. One caveat: international identification carries compliance considerations under GDPR, UK GDPR, and similar privacy laws. Check the tool's geographic coverage and run it past your legal team before you deploy.
3. **Pipe the webhook into Clay.** Every identified visit creates a Clay row. Clay enriches with firmographics, tech stack, intent data, recent job-change signals, plus any data points that are relevant to your product or offering.
4. **Score by page + fit.** A pricing-page visit from a VP Marketing at an ICP-fit company is a 5. A blog visit from an intern at a non-ICP company is a 0. One Clay rule with weighted inputs. No fancy ML.
5. **Alert and log.** Score 4 or 5 fires a Slack alert to the territory BDR. Score 2 or 3 goes into a slow drip. Every identified visit also gets logged to the contact or lead record in your CRM, so the visit history is there next time someone opens the account. Clay can fan all of this out, however you want to wire it.

The Stack

`RB2B` `Clay` `HubSpot` `Slack`

PLAY #7 · JOB-CHANGE + FUNDING SIGNALS

Job-Change + Funding Signals → Warm Outreach.

Your best champion just took a senior role at a new company. Reach out before they pick the next vendor.

1 Build champion list
Export champions, admins, power users from CRM. Load into Clay, watch.
STACK
HubSpot → Clay

→

2 Monitor job changes
Clay watches for employer changes. Webhook fires on flip.
STACK
Clay · UserGems

→

3 Track funding signals
$5M+ rounds at any ICP-fit company. 90-day budget window.
STACK
Crunchbase + Clay

→

4 Sequence both signals
Day 1: founder DM. Day 4: case study. Day 10: meeting ask.
STACK
HubSpot + HeyReach

→

5 Loop in original AE
Slack ping to the AE who worked the champion at their last company.
STACK
Slack + HubSpot

The Problem

Your customer list is the warmest network you have. When a champion leaves Company A and takes a senior role at Company B, there's a tight window where they're picking vendors and rebuilding their stack. Newly hired execs spend most of their budget in the first 100 days on the job.[1] Champions who already used your product are 3 to 5x more likely to convert versus cold outreach.[2] Most teams find out six months later when the new person books a demo. By then they've signed with someone else. The trigger is public. The outreach just has to fire automatically.

The Result

Your pipeline fills itself from the customer base you already built. Net-new logo acquisition starts with a familiar name in the inbox.

PIPELINE LIVES IN CS

Your CS team is sitting on one of your best pipeline sources. They just haven't been told it's pipeline. The "customer reference" and the "new-logo champion" are the same person. If the signal doesn't get captured and the outreach doesn't fire automatically, the window closes quietly and nobody knows it was there.

The Play

1. **Build the champion list.** Export every champion, admin, power user, and decision influencer from the CRM. Load into Clay as a monitored list with ownership attached.
2. **Monitor for job changes.** Tools like Clay or UserGems watch the list automatically. When a champion's employer changes, a webhook fires with the new company and role.
3. **Monitor funding and expansion in parallel.** Crunchbase and other resources via Clay surface new funding rounds at any ICP-fit company. New funding means new hiring and a limited budget window before the next stack gets locked in.
4. **Combine and sequence.** Both signals flow into the same Clay table with the trigger date. A 3-touch sequence auto-fires: Day 1 congrats DM. Day 4 case study email. Day 10 meeting request referencing the new role.
5. **Alert the original AE.** The AE who worked with the champion at their previous company gets a Slack ping too. They're the warmest reach-in. A familiar name in the inbox beats cold outreach 10 times out of 10.

The Stack

`Clay` `UserGems` `Crunchbase` `HubSpot` `Slack`

[1] UserGems / Salesmotion: newly hired execs spend most of their budget in the first 100 days.

[2] UserGems champion-tracking benchmarks: warm champion outreach converts at 3 to 5x cold.

PLAY #8 · COMPETITOR DISPLACEMENT

Competitor Displacement via Tech Stack Data.

Find every company running your competitor. Rank them. Displace them systematically.

