for tech founders and startup builders

the pursuit of clients

a technical playbook to land your first paying clients

> The pursuit of clients works best when you attract them rather than chase them



PLUS AUTOMATION TEMPLATES AND PROMPTS

by Ayyoub

a tech serial entrepreneur

FOR TECH FOUNDERS AND STARTUP BUILDERS

The Pursuit of Clients

A hands-on guide for tech builders who want real customers — not just followers

by Ayoub

a tech serial entrepreneur

Last updated: September 2026

Table of Contents

Part 1 — Foundations

- Chapter 1: Why Most Builders Struggle to Land Their First Clients
- Chapter 2: Finding Your Ideal Client (ICP)
- Chapter 3: Positioning in One Clear Sentence

Part 2 — The Irresistible Offer

- Chapter 4: Anatomy of an Offer People Can't Refuse
- Chapter 5: Guarantees and Reducing Perceived Risk
- Chapter 6: Pricing
- Chapter 7: Ready-to-Copy Offer Templates

Part 3 — Digital Presence (Before You Go Looking for Leads)

- Chapter 8: Why Your Presence Needs to Be Ready Before a Lead Finds You
- Chapter 9: Automating Post Creation, for Free
- Chapter 10: Automating Carousel Creation, for Free
- Chapter 11: Automating Video Creation, for Free
- Chapter 12: Posting at the Right Pace (Free Scheduling Tools)
- Chapter 13: Growing Your Audience: The Levers That Actually Work
- Chapter 14: Being Found When They Google You (Search + AI Visibility)

Part 4 — Finding Your Leads (Automated)

- Chapter 15: Building a Prospect List Without Spending Hours on It
- Chapter 16: Automating Research with n8n + LLM APIs
- Chapter 17: Qualifying and Scoring Automatically

Part 5 — Outreach

- Chapter 18: Email — Scripts + Automating Sequences and Personalization
- Chapter 19: Calls — What's Automatable vs. What Stays Human
- Chapter 20: Direct Contact and DMs — Partial Automation Without Sounding Robotic

Part 6 — Converting and Closing

- Chapter 21: Structuring the Sales Conversation
- Chapter 22: Handling Objections
- Chapter 23: Landing Your First Client with No Track Record
- Chapter 24: Closing Without Sounding "Salesy"

Part 7 — After the First Sale

- Chapter 25: Turning a Happy Client into a Referral Engine
- Chapter 26: Going from 1 to 10 Clients

Part 8 — Scaling

Chapter 27: Partnering with Micro-Influencers
Chapter 28: Breaking Into Reddit the Right Way

Bonus

Chapter 29: Toolkit and Recommended Stack
Chapter 30: Final Checklist of Mistakes to Avoid

Appendices

Appendix A: The 30-Day Launch Plan
Appendix B: Resource Links
Appendix C: Step-by-Step Setup

PART 1

Foundations

Chapter 1: Why Most Builders Struggle to Land Their First Clients

You built something. It works. You launch. You wait.

Nothing happens.

The problem isn't your product. It's the gap between building and selling. Most technical founders spend 90% of their time perfecting features and 10% finding out if anyone actually wants them. By the time they look for clients, they're out of energy, out of runway, or both.



Fig 1.1 — The 90/10 trap: most builders perfect features and neglect demand.

Here's what nobody tells you: **the first client is the hardest.** Not because your product isn't good enough, but because you haven't built the machine yet — the machine that finds people, talks to them, and brings them in. Once that machine exists, getting client 2 through 10 is a fraction of the effort.

This book is about building that machine. Not with theory. Not with "just post on LinkedIn" advice. With concrete systems, scripts, automations, and a playbook you can execute this week.

The one truth you need to internalize before reading further:

> *You don't need more features. You need a repeatable way to reach the right people, show them something they can't refuse, and make saying yes easier than saying no.*

Before we go further: **everything in this book runs through opencode** — a free, local AI coding assistant. You'll mainly use opencode to drive the system: research, posts, carousels, videos, and outreach all start with a single prompt you type into opencode.

Start here — two steps and you're ready:

1. **Download and install opencode:**

```
curl -fsSL https://opencode.ai/install | bash        # macOS / Linux
winget install opencode.opencode                     # Windows
```

2. **Open the downloaded starter-kit folder with opencode:**

```
cd /path/to/pursuit-of-clients-starter-kit
opencode
```

That's it. opencode reads the `AGENTS.md` file in the folder — it knows what to do with every tool, every skill, every prompt. From here, follow along chapter by chapter.

