Promptifi

FREE PROMPT PACK

25 prospecting prompts that fill *pipeline.*

Pulled from the Promptifi library — 2,900+ tested AI workflows for B2B sales reps. These 25 work as one system: your ICP feeds the signals, signals feed the org map, and the week plans itself around accounts that are actually in motion. Paste any prompt into ChatGPT, Claude, or Gemini, fill in the brackets, and run it.

INSIDE

How to run this pack

01 Fill every {bracket} before you run a prompt — the brackets are the fuel. Vague inputs produce vague outputs; two minutes of real context is what separates these from generic advice.

02 Chain them. These 25 are one system: the early prompts produce briefs, angles, and evidence the later prompts are built to consume. Watch for the chain references inside each prompt.

03 Trust the honesty valves. When a prompt answers WEAK, NO PLAY, or TOO SMALL, that is the prompt working — it just saved you a week of effort on a dead angle.

04 These run as-is in ChatGPT, Claude, or Gemini. Paste the whole prompt, brackets filled, and hold the model to the output format it specifies.

SECTION 01 / 05

Territory and list building

Most prospecting fails at the list, not the message. These five define who deserves your time — their outputs feed everything downstream.

PROMPT 01 · TERRITORY AND LIST BUILDING

The ICP synthesizer

Most ICPs are written from aspiration — who the founder wishes bought — instead of evidence. Building it from your five best and two worst customers grounds it in the only data that doesn't flatter: what actually closed.

> You are an elite go-to-market strategist with deep expertise in reverse-engineering ideal customer profiles from evidence instead of wishful thinking.
> `Inputs` — My product: `{two sentences}`. My five best customers and why they bought: `{list}`. My two worst-fit customers and what went wrong: `{list}`.
> `Method`: (1) Extract what the best five share that the worst two lack — industry, size, motion, trigger, team structure — and rank those traits by how early they're detectable from the outside. (2) Write the ICP as a detection checklist: five observable signals, each with where to find it (site, careers page, LinkedIn, filings). (3) Write the anti-ICP: three signals that predict a bad fit, because knowing who to skip is half the territory. (4) Flag any trait I believe in that the evidence doesn't support — label it FOUNDER FICTION.
> `Constraints`: observable signals only — no 'values innovation.' If my customer sample is too small to generalize, say so and mark which conclusions are provisional.
> `Output`: ranked ICP checklist with detection sources, the anti-ICP, any FOUNDER FICTION flags. I'll paste this into the prompts that follow.

Run this first — its output is the lens every other prompt looks through.

PRO MOVE Re-run it every two quarters; the FOUNDER FICTION flags from six months ago have a way of becoming this quarter's blind spot.

PROMPT 02 · TERRITORY AND LIST BUILDING

The lookalike expander

Your best customer is not a trophy, it's a search query. The discipline is separating why they really bought from the incidentals — the champion's personality doesn't scale, but the structural conditions do.

> You are an expert market cartographer with deep expertise in lookalike expansion — turning one closed-won account into a map of the twenty companies most likely to buy for the same reasons.
> `Inputs` — The customer: `{name, what they do, size}`. Why they really bought: `{the trigger and the pain, in one or two sentences}`. My ICP checklist: `{paste from prompt 01, optional}`.
> `Method`: (1) Decompose why they bought into transferable conditions — the structural facts any lookalike must share — versus incidental facts unique to them. (2) Define the lookalike search: industry adjacencies, size band, business model, and the one non-obvious trait that mattered most. (3) Name the categories and sub-verticals to hunt in, ordered by resemblance, with the search strings or filters I'd use to find them. (4) Describe the tell that separates a true lookalike from a surface match.
> `Constraints`: transferable conditions only — 'their VP was great' doesn't scale. Where you're guessing at an adjacency, say GUESS.
> `Output`: the transferable conditions, the ordered hunting grounds with search strings, the true-lookalike tell.

Run it when a deal closes — your best customer is a search query.

PRO MOVE Run it within a week of every closed-won, while the real reason they bought is still honest in your memory.

PROMPT 03 · TERRITORY AND LIST BUILDING

The signal stack ranker

Signal tools sell coverage; sellers need conviction. Three signals you check weekly beat fifteen you check never — the scorecard exists to give you permission to ignore most of what could theoretically matter.

> You are an elite revenue operations analyst with deep expertise in buying signals — you know most 'intent data' is noise, and that three good signals beat fifteen mediocre ones.
> `Inputs` — My product: `{one sentence}`. My sales cycle length: `{rough}`. Signals I currently watch: `{list them — hiring, funding, tech changes, leadership moves, content, whatever}`.
> `Method`: (1) Score each signal I gave you on three axes: how strongly it predicts a real buying window, how early it fires relative to my cycle, and how cheaply I can monitor it weekly. (2) Kill the signals that score low on two or more axes — name them and say why. (3) Add up to three signals I'm not watching that fit my product, with the monitoring method for each. (4) Assemble the final stack: no more than five signals, each with its check cadence.
> `Constraints`: be ruthless — a signal I can't check in under ten minutes a week is a hobby, not a system. No generic advice; every recommendation tied to my product and cycle.
> `Output`: the scorecard, the kill list with reasons, the additions, the final five-signal stack with cadences.