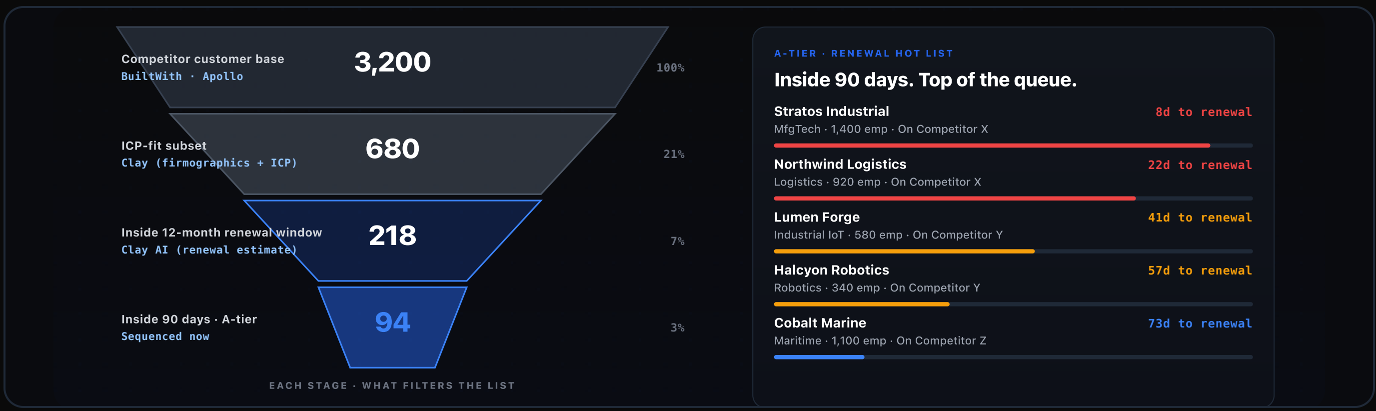

The Problem

Your biggest competitors already did the expensive part. They sold your buyer on the category. Those accounts know the problem, know the tool exists, know how to use it. They picked the wrong product, or the right product at the wrong price. Your displacement list is sitting in a data source you can query this week.

The Result

Competitor displacement stops being a one-off side project and becomes a systematic program with a renewal-aware timeline.

> **PRO TIP**
>
> **Don't message the whole list.** The temptation is to blast every competitor customer in your ICP. The math doesn't math. You want a steady monthly cohort of 40 to 80 accounts hitting the sequence so your BDR can actually personalize each touch. Quality beats volume on displacement every single time. This is the opposite of spray and pray.

The Play

1. **Pull the competitor customer list.** BuiltWith and Apollo both expose tech stack data. Filter to every company running your top 3 competitors. Start narrow.
2. **Enrich and fit-score in Clay.** Firmographics, ARR, headcount, industry against your ICP. Drop anyone who isn't a fit. What's left is your displacement universe.
3. **Estimate the renewal window.** When did the competitor tool first appear in the stack? Most B2B SaaS is on 12 or 24-month terms. A Clay AI prompt estimates the likely window. Accounts inside 90 days go to the top.
4. **Draft the displacement message.** Not "we're better." Instead: the specific limitation of the tool they're on, plus a customer who switched and what they saw. Clay pulls your case study library and matches by industry and company size.
5. **Route by tier and territory.** A-tier (fit + renewal within 90 days) gets personalized high-touch outreach. A note from the founder, the CRO, someone in the C-suite. B-tier goes into a BDR sequence. Every account has a renewal countdown on the opportunity record.

The Stack

`BuiltWith` `Apollo` `Clay` `Claude` `HubSpot` `Outreach`

PLAY #9 · PRE-MEETING BRIEFS

Automated Pre-Meeting Research Briefs.

Every AE walks into every call with a 90-second brief that used to take an hour to write.

The Problem

AEs are supposed to show up prepared. Most don't, because the research is manual and their calendars are full. The best AEs spend 20 to 40 minutes prepping for every first meeting. The worst spend zero. The delta shows up in every call: which competitors got mentioned, which pain points got surfaced, which follow-ups actually landed. Prep quality is the hidden variable most teams never measure.

The Result

Every first call starts with the AE knowing what the prospect cares about, what's on their stack, and which questions open real conversation. With more context up front, reps tailor the conversation to the prospect's actual pain and position your product as the specific answer to it. The floor rises for the whole team, not just your best reps.

> STEAL THIS FOR ANY MEETING TYPE
>
> **Works for expansion meetings, QBRs, and CS check-ins too.** Calendar trigger → enrichment → AI summary → Slack brief. The prompt changes by meeting type. Expansion prep pulls usage data. QBR prep pulls ticket volume and adoption scores. Same pipeline, swapped inputs.