Chapter 2: Finding Your Ideal Client (ICP)

Before you automate anything, before you write a single outreach message, you need to answer one question: **who exactly are you selling to?**

Your Ideal Customer Profile (ICP) is not "small businesses." It's not "anyone who needs AI." It's a specific type of business with specific characteristics:

Define your ICP with precision:

Dimension	Your answer
Industry	(e.g., real estate, healthcare, home services, legal)
Company size	(e.g., 2–50 employees)
Geography	(e.g., your country first, then expand)
Tech comfort	(e.g., low-code, owner makes decisions, not IT)
Budget	(e.g., $50–$250/month for software)
Buying trigger	(e.g., missed calls, slow replies, no-shows)



Fig 2.1 — Three filters that separate real prospects from tire-kickers.

The three filters that actually matter:

1. **Pain they feel daily.** Not a theoretical problem — something they complain about to their team, their spouse, their accountant. "I missed three leads last night because nobody answered the phone."
2. **Budget they're already spending.** If they're paying $200/month for a tool that doesn't work, switching to yours at $99/month is an easy yes. If they've never paid for software, you're educating a market — avoid this for your first clients.
3. **Access to the decision-maker.** Can you reach the owner, the manager, the person who feels the pain? If the decision requires a committee, a procurement process, or an IT review, move on. Your first clients should be one conversation away from a "yes."

Exercise: write your ICP in one sentence

> *"[Industry] businesses with [size] employees in [geography] who struggle with [specific pain] and currently spend [budget] trying to solve it."*

Example: *"Home service businesses with 3–20 employees in Austin who miss customer calls on evenings/weekends and lose leads to faster competitors."*

Another example: *"Solo medical practitioners with 1–5 staff in Chicago who lose 15% of appointments to no-shows and waste hours on manual scheduling."*

Your First Prompt

You don't have to fill out that table by hand. The delivered starter kit includes an `icp-discovery` skill that analyzes your product — its codebase, website, or a description you provide — and proposes your ICP automatically. Here's the prompt to type into opencode:

```
Analyze my product idea and workspace. Based on the codebase, website,
or description I provide, propose my Ideal Customer Profile (ICP) —
industry, company size, geography, tech comfort, budget, and buying
trigger. Save the result into brand-config.json.
```

The skill reads whatever you point it at — a GitHub repo, a deployed site, or a plain-text product description — and outputs a structured ICP in `brand-config.json`. Run it, review the output, tweak what doesn't fit, and you're ready for the next chapter.

Chapter 3: Positioning in One Clear Sentence

If you can't explain what you do in one sentence that a 12-year-old would understand, you don't have positioning yet. You have a feature list.

The formula:

> *I help [ICP] solve [pain] so they can [outcome].*

What makes this work:

- It names the audience (so they recognize themselves)
- It names the problem (so they feel understood)
- It names the result (so they see the value)

What it is NOT:

- "We use AI to optimize customer engagement across omnichannel touchpoints." (Nobody cares about your tech stack)
- "An all-in-one platform for..." (All-in-one means nothing to a busy business owner)
- "The future of..." (The future doesn't pay this month's bills)

Examples that work:

- *"I help plumbing companies answer every call and book every job — even at 2 AM."*
- *"I help clinic owners fill cancellation slots automatically so they stop losing revenue."*
- *"I help agencies send proposals in minutes instead of days, so they close faster."*

Your homework: Write 5 variations. Test them by saying each one out loud to someone who isn't a tech person. The one that gets "oh, that's cool" or "I know someone who needs that" is your winner.

The kit does this for you. The `offer-builder` skill in the delivered starter kit reads your project description, analyzes your market, and gives you exact positioning lines and a pricing structure. Here's the prompt:

```
Analyze my product, its features, and its market. Propose a positioning
sentence, a pricing ladder (free trial → monthly → annual), and three
offer angles I can test in outreach.
```

It reads `brand-config.json` (your ICP from Chapter 2), combines it with your product description, and outputs a complete positioning brief. You don't have to guess.

PART 2

The Irresistible Offer

Chapter 4: Anatomy of an Offer People Can't Refuse

A great product with a mediocre offer loses to a mediocre product with a great offer. Every time.

Your offer is not your product. Your offer is the **bundle of value, risk reduction, and framing** that makes the decision easy.

