

The AI Marketing Ops Manual

The execution companion to the Growth Tip Sheet - 15 plays with numbered steps and copy-paste prompts: setup, the weekly content engine, scheduled automations, launch and outreach, quality control. One hour of setup, then about 90 minutes a day.

PERCEPTION



REALITY

Ben Wynn · CMO, GOAT Network

X @0x1164 · Telegram @OX1164

This is the execution companion to the Growth Tip Sheet. The tip sheet tells you what works and why, with sources. This manual tells you exactly how to run it - as 15 plays, each with a goal, a time estimate, numbered steps, and a copy-paste prompt.

Two rules before you start:

- 1. You are the editor.** Every prompt here produces a draft, not a publication. You review, you decide, you own the output. Never publish an AI-generated statistic or claim without checking it yourself.
- 2. Do the setup plays first.** Plays 1-3 create the context files that every other play depends on. One hour of setup makes every future task faster and dramatically better.

Works with any capable AI assistant (Claude, ChatGPT, or similar). Where a play needs web browsing or scheduling, it says so.

Part 1 - One-time setup (about an hour, do this week)

Play 1: Write your Positioning Doc

Goal: one short document that tells any AI (and any human) exactly what your project is. **Time:** 25 minutes. **You'll reuse this in every other play.**

Steps: 1. Copy the template below into a document. 2. Fill in every bracket. Keep answers short and concrete - one to three sentences each. 3. Save it as `positioning.md`. Update it whenever your positioning changes.

TEMPLATE - Positioning Doc

```
PROJECT: [name - use exactly one spelling everywhere]
EDGE: [the one thing this project can credibly win on - product, brand association,
      distribution, community, speed, focus or trust - in one sentence]
ONE-LINER: We help [specific person] achieve [concrete outcome] without [painful thing].
WHO, EXACTLY: [the specific person - role, situation, where they hang out online]
THEIR PAIN: [what costs them time or money today, in their words]
WHY US: [what we can credibly claim that competitors can't]
PROOF: [numbers, users, results, partners - only things that are true and checkable]
ACTIVATION MOMENT: [the one completed task that means a new user "got it"]
COMPETITORS: [2-4 names and one line on how we differ from each]
LINKS: [site, docs, X handle, community]
```

If you can't fill a bracket, that's useful information - it means the gap is clarity, not marketing. Fix the bracket before running any content plays. The EDGE line matters most: every other play in this manual communicates or compounds it.

Play 2: Write your Brand Guardrails

Goal: a dos-and-don'ts file that keeps every AI draft on-brand, so your review takes two minutes instead of twenty. (This is how Anthropic's own growth team reduces human review on automated work - Tip Sheet §13.) **Time:** 20 minutes.

Steps: 1. Copy the template. Fill it in honestly - the "never" list matters most. 2. Save as `guardrails.md`.

TEMPLATE - Brand Guardrails

```
VOICE: [3 adjectives, e.g. plain, direct, technical-but-friendly]
WE SOUND LIKE: [one example sentence in your real voice]
WE NEVER SOUND LIKE: [one example of what to avoid, e.g. hype-speak]
ALWAYS: [e.g. concrete numbers over adjectives; short sentences; show the product]
NEVER: [e.g. superlatives without a source; price/token talk; jargon on the homepage;
       em dashes if you hate them; claims we can't verify]
FORMATTING: [e.g. hyphens only; links go in the first reply on X; captions on video]
DISCLOSURE: [e.g. sponsored content is always labelled; AI-assisted drafts are fine,
             fabricated quotes or data are never fine]
```

Play 3: Set up your AI workspace

Goal: stop re-explaining your project in every chat. **Time:** 10 minutes.

Steps: 1. In your AI tool, create a dedicated project/space for marketing. 2. Add `positioning.md` and `guardrails.md` to it (as project files, custom instructions, or simply by pasting both at the top of each session - whatever your tool supports). 3. Save this manual's prompts somewhere you can copy from in seconds - a notes file is fine. Treat the prompt library like code: when a prompt produces something good, save the improved version.

Part 2 - The weekly content engine

Play 4: One source asset, five outputs

Goal: run the create-once-publish-five-times pipeline (Tip Sheet §11) in about 30 minutes instead of a full day. **When:** once a week, after you've shipped or demoed something real.