The Play

1. **Pull the meeting list from the calendar.** A Clay webhook watches the AE team's calendars for external prospect meetings. 60 minutes before the call, brief generation fires.
2. **Enrich the company and attendees.** Clay grabs firmographics, recent news, funding, tech stack, hiring signals. Each attendee's LinkedIn, last 3 jobs, recent posts, mutual connections.
3. **Summarize with AI.** Claude runs a structured prompt built off your own template, tuned to your industry and how your product positions. Output is fixed: 3-sentence company summary, top 3 likely pain points, stack relevance, one-sentence attendee bios, 5 open-ended questions. No freeform essays.
4. **Deliver to Slack 30 minutes before.** Each AE gets a DM. Brief is ~250 words, readable in 90 seconds, with a one-click link back to the account or deal record in HubSpot for anyone who wants the full context.
5. **Log to the deal record in HubSpot.** The same brief attaches to the account and opportunity so it's referenceable later, not lost in a Slack thread. Managers can audit prep consistency. The next rep in the deal gets the context without asking.

The Stack

`Google Calendar / Outlook` `Clay` `Claude` `Slack` `HubSpot / Salesforce`

PLAY #10 · ALWAYS-ON EMAIL ENGINE

Always-On Email Engine.

Built once. Always generating revenue.

The Problem

Most lifecycle and nurture programs are set-and-forget. Every customer in a tier gets the same generic email. Every new signup drops into the same onboarding sequence. The work of figuring out which customer types you actually have and what each one needs to hear right now never happens, so the automation runs but nothing resonates. The engineered version segments the audience, maps the message, and lets it run.

The Result

For one fintech client, the engine is generating **$1.2M per month in originations** (in fintech, originations means the money lent or sent to customers). Marketing automation done well. Customers are segmented by type, messaging is mapped to where they are in the lifecycle, and new customers get the right message at the right time. The system runs on autopilot because the segmentation and the message map are right.

PRO TIP

The segments and the message map are the moat, not the copy. Most teams obsess over subject lines and skip the work of figuring out which customer types they actually have and what each one needs to hear. Get the map right and the automation runs itself. The clever copy can come later.

The Play

1. **Map your customer types.** Not just plan tier or company size. Real behavioral segments: who they are, where they are in the lifecycle, what they've activated or bought, what they haven't, what they're likely to need next.
2. **Map messaging to each segment.** What does this customer need to hear right now? Onboarding nudge, expansion offer, cross-sell, win-back, re-engagement. One message map per segment, written for that customer's actual moment.
3. **Build the lifecycle automation in HubSpot.** Triggers fire on lifecycle events: signup, activation, usage threshold, dormancy, account anniversary. Each event drops the customer into the right message track for their segment.
4. **Layer in account context from Clay.** Firmographics, tech stack, funding, hiring signals enrich the segment so the message lands with real context, not generic copy.
5. **Monitor and refine, don't rebuild.** Watch what's resonating per segment. Refresh messaging, add new segments as the customer base evolves. The system produces without quarterly tear-downs.

The Stack

`HubSpot` `Clay` `Outreach`

PLAY #11 · CRM HYGIENE + SIGNAL SCORING

CRM Data Hygiene + Custom-Event Signal Scoring.

The play that decides whether the other ten actually work.