The five components of an irresistible offer:

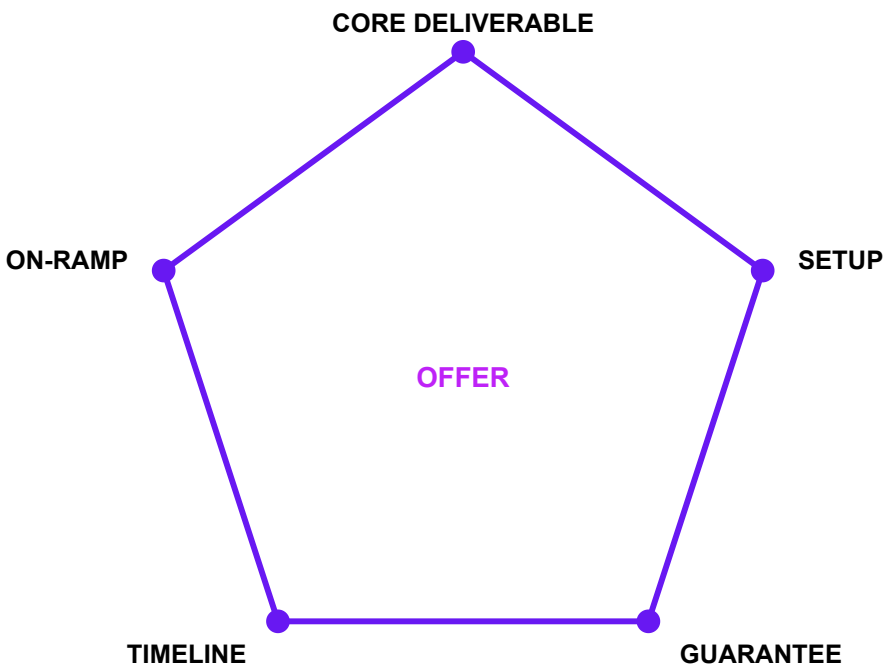


Fig 4.1 — Five components that turn a product into an irresistible offer.

1. **The core deliverable.** What you actually provide. Keep it concrete. "A ready-to-use AI assistant configured for your business" beats "access to our AI platform."
2. **The setup.** You do the work. The client doesn't lift a finger. This is your biggest differentiator. Most SaaS tools require the customer to do the heavy lifting. Flip that.
3. **The guarantee.** Take away their risk. If they don't see value, they walk away for free. This isn't being weak — it's being confident.
4. **The timeline.** When will they see results? "Ready in 24 hours" beats "setup takes a few weeks."

5. **The on-ramp.** What happens after the free period? A clear, no-pressure path to continued use.

Chapter 5: Guarantees and Reducing Perceived Risk

The biggest objection isn't price. It's: **"What if this doesn't work and I've wasted my time?"**

Your job is to make the answer to that objection: "You can't lose."

Risk-reduction tools ranked by power:

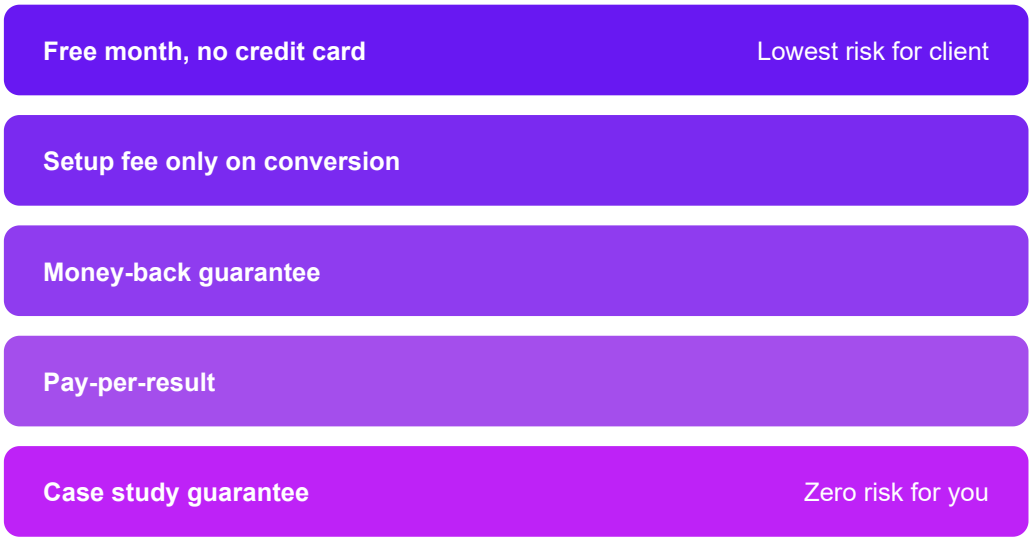


Fig 5.1 — Risk-reduction tools ranked by power. Start at the top.