Steps: 1. Capture a source asset: a demo recording transcript, your build notes, a real user result - or the strongest entry in your idea inbox (Play 14). Rough is fine. 2. Paste the prompt below with your context files and the asset. 3. Review the five drafts against your guardrails. Edit in your voice. Publish across the week, not all at once - front-load the strongest pieces early-to-midweek (Tue-Wed peak, never weekends), timed to your audience's timezone; for US + Europe reach that's early-to-mid afternoon UTC (Tip Sheet §1.3).

PROMPT - Content Engine

Here is my positioning doc and brand guardrails: [paste both]
Here is this week's source asset: [paste transcript / notes / result]

Produce five drafts from this one asset:

1. X THREAD (3-5 posts): first post must work standalone with a strong hook and NO link. End post 1 with "Link below". Post 2 is a reply carrying the link.
2. 30-SECOND VIDEO SCRIPT for a screen recording: outcome shown in the first 3 seconds, caption text included, URL spoken and shown on the end card.
3. BLOG/DOCS POST (500-800 words): question-style H2 headings, the direct answer in the first 60 words, at least two concrete numbers from the asset, and a 3-question FAQ section at the end.
4. NEWSLETTER SECTION (100-150 words): what shipped, why it matters, one link.
5. FAQ/DOCS ENTRY: the single most useful Q&A from this asset, written so an AI search engine could quote it verbatim.

Rules: follow my guardrails exactly. Use only facts from the asset and positioning doc - do not invent numbers. Adapt tone to each format; do not reuse identical sentences across outputs.

Play 5: The hook factory

Goal: the first line is 90% of the post - generate ten, pick one. **When:** before publishing anything on X.

PROMPT - Hooks

Context: [paste positioning doc]. I'm posting about: [one sentence].
Write 10 first lines for an X post. Mix these shapes: a specific number/result, a contrarian take, a before/after, an open loop, a direct question to my exact target user. Max 12 words each. No hashtags, no emojis, no jargon.
Then mark which 2 you'd bet on and say why in one line each.

Pick the one that would stop *you* mid-scroll (Tip Sheet mindset: study yourself). If none would, run it again with feedback.

Play 6: The link-in-reply formatter

Goal: never lose reach to the link penalty again (Tip Sheet §1.1). **When:** any post that needs a link.

PROMPT - Link-in-Reply

Take this draft post: [paste]
Restructure it as:
POST 1 - hook + body, NO link anywhere, ending with "Link below" (or a natural equivalent). If there's a video, note that the URL belongs on the end card.
POST 2 - a one-line self-reply carrying the link with a reason to click.
Keep my guardrails: [paste guardrails]

Part 3 - Scheduled automations

Play 7: The morning watchdog

Goal: your metrics reviewed daily without you staring at dashboards (Tip Sheet §13 - this is the exact pattern Anthropic's head of growth runs on 20-25 charts every morning). **Time:** 15 minutes to set up, then ~2 minutes a day.

Steps: 1. Decide your inputs: your analytics export, a screenshot of your dashboard, or pasted numbers for the funnel (visitors, signups, activated, retained, referring) plus X analytics. 2. If your AI tool supports scheduled/recurring tasks, schedule the prompt below each morning. If not, set a daily calendar reminder and paste it manually - it still takes two minutes. 3. Act only on what it flags. Keep eyeballing your two or three favourite numbers yourself.

PROMPT - Watchdog

```
Here are today's numbers: [paste or attach]
Yesterday's / last week's for comparison: [paste]
My north-star metric is: [metric].

Report in exactly this format:
1. NORTH STAR: current value, change, one-line read.
2. NEEDS ATTENTION: anything that moved >20% or broke a trend - what and
   one plausible why. If nothing, say "Nothing needs attention."
3. FUNNEL LEAK: the single stage losing the most people right now.
Do not summarise everything. Flag only what deserves my attention.
```

Play 8: The weekly GEO monitor

Goal: know what the AI engines say about you, and fix wrong answers with content (Tip Sheet §3.3). **When:** same time every week, 15 minutes.

Steps: 1. Ask ChatGPT, Claude, Perplexity and Gemini the same four questions: - "What are the best tools for [your category]?" - "What is [your project name]?" - "How does [your project] compare to [main competitor]?" - "[The exact question your target user would ask]" 2. Log the answers in a simple table: date, engine, mentioned yes/no, accurate yes/no, sources cited. 3. Every wrong or missing answer is a bug. Fix bugs with content: a docs page, an FAQ entry, a comparison page - then feed the gap into Play 4 as next week's asset.