Run it quarterly — it decides which signals earn a place in your daily scan.

PRO MOVE Put the surviving five signals in your calendar as one 30-minute Friday block — a signal stack without a check cadence is a wish.

PROMPT 04 · TERRITORY AND LIST BUILDING

The list triage engine

A raw list treats every row as equal, and equal treatment is how reps spend March on accounts that were never going to buy. Tiering by evidence — ICP match plus live signal — is the single highest-return sort in prospecting.

> You are an expert list surgeon with deep expertise in triage — you know that treating a raw list as equal rows is how reps burn weeks on accounts that were never going to buy.
> `Inputs` — The list: `{paste rows — company, industry, size, plus whatever fields you have}`. My ICP checklist: `{paste from prompt 01}`. My capacity this week: `{number of accounts I can genuinely work}`.
> `Method`: (1) Sort every row into three tiers — WORK NOW (matches ICP and shows a live signal), WATCH (matches ICP, no signal yet), DISCARD (fails the ICP or trips the anti-ICP) — citing which checklist items drove each call. (2) Within WORK NOW, rank to my capacity number and cut the rest to WATCH. (3) For each DISCARD, one clause of reasoning, so I can audit your judgment and recalibrate the checklist if you're wrong.
> `Constraints`: when data is missing for a row, say UNKNOWN rather than assuming — and tell me the single field worth enriching first. Never inflate WORK NOW to be polite; an honest short list beats a flattering long one.
> `Output`: the three tiers with reasoning, the capacity-cut ranking, the enrichment recommendation.

Run it on any raw list before a single email goes out.

PRO MOVE Audit five of its DISCARD calls by hand the first time; trust built on a spot-check beats blind faith or blind doubt.

PROMPT 05 · TERRITORY AND LIST BUILDING

The disqualification pass

Pipeline that looks full and doesn't move is worse than thin pipeline — it hides the problem and eats the hours that would fix it. The hours-of-hope number turns a vague feeling into a cost you can act on.

> You are an elite qualification skeptic with deep expertise in the most profitable skill in prospecting: walking away early. Your bias is disqualification — accounts must earn their way in.
> `Inputs` — The accounts I'm actively working: `{list with one line of status each}`. My ICP and anti-ICP: `{paste from prompt 01}`. My average deal size and cycle: `{rough numbers}`.
> `Method`: (1) Interrogate each account with three questions: is there evidence of a problem my product solves, evidence anyone there owns that problem, and evidence they can buy in my timeframe? (2) Verdict each: KEEP, PROBATION (name the single proof that would keep it, and the date by which I need it), or CUT. (3) Total the hours per week I'm spending on CUT and PROBATION accounts, using reasonable estimates, so the cost of hope is a number.
> `Constraints`: 'they seemed interested' is not evidence. Be harsher than I would be — I can always overrule you, but I can't get the weeks back. If every account survives, re-run with a higher bar and tell me you did.
> `Output`: verdicts with reasoning, probation proofs and dates, the hours-of-hope number.

Run it when the pipeline looks full but nothing is moving.

PRO MOVE Run it the week before your manager's pipeline review — walking in having already cut the zombies changes the entire conversation.

SECTION 02 / 05

Reading the signals

An account in motion is worth ten at rest. These five decode the public record — filings, hires, leadership moves — into timing.

PROMPT 06 · READING THE SIGNALS

The hiring velocity decoder

A careers page is the most honest public document a company publishes — nobody hires six people for a problem they don't have. Velocity is the tell: backfill is noise, a surge is strategy with a budget attached.

> You are an expert organizational analyst with deep expertise in reading hiring pages the way an investor reads them — as the most honest public statement of what a company is about to attempt.
> `Inputs` — The account: `{name}`. Their current openings: `{paste titles and teams, or the careers page text}`. My product: `{one sentence}`.
> `Method`: (1) Group the openings into initiatives — what are they building, scaling, or fixing, and which teams carry it. (2) Read the velocity: is this steady-state backfill or a surge, and what does the mix of senior versus junior roles say about timing? (3) Identify where my product intersects an initiative — either accelerating it or covering the gap while roles sit unfilled. (4) Name the person most likely to own the initiative, by title, and the quarter the pressure peaks.
> `Constraints`: distinguish signal from churn — replacing one AE is noise; opening a six-person team is strategy. If the openings tell no coherent story, say NO STORY and move on.
> `Output`: the initiatives, the velocity read, the intersection with my product, the likely owner and the pressure quarter.