Type	What it says	Risk to you
Free month, no credit card	"Try it. If you like it, pay. If not, walk away."	Low — you invested setup time
Pay-per-result	"Only pay when it works"	Medium — you need volume
Money-back guarantee	"Pay now, get your money back if not happy"	Low — most won't ask
Setup fee only on conversion	"Free setup + first month; pay only when you stay"	Low-medium
Case study guarantee	"I'll publish a case study of your results"	None — you get marketing material

The power move for your first clients:

> *"I'll set up everything for you — your account, your AI agent, your workflows, your integrations. You don't pay anything. No credit card required. Just use it for a month. If you like it, we continue. If not, you keep everything I built, no hard feelings."*

Why this works:

- **Zero risk for the client.** There is literally no downside.
- **You're investing your time, not their money.** This shows confidence.
- **The setup creates switching costs.** Once their agent is trained on their data, configured for their workflow, and handling their conversations — they won't want to start over.
- **Nobody refuses.** If the tool is good enough to survive a month of real use, you'll convert.

Chapter 6: Pricing

Pricing for your first clients is about one thing: **getting them in the door.**

The ladder:

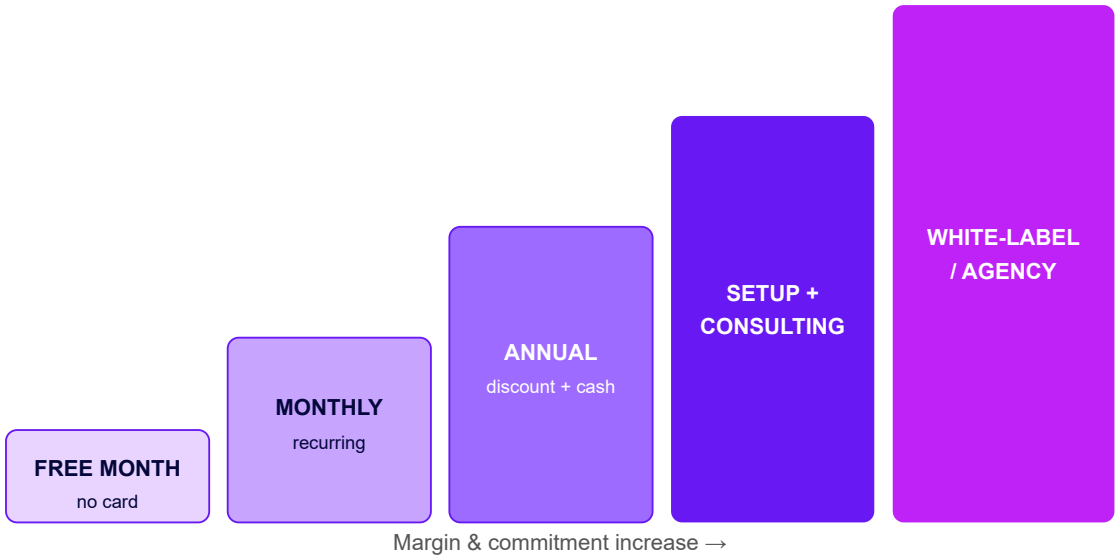


Fig 6.1 — The pricing ladder: acquire cheap, monetize deep.

1. **Free month (no card).** This is your acquisition cost. Think of it as marketing spend — except instead of paying Facebook, you're paying with setup time.
2. **Monthly subscription (after free month).** Standard pricing. No commitment. They can cancel anytime. This is where you prove ongoing value.
3. **Annual plan (discount).** After 2-3 months of happy use, offer a discount for annual payment. This is where you get cash flow.
4. **Setup + consulting (upsell).** Once they're a paying client, offer additional services: custom agent training, team onboarding, workflow optimization. This is high-margin.
5. **White-label / agency (scale).** For clients who want to resell your solution to their own customers.

Pricing psychology tips:

- Don't start high. Your first goal is case studies and testimonials, not profit.
- The free month isn't charity — it's the most efficient sales tool you'll ever have.
- Always have a clear "next step" pricing. The client should never wonder what they'll pay.