PROMPT - GEO Gap Fix

```
An AI engine answered "[paste the wrong/missing answer]" when asked
"[the question]". Here's my positioning doc: [paste].
Draft the single page most likely to fix this: question-style title, direct
answer in the first 60 words, 2-3 verifiable facts with numbers, an FAQ block,
and suggested schema markup type. Write for extraction, not for cleverness.
```

Play 9: The monthly competitor teardown

Goal: know exactly what rivals' onboarding and pricing do better, without spending an afternoon on it. **When:** monthly. **Needs:** an AI with web browsing.

PROMPT - Teardown

Visit [competitor URL] and walk their signup/onboarding flow and pricing page. Report: 1) their one-liner and who it targets, 2) every onboarding step and where friction is added or removed, 3) what their first-run experience routes new users toward, 4) pricing structure and what it signals, 5) the 2 things they do better than my flow [describe or link yours], 6) the 1 thing I should copy this month and the 1 trap to avoid. Be specific. Quote their actual wording where relevant.

Part 4 - Campaign plays

Play 10: Launch week runbook

Goal: stack every free amplifier into the same week so each multiplies the others (Tip Sheet §12).

When: Week 3 of the 30-day plan, and every major release after.

Day-by-day checklist: - **Monday:** submit to directories and trackers relevant to your niche (DeFiLlama, DappRadar, agent directories). Confirm your docs and llms.txt are live and crawlable. - **Tuesday:** publish the launch blog post (Play 4 output #3). Line up your Product Hunt assets if using it. - **Wednesday - launch day:** launch thread (link in reply - Play 6) posted in your audience-overlap window - early-to-mid afternoon UTC for US + Europe - demo video with URL on the end card, ecosystem cross-posts go out, Product Hunt goes live if scheduled. Check the calendar first: if a major industry event or news moment lands that day, shift a day rather than get buried. Be online all day answering every reply. - **Thursday:** personally DM your first 20 target users. Use Play 11's outreach prompt to draft, then personalise each one yourself - a template that reads like a template converts like one. - **Friday:** pitch ecosystem media and newsletters (Play 11). Coordinate a cross-post with another bootcamp team. - **Weekend:** run the Watchdog (Play 7) on launch metrics. Note what pulled signups and what didn't - that's next week's double-down.

Play 11: Partnership and outreach assistant

Goal: run the "perception through partnerships" play (Tip Sheet §16) with the cold-outreach craft from §15. **When:** ongoing - two or three genuine attempts a week beats a monthly blast.

Steps: 1. List 10 targets: ecosystem teams, niche newsletters, podcasts, micro-KOLs, one reach partner. 2. For each, run the research prompt, then the draft prompt. Personalise before sending. Follow up until told to stop - most replies come from follow-ups.

PROMPT - Outreach Research

Research [person/org] using their site and recent public posts: [links]. Tell me: what they're currently working on or promoting, what they seem to care about, and 3 specific ways my project [paste one-liner] could offer THEM value first (content, integration, data, audience, effort). Rank by realism.

PROMPT - Outreach Draft

Using that research and my guardrails [paste], draft a cold message:
3 sentences max. Sentence 1: specific and personal (proves I did the work).
Sentence 2: the value I'm offering them, concretely. Sentence 3: one clear, small ask. Also give me 3 candidate subject lines under 6 words.
Then a polite follow-up for 4 days later that adds one new piece of value.

Play 12: KOL vetting assistant

Goal: never pay for a fake audience (Tip Sheet §6). **When:** before any paid collaboration. **Needs:** browsing helps, but you can paste screenshots instead.

PROMPT - KOL Vet

Vet this account for a paid collaboration: [handle/link, or paste their recent posts and stats]. Run the 5 checks: 1) engagement rate math on recent posts, 2) reply quality - real questions vs emoji strings, 3) follower-to-view ratio red flags, 4) how their past sponsored posts performed vs organic, 5) any history with failed or scam projects.
Verdict: proceed / proceed with caution / avoid - with the one decisive reason.
If proceed: suggest a deal structure that pays on results, and the tracking link setup I need before anything goes live.

Part 5 - Quality control (every publish)

Play 13: The pre-publish review

Goal: a two-minute final gate that catches the mistakes that cost reach and trust.