Run it monthly on tier-one accounts — headcount is strategy you can read for free.

PRO MOVE Screenshot the openings monthly for tier-one accounts; the diff between snapshots is often a clearer signal than any single read.

PROMPT 07 · READING THE SIGNALS

The earnings and filings miner

Executives tell you exactly what they care about four times a year, under legal obligation to be accurate — and almost no rep reads it. Quoting their own stated priority back to them is the cheapest credibility available in outbound.

> You are an elite financial-document analyst with deep expertise in mining earnings calls and filings for the sentences that matter to a seller: stated priorities, admitted problems, and budgeted initiatives.
> `Inputs` — The excerpt: `{paste earnings call section, 10-K risk factors, or shareholder letter}`. My product: `{one sentence}`. My buyer: `{title}`.
> `Method`: (1) Extract every stated priority and admitted pain, quoting the exact language — executives' own words are outreach gold because repeating them back proves you did the reading. (2) Map each to my product: STRONG, STRETCH, or NONE — and discard the stretches, because forced relevance reads as spam to anyone who knows their own business. (3) For the strongest match, write the one-sentence outreach angle that quotes or closely paraphrases their language. (4) Note the fiscal timing: when budgets tied to this priority likely open.
> `Constraints`: quote precisely — a mangled quote is worse than none. If nothing clears STRONG, say NOTHING HERE.
> `Output`: the priority and pain inventory with quotes, the mapping, the angle, the fiscal timing.

Run it on public-company targets before any executive outreach.

PRO MOVE The risk-factors section of the 10-K is the underrated gold: it's the company's own list of what keeps them up at night.

PROMPT 08 · READING THE SIGNALS

The leadership change playbook

New executives arrive with mandates, budgets, and zero loyalty to inherited tools — a 90-day window where the switching cost conversation is easiest. Their previous company's stack is your best predictor of their next moves.

> You are an expert in executive transitions with deep expertise in the 90-day window — new leaders arrive with mandates, budgets, and no loyalty to the tools they inherited.
> `Inputs` — The new executive: `{name, new title, company}`. Where they came from: `{previous company and role}`. My product: `{one sentence}`.
> `Method`: (1) Infer the mandate: why does a company hire this role from that background — what's broken or ambitious enough to justify the hire? (2) Read their history: what tools, motions, or vendors did their previous company visibly use that they may import or replace? Mark each inference with confidence. (3) Place them on the 90-day clock: listening tour, diagnosis, first moves — and name which phase favors my outreach and why. (4) Write the outreach angle that positions my product as ammunition for their mandate, never as homework on their pile.
> `Constraints`: no congratulations openers. If the mandate is unclear, offer two hypotheses rather than one confident guess.
> `Output`: the mandate, the import/replace read with confidence, the timing call, the angle.

Run it within two weeks of a new executive landing at a target account.

PRO MOVE Set a recurring search for title changes at your tier-one accounts; being second to congratulate but first with a relevant angle wins.

PROMPT 09 · READING THE SIGNALS

The expansion signal reader

Growth breaks things — that's not cynicism, it's physics. New markets, offices, and product lines put load on processes built for the old size, and the company that just announced good news is quietly living the strain of it.

> You are an elite growth-signal analyst with deep expertise in expansion moves — you know that growing companies buy differently: faster, messier, and with new problems their current stack wasn't built for.
> `Inputs` — The expansion: `{what they announced or what you observed}`. Company: `{name, size}`. My product: `{one sentence}`.
> `Method`: (1) Name what the expansion breaks: which existing processes, teams, or tools now face load they weren't designed for. (2) Rank those breakage points by how directly my product addresses them. (3) Estimate the window: expansion pain peaks at a predictable distance from the announcement — place my outreach inside it and justify the timing. (4) Write the angle that leads with the breakage, not the congratulations — 'new markets usually strain X' beats 'saw the big news.'
> `Constraints`: if my product doesn't credibly address any breakage point, say WATCH ONLY and give me the future signal that would change the answer. No adjectives about their exciting growth.
> `Output`: the breakage map ranked, the window with justification, the angle.

Run it when an account opens an office, enters a market, or launches a product line.

PRO MOVE Lead the outreach with the breakage, never the congratulations — 'new regions usually strain X' lands where 'saw the big news' dies.

PROMPT 10 · READING THE SIGNALS

The competitor-customer hunter

A competitor's customer already believes in the category — the hardest part of the sale is pre-done. The discipline is honest energy assessment: a price hike creates switching motion, a mildly bad quarter creates only schadenfreude.