Chapter 7: Ready-to-Copy Offer Templates

Template 1: The Zero-Risk Setup

> *Hi [Name],*
>
> *I saw [something specific about their business]. I built an AI assistant that handles [their specific pain — e.g., "customer inquiries on WhatsApp and Instagram"].*
>
> *I'd like to set it up for your business — completely free. No credit card, no commitment. I'll configure everything: the agent, the knowledge base, the integrations. You just use it for a month and tell me what you think.*
>
> *If it works, we talk about continuing. If not, you keep everything.*
>
> *Want to try it? Takes about 20 minutes of your time to get started.*

Template 2: The Specific Result Promise

> *Hi [Name],*
>
> *I help [ICP type] stop losing [leads/customers/appointments] because they [specific problem — e.g., "can't answer messages at night"].*
>
> *I'll set up an AI assistant for your business — handles WhatsApp, Instagram, your website chat, even phone calls. Qualified it to [specific result — e.g., "book appointments directly into your calendar"].*
>
> *First month free. No card. I do all the setup. You just see it work.*
>
> *Interested?*

Template 3: The Competitor Switch

> *Hi [Name],*
>
> *I noticed you're using [competitor tool]. Most businesses I work with tell me [common complaint — e.g., "it's too expensive for what it does" or "the setup is a nightmare"].*
>
> *I built something that does [core value] — set up in minutes, not weeks. And I'll do the setup for you, free, for a month.*
>
> *If it's better, you switch. If not, no hard feelings.*
>
> *Worth a 15-minute call?*

The kit finishes this for you. The `offer-builder` skill in the delivered starter kit does this for your specific product — it reads your `brand-config.json` and produces a tailored offer with exact pricing, guarantee language, and three outreach angle templates. You don't have to write any of this from scratch.

PART 3

Digital Presence

Before You Go Looking for Leads

Chapter 8: Why Your Presence Needs to Be Ready Before a Lead Finds You

You're about to start reaching out to people. Some of them will Google you before replying. Some will check your social profiles. Some will look for evidence that you're real.

If your digital presence is empty, you're invisible.

But here's the thing: you don't need to spend months building a following. You need **enough content that looks professional and credible** when someone checks you out.

The minimum viable presence:

1. **A simple landing page** (or website). One page. What you do, who you help, proof it works, how to get started. That's it.
2. **Social profiles that look alive.** Not necessarily huge — just not empty. 10-15 posts that show you know what you're talking about.
3. **Consistent branding.** Same colors, same tone, same name across platforms. This signals professionalism.
4. **Social proof.** Even one screenshot of your tool working, one testimonial, one case study. Something that says "this is real."

The goal is not to go viral. The goal is: when someone checks you out after your outreach message, they think "this person looks legit."

Chapter 9: Automating Post Creation, for Free

You don't have time to manually create posts for every platform. Good news: you don't have to.

The system:

One AI agent handles the entire pipeline. You give it a topic, it produces everything:

- One image per platform (correct dimensions, correct style)
- Captions tailored to each platform's tone
- Hashtags that rotate to avoid shadowbans
- All stored in organized folders, ready to schedule

The tools (all free or open-source):

Tool	What it does	Cost
web-image	Renders HTML/CSS to PNG at exact platform dimensions	Free (local)
Playwright	Screenshots, page automation	Free (local)
Openverse / Picsum	Free stock photos, no API key	Free
ffmpeg	Image processing, format conversion	Free

The workflow:

```
You (topic) → AI agent (content + design) → HTML template → PNG export (4 platforms) → Captions (.txt)
```

What you get per post:

- Instagram/Facebook: 1080×1350 (4:5)
- TikTok: 1080×1920 (9:16)
- LinkedIn: 1080×1080 (1:1)
- X/Twitter: 1600×900 (16:9)

Time investment: 5 minutes per post (you give the topic, the system does the rest).

The rotation rule: Never post the same visual layout twice. Alternate between editorial, conversation bubbles, stat-hero, before/after, cards, and icon-grid layouts. Your audience should never feel like they're seeing the same post again.

What You Get In The Starter Kit

This book ships with a ready-to-launch starter kit — a zip file containing all the skills, templates, scripts, and a master prompt (AGENTS.md) that tells opencode exactly what to do with each tool. You don't need to build any of this. You download it, open it in opencode, and start. The `posts-factory` skill in the kit generates single-image posts for all platforms from a single topic prompt. The `carousel-factory` skill builds multi-slide carousels. The `video-factory` skill creates motion-graphics videos with AI voiceover. All of them read your brand settings from `brand-config.json` so every piece of content looks consistent.