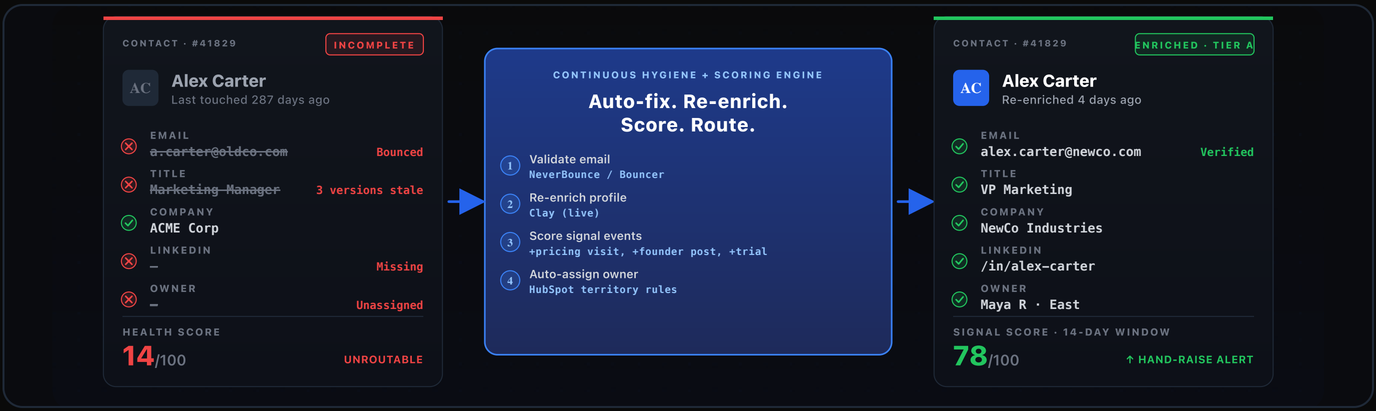

The Problem

Every play in this book falls apart on dirty data. Signal-based outbound fires on bad contacts. Dynamic emails render against junk. Pre-meeting briefs surface the wrong attendee's LinkedIn. Most teams know this and still treat hygiene as a once-a-year cleanup project. The math doesn't math. Hygiene has to run daily or the whole engine degrades.

The Result

The other 10 plays fire on clean data. Signal scoring becomes trustworthy instead of noisy. Routing moves from "title matches ICP" to "behavior plus fit match ICP."

> **THE UNCOMFORTABLE TRUTH**
>
> **This is the least exciting play in the book and the one that decides whether the other ten work.** The teams that build the flashy plays without the hygiene foundation spend year two tearing everything out and rebuilding. Do this first. The rest compounds.

The Play

1. **Define hygiene rules as data, not vibes.** Every object in HubSpot gets a health score: required fields present, email validity, LinkedIn URL valid, last-enriched date under 90 days, owner assigned. Visible on the record.
2. **Auto-fix what's auto-fixable.** Clay watches every new and modified record. Missing LinkedIn URL: auto-append. Stale title: re-enrich. Invalid email: flag and validate via NeverBounce or Bouncer. Runs continuously, not in batches.
3. **Score custom events, not just fields.** Pricing visit, demo request, trial signup, feature usage, third-party intent. Each event has a weight. Accounts get a rolling signal score that decays over time.
4. **Route by signal pattern, not title alone.** A VP Marketing at a 500-person company who just visited pricing twice and engaged with a founder post: one action. Same title + 6-month-old ebook download: different action.
5. **Operate the system weekly.** One RevOps owner. 30 minutes a week to review hygiene scores, event fire rates, and enrichment queue health. Anything red becomes a Monday morning ticket.

The Stack

`HubSpot` `Clay` `NeverBounce / Bouncer` `Custom Workflows`

Fractional Demand

ABOUT FRACTIONAL DEMAND

Your GTM Team. Ready from Day One.

Fractional Demand is a Go-To-Market partner that embeds senior operators across Paid Media and RevOps to build and run a scalable demand engine that accelerates qualified pipeline without the overhead of a full team.

PAID MEDIA

Fractional Paid Media

Senior performance marketers running LinkedIn, Meta, Google, X, and Reddit campaigns tied directly to your pipeline.

REVOPS + GTM ENGINEERING

Fractional RevOps

CRM optimization, data enrichment, automation, and GTM engineering so your paid media, outbound, and sales all work together as one system.

We don't just hand you a playbook. We build the full system: signal-based outbound, CRM enrichment workflows that turn engaged accounts into sales conversations, paid media that compounds the same data, and the reporting infrastructure that connects first click to closed-won. One partner. RevOps and paid media working together.

Why Fractional Demand

Senior-Level Only

Every person on your account has 7+ years of B2B experience. No juniors.

Dedicated Focus

Each manager handles 5 to 7 clients. You get in-house-level attention.

Full-Funnel Thinking

We think about your entire funnel from first click to closed-won.

Ready from Day One

No ramp-up. We plug in, execute, and start driving pipeline immediately.

Get a GTM Team That Builds Pipeline Fast.

Senior-only · Embedded operators · B2B SaaS focus · fractionaldemand.com