> You are an expert displacement strategist with deep expertise in competitor-customer prospecting — the highest-converting list in outbound, because these buyers already believe in the category.
> `Inputs` — The competitor: `{name}`. What just happened to them: `{price hike, outage, acquisition, layoffs, sunset}`. My product's honest edge over them: `{one or two sentences}`. How I can find their customers: `{review sites, case studies, job postings mentioning the tool, communities}`.
> `Method`: (1) Assess the event honestly: does it create real switching energy, or just schadenfreude? If the latter, say LOW ENERGY and stop. (2) Profile which of their customers feel this event most — by size, use case, or contract timing — because pain is never uniform. (3) Build the finding method: the three fastest ways to assemble a list of affected customers from my available sources. (4) Write the outreach angle that names the situation without gloating — respectful of their original choice, concrete about what's now different.
> `Constraints`: never trash the competitor; the reader chose them once. No fake urgency.
> `Output`: the energy verdict, the most-affected profile, the list-building method, the angle.

Run it when a competitor stumbles — pricing change, outage, acquisition, bad quarter.

PRO MOVE Build the affected-customer list the same week the event breaks; switching energy decays in about a month.

SECTION 03 / 05

Mapping the org

Deals die when you're single-threaded into the wrong person. These five map power before you spend a single touch.

PROMPT 11 · MAPPING THE ORG

The org chart inferencer

Deals stall in the gap between who you know and who decides — and that gap is usually visible from public fragments before your first call. Marking every link OBSERVED or INFERRED keeps you strategic instead of confidently wrong.

> You are an elite organizational cartographer with deep expertise in reconstructing reporting structures from public fragments — LinkedIn titles, team pages, conference bios, press quotes.
> `Inputs` — The company: `{name, rough size}`. The function I sell into: `{department}`. Fragments I've collected: `{paste names, titles, anything}`. My product: `{one sentence}`.
> `Method`: (1) Reconstruct the likely structure of the target function: who reports to whom, where the budget probably sits, and which titles are decoration versus power — marking every link as OBSERVED or INFERRED. (2) Identify the three roles that matter for my deal: the person with the pain, the person with the budget, and the person with veto power — they are rarely the same person. (3) Name the gaps: which people I haven't found yet whose existence the structure implies. (4) Recommend the entry point and the sequence — who first, who second, and why that order.
> `Constraints`: title inflation varies wildly by company size — calibrate power to headcount, not title words. Where the structure is genuinely unknowable from my fragments, say so.
> `Output`: the mapped structure with OBSERVED/INFERRED tags, the three roles, the gaps, the entry sequence.

Run it before choosing your entry point at any account over ~200 people.

PRO MOVE Update the map after every call; each conversation confirms or kills an inference, and a living map compounds.

PROMPT 12 · MAPPING THE ORG

The champion identifier

The most expensive mistake in enterprise sales is mistaking a fan for a champion — fans like you, champions spend internal capital on you. The three tests exist because enthusiasm is free and calendar invites aren't.

> You are an expert in champion dynamics with deep expertise in the difference between a fan and a champion: fans like you; champions spend internal capital on you. Confusing them kills deals.
> `Inputs` — My contacts and every interaction so far: `{paste names, titles, and what each has said or done}`. What my product changes for this company: `{one sentence}`.
> `Method`: (1) Score each contact on the three champion tests: do they personally win if this succeeds, can they access power, and have they already done anything that cost them effort or risk on my behalf? Cite the behavioral evidence for each score — words are cheap, calendar invites aren't. (2) Verdict each: CHAMPION, FAN, BLOCKER-IN-DISGUISE, or UNKNOWN. (3) For the best candidate, design the champion test: a small, reasonable ask that requires internal effort — and name what a pass and a fail each look like. (4) If nobody clears the bar, say NO CHAMPION and name which contact is closest and what's missing.
> `Constraints`: enthusiasm is not evidence. Weight actions ten times heavier than words.
> `Output`: the scorecard with evidence, verdicts, the champion test, the honest gap if there is one.

Run it once you have two or more contacts showing any engagement.

`PRO MOVE` Run the designed champion test within a week; an untested champion in your forecast is a guess wearing a title.

PROMPT 13 · MAPPING THE ORG

The power-line tracer

Org charts show reporting lines; money moves along different wires. Deals die in the invisible layer — finance, security, procurement — that activates at your price point whether or not anyone mentioned it.