> ***How to set this up and run it:*** *See* ***Appendix C: Step-by-Step Setup*** *at the end of this book.*

Chapter 10: Automating Carousel Creation, for Free

Carousels are the highest-engagement content format on Instagram and LinkedIn. They're also the most time-consuming to create manually. Not anymore.

The system:

One AI agent generates a complete carousel package:

- HTML slides in your brand identity (light theme, consistent palette)
- Multi-platform exports (Instagram 4:5, TikTok 9:16, LinkedIn 1:1, X/Twitter 3:4)
- LinkedIn PDF (native carousel format — required by LinkedIn)
- Per-platform captions with hashtags
- Automated quality checks (dimensions, text overflow, logo presence)

Your brand, not hardcoded tokens:

The kit's HTML templates read your colors, fonts, and logo path from a single file: `brand-config.json`. You define your palette once — background, primary, secondary, text color, gradient — and every template, every layout pack, every export uses those tokens. No hardcoded color values to hunt down. No inconsistent visuals across platforms.

The kit ships with 4-5 layout packs to rotate between carousels — editorial, conversation bubbles, stat hero, before/after, cards, and icon grid. Each pack is a distinct visual style that shares your brand palette. Rotate them so your audience never sees the same layout twice.

Step-by-step setup and first use

> *For the full setup walkthrough — installing opencode, configuring your brand, and running your first prompt — see **Appendix C: Step-by-Step Setup** at the end of this book.*

The rules:

- Slide 1 must stop the scroll (hook, not intro)
- Never show internal tracking or day numbers in slides
- One strong visual proof per slide (not two competing elements)
- X/Twitter: max 4 images per post (hard platform limit)
- LinkedIn: export as PDF document, not separate images
- Quality check every export before publishing

Output per carousel: 4 platform exports + LinkedIn PDF + 6 caption files.

Chapter 11: Automating Video Creation, for Free

Video is the most engaging format, but also the most intimidating to produce. The system below lets you produce professional motion-graphics videos without a video editor.

The stack:

Tool	Role	Cost
HyperFrames	HTML → MP4 renderer (frame-by-frame)	Free (local)
GSAP	Animation engine (timeline-based)	Free (CDN)
ElevenLabs	AI voice-over	Pay-per-use (~$0.01/word)
ffmpeg	Audio normalization, format conversion	Free
Lottie	Animated icons	Free

Each tool in the kit does a specific job:

- `icp-discovery` skill → analyzes your product and defines who to target
- `offer-builder` skill → creates your irresistible offer and pricing
- `posts-factory` skill → generates single-image posts for all platforms
- `carousel-factory` skill → builds multi-slide carousel exports
- `video-factory` skill → creates motion-graphics videos with AI voiceover
- `seo-factory` skill → runs your search + AI visibility cadence (Search Console ground truth, Google Trends + Autocomplete demand checks, GEO audit, weekly score tracking) — see Chapter 14
- `design-quality` skill group → the craft layer: `motion-graphics-designer` (motion design taste, choreography to narration & SFX), `cinematic-camera` (camera moves), `motion-graphics` (strict visual QC gate), `apple-design` + `emil-design-eng` (UI polish philosophy). The agent reads these before building anything video so the output looks designed, not templated.
- `prospect-research` skill → finds and scores potential clients (runs on Agent Reach)
- `agent-reach` skill → the research engine: searches LinkedIn, X, Reddit, Instagram, Facebook and the web from one CLI (github.com/Panniantong/agent-reach)
- `web-image` skill → the HTML → PNG renderer the posts and carousel factories rely on
- `outreach-dashboard` skill → runs a local dashboard that collects and qualifies leads

- `AGENTS.md` → the master prompt that chains everything together

You don't wire these tools together yourself. AGENTS.md is the glue — it tells opencode which skill to call, in what order, and what to do with the output. Each skill reads from `brand-config.json` for visual consistency and writes its outputs to organized folders so the next skill in the chain can pick them up.