PROMPT - Pre-Publish Check

Review this draft against my guardrails [paste both draft and guardrails].
Check and answer each in one line:
1. Any statistic or claim I haven't verified? List them.
2. Any guardrail violations (voice, formatting, banned words/claims)?
3. If it's an X post: is the link in the reply, not the body? Is the first line a real hook?
4. Jargon test: would a non-technical friend understand every sentence?
5. One improvement that would make this land harder - then stop.

If check 1 lists anything, verify it or cut it. A wrong number can earn a Community Note, and that carries a real reach penalty (Tip Sheet §1.2).

Part 6 - The standing habit

Play 14: The idea inbox

Goal: marketing runs on inspiration - ideas form while you consume and while you do, and they evaporate fast. This habit catches them and turns them into output. **Time:** 2 minutes to set up. Seconds per capture. Ongoing forever.

Steps: 1. Create one note called `idea-inbox` on whatever is always within reach (phone notes app is perfect). One file, not many. 2. The moment an idea arrives - mid-scroll, mid-build, mid-conversation - write one line and get back to what you were doing. No formatting, no judging. 3. **Act on the strong ones fast, while the energy is high.** Acting is what develops the idea: draft the post, sketch the campaign, message the person - the same day if you can. A rough version shipped today beats a perfect version imagined forever. 4. Weekly, when you run the Content Engine (Play 4), skim the inbox: the best entry often becomes the week's source asset; delete anything that no longer excites you. An empty-ish inbox is a healthy inbox.

PROMPT - Idea Developer (for when an idea feels promising but vague)

```
Here's my positioning doc: [paste]. Here's a raw idea from my inbox:
"[one line]".
Develop it three ways: 1) the smallest version I could ship today,
2) the full version if it works, 3) the X post that tests whether anyone
cares, written to my guardrails [paste]. Then tell me which to do first
and why, in two lines.
```

Play 15: The search starter (SEO + GEO in one afternoon)

Goal: get found by both kinds of search engine - human and AI - with one afternoon of setup and one hour a week after (Tip Sheet §3-4). **Time:** ~3 hours once, then it compounds.

Steps: 1. Set up Google Search Console and Bing Webmaster Tools (Bing feeds ChatGPT's retrieval). Submit your sitemap to both. Add an `llms.txt` - a plain-text index of your key pages and docs for AI crawlers. 2. Run the keyword prompt below to build a 10-query list. 3. Publish your first comparison page using the page prompt - comparison pages rank for humans and get quoted by AI engines. 4. Turn your 5 most-asked community questions into FAQ/docs pages (feed them through Play 4's blog format). 5. Weekly after: one post against the query list, plus the GEO monitor (Play 8). Judge early progress on Search Console impressions, not clicks - impressions move first.

PROMPT - Keyword List

```
Here's my positioning doc: [paste]. Here are autocomplete and "People also ask"
suggestions I harvested for my topic: [paste]. Here are my main competitors'
top pages/sitemaps: [paste or list].
Build me a list of 10 long-tail queries my exact target user would type,
ranked by buying intent. For each: the query, the intent behind it, and the
content type most likely to win it (comparison page / how-to / FAQ / glossary).
Exclude anything too broad to win as a small site.
```

PROMPT - Comparison Page

Using my positioning doc and guardrails [paste both], draft a "[my project] vs [competitor]" page: question-style title, the direct answer in the first 60 words, an honest feature/approach comparison (concede what they do well - honesty gets cited), 2-3 verifiable numbers, a 3-question FAQ block, and suggested FAQPage schema. Write for extraction: short paragraphs, no hype. Flag any claim I need to verify before publishing.

The operating rhythm (what this looks like per week)

- **Daily (10-15 min):** Watchdog (Play 7) · replies on X · publish one piece from this week's Play 4 batch · capture ideas the moment they arrive (Play 14).
- **Weekly (90 min):** one source asset through the Content Engine (Play 4-6) · GEO monitor (Play 8) · 2-3 outreach attempts (Play 11).
- **Monthly (1 hr):** competitor teardown (Play 9) · one new page against your Play 15 query list · re-read your Positioning Doc and Guardrails and update them.
- **Per launch:** the runbook (Play 10).

Total: roughly 90 minutes a day, most of it reviewing rather than producing. That's the point - the AI supplies volume and speed; you supply judgement, voice and the decisions. One founder running these plays consistently covers what an early marketing hire would - and you'll know your numbers better than most funded teams do.

Companion to the Growth Tip Sheet - section references (\$) point there. Both documents: Ben Wynn, CMO, GOAT Network · X @0x1164 · Telegram @OX1164 · July 2026.