> You are an elite deal strategist with deep expertise in tracing economic power — you know the org chart shows reporting lines, but money flows along different wires.
> `Inputs` — The deal: `{what I'm selling, rough price}`. The org as I understand it: `{paste from prompt 11 or describe}`. Who I'm currently talking to: `{names, titles}`. Any clues about budget: `{paste anything said about approval, procurement, spending}`.
> `Method`: (1) At this price point and company size, name where approval power almost certainly sits — and compare it to where I'm currently connected; the gap is my risk. (2) Trace the likely approval chain: recommender, approver, and the silent stakeholders who can slow or kill it — finance, security, procurement, legal — with a guess at which ones activate for this deal size. (3) For each gap between my contacts and the chain, name the bridge: who I know that can introduce, or what artifact (business case, security doc) travels where I can't. (4) Flag the single most dangerous unknown.
> `Constraints`: calibrate to deal size honestly — a $5k deal doesn't need a CFO strategy, and pretending it does wastes everyone's time.
> `Output`: the power location versus my position, the chain, the bridges, the dangerous unknown.

Run it before any proposal goes upstairs — know whose money it actually is.

PRO MOVE Ask your champion 'who would be annoyed to learn about this late?' — it maps veto power faster than any org question.

PROMPT 14 · MAPPING THE ORG

The multi-thread architect

Single-threaded deals die of ordinary things — a resignation, a reorg, a vacation at the wrong time. Multi-threading isn't networking for its own sake; it's making your deal survivable by any one person's bad month.

> You are an expert in deal resilience with deep expertise in multi-threading — building enough relationships inside an account that no single departure, reorg, or bad day can kill the deal.
> `Inputs` — Current threads: `{who I talk to and how warm each is}`. The org map: `{paste from prompts 11–13 or describe}`. Deal stage: `{where we are}`. My team: `{anyone on my side who could own a thread — SE, exec, CSM}`.
> `Method`: (1) Audit my thread coverage against the three roles that matter — pain owner, budget owner, veto holder — and name which are covered, thin, or absent. (2) For each gap, design the opening move: the natural, non-forced reason for contact — a relevant artifact, a peer-to-peer exec match, a technical validation — because 'looping in' with no pretext reads as going around someone. (3) Sequence the moves so my current contact feels included, not circumvented: what I say to them about each new thread. (4) Assign each thread an owner on my side.
> `Constraints`: never open a thread behind my champion's back without naming the risk first. Two strong threads beat five limp ones.
> `Output`: the coverage audit, the opening moves, the sequencing script, the ownership map.

Run it the moment a deal gets real — single-threaded is one resignation from dead.

PRO MOVE Give every new thread a real artifact as its pretext — a benchmark, a security doc — 'looping in' with nothing in hand reads as going around someone.

PROMPT 15 · MAPPING THE ORG

The warm path finder

A warm introduction converts at multiples of cold outreach, yet most reps check for one only after cold fails. The forwardable ask is the whole trick: every word the connector must write themselves halves the odds they send it.

> You are an elite relationship strategist with deep expertise in path-finding: the discipline of checking every warm route into an account before defaulting to cold outreach.
> `Inputs` — Target account and buyer: `{company, name, title}`. My assets: `{customers in the same industry, investors, partners, former colleagues, communities I belong to, my company's advisors}`. What I can offer a connector: `{the favor economy runs both ways}`.
> `Method`: (1) Enumerate every plausible path from my assets to the target: shared employer history, customer-to-peer intro, investor or board overlap, partner channels, community co-membership. (2) Score each path on strength of the connector's relationship to the target and their willingness to spend capital for me. (3) For the best path, write the exact forwardable ask: three sentences the connector can send without editing — because every word they have to write themselves halves the odds they send it. (4) Define the fallback: if no path clears the bar, say GO COLD and don't manufacture a weak intro, which burns the connector and the account.
> `Constraints`: never ask a connector to endorse what they haven't seen. The forwardable ask must make the connector look good, not just me.
> `Output`: the path inventory scored, the forwardable ask, the verdict.

Run it before going in cold — a warm path converts at multiples of cold.

PRO MOVE Spend ten minutes on paths before any tier-one cold sequence — the map usually surprises you at least once a month.

SECTION 04 / 05

Prioritization and timing

Effort is finite; accounts aren't equal. These five decide where this week's hours actually go.

PROMPT 16 · PRIORITIZATION AND TIMING

The buying-signal stack ranker

Monday morning is where prospecting weeks are won or lost, and 'who feels promising' is not a ranking method. Scored signals with visible arithmetic turn the queue into something you can trust — and audit when it's wrong.

> You are an elite pipeline mathematician with deep expertise in stack-ranking accounts by evidence of buying motion, not by which logo would look best on a slide.
> `Inputs` — The accounts: `{list with whatever signal data you have per account — hires, news, engagement, tech changes}`. My signal stack and weights: `{paste from prompt 03, or let me weight}`. My capacity: `{accounts I can work well this week}`.
> `Method`: (1) Score every account against the signal stack, showing the arithmetic per account so I can audit it. (2) Rank, then cut at my capacity line — everything below the line gets a WATCH tag and the specific signal that would promote it. (3) For the top three, state the reason-to-reach-out-this-week in one sentence each — the signal, not the wish. (4) Name any account I'm emotionally attached to that the math demoted, so the bias is visible.
> `Constraints`: show your work — a ranking I can't audit is a ranking I won't trust next week. Ties break toward recency of signal.
> `Output`: the scored ranking with arithmetic, the capacity cut, the top-three reasons, the attachment flag.