The workflow:

```
Script (beats) → VO (one MP3 per scene) → Music pad → HTML composition → Render →
Normalize → QC → Export per platform
```

Critical rules:

- One MP3 per scene (never split a single narration file — it cuts sentences mid-phrase)
- Pad each VO with +0.20s trailing silence
- Total video: 1080×1920 @30fps, H.264 yuv420p
- Audio: -14 LUFS integrated, ≤ -1 dBTP
- Central 1080×1080 safe zone (for IG 4:5 and LinkedIn 1:1 crops)
- Quality check mandatory before any export

Visual packs (rotate between videos):

Pack	Style
V1	Kinetic typography
V2	Stacked cards
V3	Convergence (elements merge into logo)
V4	Before/After
V5	Stat hero (count-up)
V6	Animated icons
V7	Dynamic background

Exports per video: TikTok/Reels (9:16), IG/Facebook (4:5), LinkedIn (1:1), X/Twitter (9:16), YouTube Shorts (9:16) + 6 caption files.

> ***How to set this up and run it:*** *See* ***Appendix C: Step-by-Step Setup*** *at the end of this book.*

Chapter 12: Posting at the Right Pace (Free Scheduling Tools)

Creating content is half the job. Publishing consistently is the other half.

Recommended posting cadence:

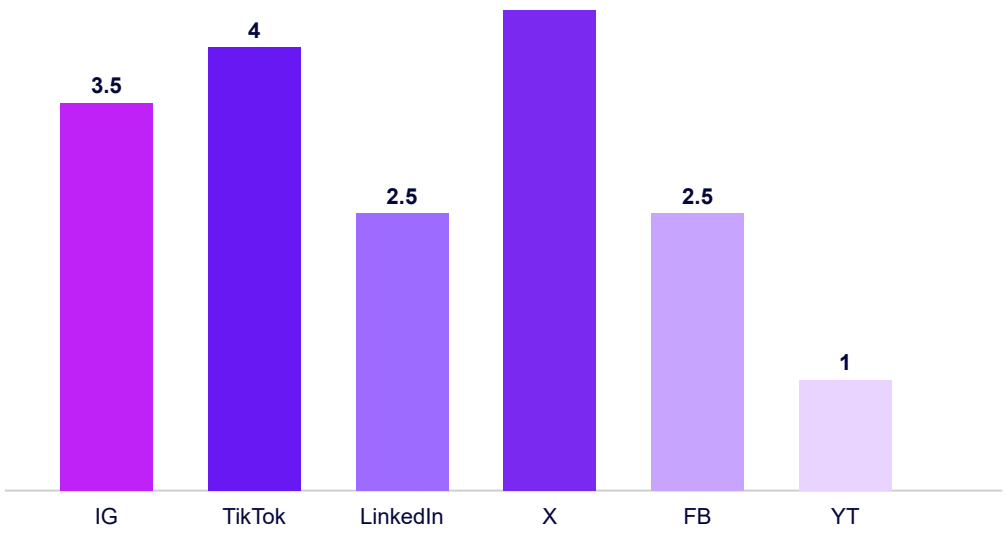


Fig 12.1 — Recommended weekly posting cadence per platform.

Platform	Posts/week	Format
Instagram	3-4	Mix of carousels, single images, reels
TikTok	3-5	Short-form video, 15-60 seconds
LinkedIn	2-3	Carousels, text posts, articles
X/Twitter	5-7	Threads, single images, short video
Facebook	2-3	Repurpose from Instagram
YouTube	1	Shorts (repurposed from TikTok)

Free scheduling tools:

Tool	Platforms	Limits	Cost
Meta Business Suite	IG, FB, Threads	30 days ahead	Free
Buffer	Multiple	3 channels, 10 posts/channel	Free tier
Later	IG, TikTok, FB, Pinterest	5 posts/month per profile	Free tier
TweetDeck / X Pro	X/Twitter	Full scheduling	Free
n8n + APIs	All	Unlimited (self-hosted)	Free (self-hosted)

The n8n automation (advanced):

Build a workflow that:

1. Reads from a folder of exported images + captions
2. Posts to each platform via their APIs
3. Logs success/failure to a tracking sheet
4. Rotates hashtags automatically

Pro tip: Batch your content creation. Create all posts for the week in one session. Schedule them. Then focus on outreach — which is where clients actually come from.

Chapter 13: Growing Your Audience: The Levers That Actually Work

Posting content alone won't grow your audience. You need to be intentional about distribution.