Run it Monday morning — it turns the account list into a ranked queue.

PRO MOVE Name the emotionally-attached account out loud each week; the bias you can see is the bias that can't steer.

PROMPT 17 · PRIORITIZATION AND TIMING

The territory heat map

Account-level grinding hides territory-level patterns — you can work hard all quarter in the wrong pond. Portfolio thinking asks the fund manager's question: where is effort concentrated, and is that where returns are?

> You are an expert territory strategist with deep expertise in portfolio thinking — seeing a patch of accounts the way a fund manager sees positions: concentration, coverage, and neglect.
> `Inputs` — My territory: `{industries, geography, size band}`. My accounts by status: `{rough counts — active, watched, untouched}`. Where my last two quarters of wins came from: `{describe}`. Quota context: `{what I need this year}`.
> `Method`: (1) Map the heat: which segments of my territory produce wins, which consume effort without producing, and which I've never seriously touched. (2) Diagnose the pattern — am I over-fishing one pond because it's comfortable? Quantify the imbalance. (3) Recommend the reallocation: what percentage of prospecting hours moves where, and what I stop doing to fund it. (4) Name the experiment: one untouched segment worth a two-week structured test, with the result that would justify continuing.
> `Constraints`: recommendations in hours and percentages, not sentiments. If my win data is too thin to diagnose, say THIN DATA and design the tracking that fixes it.
> `Output`: the heat map in words, the imbalance number, the reallocation, the experiment.

Run it monthly — the territory-level view your account-level grind can't see.

`PRO MOVE` Fund every experiment by naming what you stop doing — reallocation without subtraction is just a longer to-do list.

PROMPT 18 · PRIORITIZATION AND TIMING

The trigger calendar builder

Some buying windows are printed on the calendar months ahead — fiscal year ends, renewal anniversaries, planning seasons — and still surprise reps every year. Back-scheduling by your cycle length turns timing from luck into logistics.

> You are an elite timing strategist with deep expertise in predictable triggers: fiscal year ends, renewal anniversaries, planning seasons, conference cycles — the buying windows you can schedule outreach against months ahead.
> `Inputs` — My top accounts: `{list}`. Known dates: `{fiscal year ends, contract anniversaries, industry events — whatever you have}`. My product: `{one sentence}`. My sales cycle: `{rough length}`.
> `Method`: (1) For each account, place every known and inferable date on a timeline — and infer the ones I don't have: public companies telegraph fiscal years; industries share planning rhythms. Mark inferred dates as such. (2) Back-schedule from each window by my sales cycle length: the date outreach must start for the window to be reachable. (3) Compress into a quarter-view calendar: which accounts get touched in which weeks, and why. (4) Flag the windows I've already missed this cycle, with the re-entry date for next cycle.
> `Constraints`: an inferred date is a hypothesis — pair each with the question that would confirm it on a first call.
> `Output`: the per-account timeline, the back-scheduled start dates, the quarter calendar, the missed-window list.

Run it quarterly — some buying windows are printed on the calendar in advance.

PRO MOVE For every inferred date, ask the confirming question on your next first call; a confirmed fiscal calendar is territory gold.

PROMPT 19 · PRIORITIZATION AND TIMING

The graceful deprioritizer

Persistence closes deals; stubbornness burns quarters — and from the inside they feel identical. The evidence-versus-attachment split is the tiebreaker, and the exit touch matters because parked accounts come back to reps who left well.

> You are an expert in strategic patience with deep expertise in the deprioritization decision — the difference between persistence that closes deals and stubbornness that burns quarters.
> `Inputs` — The account: `{name}`. The history: `{touches, responses, how long}`. Why I keep trying: `{honest answer}`. What working it costs me weekly: `{rough hours}`.
> `Method`: (1) Separate the evidence from the attachment: list what has actually happened versus what I keep hoping will happen. (2) Verdict: PERSIST (name the concrete evidence that justifies it), PARK (define the exact signal that reactivates it and set the watch), or RELEASE. (3) If PARK or RELEASE, write the graceful exit touch — the final message that leaves the door open with dignity: no guilt, no 'last attempt,' just a clean close with a standing offer. (4) Redirect the recovered hours: name where they go, from my WATCH list, so the decision creates something instead of just ending something.
> `Constraints`: 'big logo' is attachment, not evidence. The exit touch must be one I'd be glad they screenshot.
> `Output`: evidence versus attachment, the verdict with reasoning, the exit touch if applicable, the redirect.

Run it when an account keeps absorbing touches without progressing.

PRO MOVE Write the exit message you'd be glad they screenshot — parked is a relationship state, not a failure state.

PROMPT 20 · PRIORITIZATION AND TIMING

The re-entry timing engine

Closed-lost is the cheapest pipeline you'll ever get — the education is paid for, the need was real, and something has changed since. Sorting losses by why they ended tells you exactly when each becomes workable again.

> You are an elite re-engagement strategist with deep expertise in closed-lost mining — the cheapest pipeline in the business, because the education is already paid for.
> `Inputs` — The closed-lost or gone-cold accounts: `{list with why each ended, as honestly as I know it}`. Time since each ended: `{rough}`. What's changed on my side since: `{new features, pricing, proof points}`.
> `Method`: (1) Sort the losses by reason into re-entry classes: TIMING (highest value — they said not yet), PRODUCT GAP (re-enter only if I've closed the gap — check against my changes), COMPETITOR WIN (re-enter near their renewal anniversary), and BAD FIT (leave dead). (2) For each live class, compute the re-entry date from the ending date and the class logic. (3) Write the re-entry angle per account: it must acknowledge the history in one graceful clause and lead with what's different now — never pretend the past didn't happen. (4) Rank all re-entries into one queue by expected value.
> `Constraints`: if nothing has changed on my side, the TIMING class is the only honest re-entry — say so. No 'checking in.'
> `Output`: the classed losses, the dated queue, the per-account angle.

Run it on closed-lost and gone-quiet accounts every quarter — the second bite converts.

PRO MOVE Mine the TIMING class first every quarter; 'not yet' accounts convert at rates cold lists dream about.

SECTION 05 / 05

The prospecting operation

Prospecting compounds when it's a system, not a mood. These five run the machine — daily blocks, clean data, and the patterns your own history is trying to teach you.

PROMPT 21 · THE PROSPECTING OPERATION

The daily block architect

Prospecting doesn't fail from laziness — it fails from interruption, one Slack at a time. Blocks with assigned accounts remove the two decisions that kill momentum: when to prospect and who to contact.

> You are an elite sales productivity architect with deep expertise in time design — you know prospecting dies by interruption, and that reps who block it like a meeting outproduce reps who fit it in.
> `Inputs` — My ranked queue: `{paste from prompt 16}`. My calendar's fixed commitments this week: `{describe}`. My honest energy pattern: `{when I'm sharp, when I'm useless}`. Daily activity expectations: `{calls, emails, whatever my org demands}`.
> `Method`: (1) Place the prospecting blocks where my energy is highest and my calendar can protect them — research in one block, execution in another, because switching between them is where the hour leaks. (2) Assign specific accounts from the queue to specific blocks, so no block starts with 'who should I contact?' (3) Build the interruption rule: what qualifies as urgent enough to break a block, and what waits. (4) Set the day's done-line: the number that means I can stop without guilt — sustainable beats heroic.
> `Constraints`: no block over 90 minutes. If my fixed commitments make the plan impossible, say IMPOSSIBLE and show me the smallest calendar change that fixes it.
> `Output`: the block schedule with assigned accounts, the interruption rule, the done-line.

Run it Sunday night or Monday morning — the week gets designed once, then executed.

PRO MOVE Separate research blocks from execution blocks; the context switch between them is where the hour quietly leaks.

PROMPT 22 · THE PROSPECTING OPERATION

The CRM note distiller

The note you write today is a letter to a stranger — you in ninety days, with no memory of the call. Four fields, with facts separated from interpretations, is the difference between a CRM and a diary.

> You are an expert sales data hygienist with deep expertise in the note that matters: the one that lets anyone — including you in three months — pick up the account cold and act intelligently in sixty seconds.
> `Inputs` — My raw notes, voice-memo transcript, or brain dump from today's touches: `{paste it, messy is fine}`.
> `Method`: (1) Per account touched, distill to the four fields that matter: what happened (fact, one line), what it means (interpretation, one line), next action with a date, and the risk or unknown worth remembering. (2) Strip everything that won't matter in two weeks — the weather, the pleasantries, my mood. (3) Flag any note where I recorded an interpretation as a fact — 'they're interested' versus 'she asked for pricing by Friday' — and correct the record to the observable. (4) Surface any commitment I made that isn't yet on my task list.
> `Constraints`: every next action gets a date or it's a wish. Preserve exact quotes where I captured them — the prospect's own words outrank my summary.
> `Output`: the per-account four-field notes, ready to paste, plus the commitments list.

Run it at day's end — thirty seconds per account, so future-you inherits intelligence.

PRO MOVE Dictate raw notes in the parking lot and let the distiller structure them — capture beats composition every time.

PROMPT 23 · THE PROSPECTING OPERATION

The win-loss pattern miner

The best prospecting advice you'll ever receive is encoded in your own last two quarters — where winners entered, what triggered them, where losers stalled. The PATTERN-versus-ANECDOTE tag is what keeps three coincidences from becoming a strategy.

> You are an elite pattern analyst with deep expertise in mining a rep's own win-loss history for the strategy hiding in it — because the best prospecting advice you'll ever get is encoded in your last two quarters.
> `Inputs` — My recent history: `{list wins and losses with whatever I know: source, industry, size, who I entered through, cycle length, why it ended}`.
> `Method`: (1) Hunt for patterns across five dimensions: where winners came from, who I entered through, what trigger preceded them, how fast they moved, and where losers stalled. (2) Rank the patterns by confidence given my sample size — and be honest about which are real signal versus small-number noise, marking each PATTERN or ANECDOTE. (3) Translate the top two PATTERNS into prospecting rules: specific changes to who I target and how I enter. (4) Name the one thing I keep doing that the data says isn't working.
> `Constraints`: never build a rule on fewer than three consistent cases. If the whole sample is too small, say so and specify what to log going forward so next quarter's version of this prompt has teeth.
> `Output`: the pattern inventory with PATTERN/ANECDOTE tags, the two rules, the stop-doing flag.

Run it quarterly on your own history — your funnel has been leaving you notes.

PRO MOVE Log the five dimensions religiously for one quarter; the miner is only as sharp as the history you feed it.

PROMPT 24 · THE PROSPECTING OPERATION

The referral engine

Referral asks fail from vagueness and timing, not from customer unwillingness — 'anyone who might benefit' converts at nearly zero because it's unanswerable. A named profile, asked at the goodwill peak, is a different sport.

> You are an expert in referral mechanics with deep expertise in the ask most reps fumble: specific, easy, and timed to the moment of maximum goodwill.
> `Inputs` — The happy customer: `{name, company, what outcome they got}`. When the goodwill peaked: `{go-live, big result, renewal, kind email}`. Who I'm hoping they know: `{the profile from my ICP, or specific names from their network if I've looked}`.
> `Method`: (1) Time-check: is the goodwill moment still warm? Referral asks decay fast — if it's stale, name the next natural moment to wait for instead. (2) Sharpen the ask from 'anyone who might benefit' — which converts at nearly zero — to one or two specific profiles or names, because specific asks are answerable. (3) Write the ask message: short, explicitly permission-based ('would you be open to...'), with the forwardable blurb included so their effort is one forward, not an essay. (4) Write the thank-you-regardless follow-up, because the relationship outranks the referral.
> `Constraints`: never make the customer feel like a lead source. One ask per moment.
> `Output`: the timing verdict, the sharpened ask, the message with forwardable blurb, the follow-up.

Run it within two weeks of any closed-won or any delighted-customer moment.

`PRO MOVE` Include the forwardable blurb every single time; making their effort one forward instead of one essay is the whole conversion.

PROMPT 25 · THE PROSPECTING OPERATION

The quarterly territory review

Territories drift — what worked in Q1 quietly stops, and without a written review you keep doing it through Q3. Grading execution failures separately from strategy failures matters because they have opposite fixes.

> You are an elite sales operator with deep expertise in the quarterly review — the discipline of interrogating a territory like a business unit: what worked, what didn't, and what changes, in writing, before the new quarter's noise begins.
> `Inputs` — Last quarter's numbers: `{pipeline created, meetings booked, win rate, whatever I track}`. My rules from last review: `{what I said I'd do differently}`. The pattern-miner output: `{paste from prompt 23 if run}`. This quarter's target: `{number}`.
> `Method`: (1) Grade last quarter's rules honestly: which did I actually follow, and of those, which moved a number? Separate execution failures from strategy failures — they have different fixes. (2) Work the math backward from this quarter's target: given my conversion rates, how many accounts, touches, and meetings does the target imply per week? State whether that math is achievable or requires a rate improvement, and which rate is the cheapest to improve. (3) Set no more than three rules for the new quarter, each with a number attached and a weekly check. (4) Write the one-paragraph territory thesis: where this quarter's pipeline will come from and why.
> `Constraints`: three rules maximum — a ten-point plan is a zero-point plan. Every claim gets a number or a date.
> `Output`: the graded rules, the backward math with the verdict, the three new rules, the thesis.

Run it in week one of each quarter — the operating review your territory deserves.

PRO MOVE Keep every quarter's thesis paragraph; four of them in a row is a promotion case that writes itself.

THE FULL LIBRARY

This is 25. The library has *2,900+.*

Every stage of the deal — prospecting, discovery, objections, negotiation, closing, renewals — curated and tested for AEs, SDRs, and BDRs. New workflows added continually. $4.99 a month, or $99 once for lifetime access.

Browse the library

PROMPTIFI.CO