The levers:

1. **Consistency.** Post at least 3x/week on your primary platform. Algorithms reward regularity.
2. **Engagement.** Comment on other people's posts. Answer questions in communities. Be helpful before you ask for anything.
3. **Hashtags.** Rotate 2-3 sets per platform. Never reuse the same set twice in a row (shadow-ban risk). Mix broad and niche tags.
4. **Cross-posting.** One piece of content should live on 3-4 platforms. Adapt the format (carousel for IG, PDF for LinkedIn, video for TikTok), not the message.
5. **Social proof.** Share results. "This client stopped missing calls since last week." Screenshots, numbers, real stories.
6. **DM strategy.** When someone engages with your content (likes, comments), that's a warm signal. Reach out.

What doesn't work:

- Buying followers
- Follow/unfollow
- Generic "great post!" comments
- Posting only about yourself
- Inconsistent posting (3 posts one week, silence the next)

Chapter 14: Being Found When They Google You (Search + AI Visibility)

Before a single outreach message goes out, a prospect is going to look you up. They'll search your product name, your position, your problem area. If the search results answer "nothing here — this person doesn't exist," your outreach just lost to the empty page.

Here's the number most builders miss: **your posts and videos don't have to do all the work.** Search is the channel that works while you sleep — and in 2026 it has two fronts:

1. **Classic search (SEO).** Someone types your pain point into Google and your landing page shows up.
2. **AI answers (GEO — Generative Engine Optimization).** Someone asks ChatGPT, Perplexity, or Gemini "is there a tool that does X?" and your product gets cited in the answer.

The good news: both fronts are **free to work on**. No SEO suites to buy, no AI-marketing SaaS to subscribe to. Every signal in this chapter costs your time, not money.

The free signals, and what each one tells you:

Signal	Tool	What it proves
Search Console	Google Search Console (free)	Ground truth — the actual queries, impressions, clicks and positions for your domain, last 28 days
Demand	Google Trends (free)	Whether a topic is growing or flat in your market
Real user phrasing	Google Autocomplete (free, no API key)	The exact bare-queries people type
Structure	JSON-LD schema markup (free)	How clearly Google understands your page
AI-friendliness	`llms.txt` + QA answers (free)	Whether AI assistants can read and cite you

The rule that saves you from guessing:

> ***Search Console is ground truth; everything else is a hypothesis.*** *Never "optimize" a page in the dark. Pull the last 28 days of queries and positions first — find pages that are ranked 8-15 (visible, but not clicked), match them against autocomplete and trends to confirm demand, then improve exactly those.*

Your minimum search-visible surface (one afternoon, free):

1. **One landing page that answers a real question.** Your positioning sentence from Chapter 3, written like a headline, not a feature list.
2. **One clean `<title>` and description** per page. The title's job is to say what the page is; the description's job is to make someone click.
3. **JSON-LD schema.** Mark up your product (name, what it does, price, reviews). This is what makes rich results and AI answers possible. You can generate it by hand in 20 minutes — no plugin needed.
4. **A sitemap submitted to Search Console** so your few pages actually get indexed.
5. **An AI-readable surface** (`/llms.txt`), summarizing your product in plain, quotable language — this is the fastest single lever for GEO.
6. **Prove freshness.** Search and AI systems both favor pages that are maintained. Every post, carousel and video you ship from the previous chapters is a freshness signal — make sure the ones that answer a search query link back to your landing page.

The 30-minute weekly cadence:

```
1. Search Console → last 28 days: top queries, top pages.
2. Compare positions vs. last week. Find the "almost there" cluster (position 5-15).
3. Autocomplete + Google Trends on those seeds → is demand real or a mirage?
4. One GEO sanity check: do current AI assistants cite your domain when asked
   about your problem area? If not, note it.
5. Improve the ONE page with the best ROI (the fix with the strongest GSC data
   behind it). Never five fixes at once.
```

This is the whole system. It's 30 minutes a week, and it compounds.

The kit does this for you — `seo-factory`. The `seo-factory` skill in the starter kit is a ported, product-agnostic version of the same SEO/GEO command center this book was built with. It reads your `brand-config.json` (domain, product keywords, language), keeps a `TRACKING.md` that logs your visibility score week over week, and runs the cadence above. It uses only free signals — Search Console, Google Trends, Google Autocomplete — and it refuses to recommend paid tools. One prompt starts it:

```
Read brand-config.json and run the seo-factory skill. Pull my last 28 days
of Search Console data, snapshot my visibility score, validate the top 3
queries with autocomplete + trends, and propose the SINGLE page fix with
the best ROI. Update content/seo/TRACKING.md.
```

Every time you run it, the score is logged. The rule is simple: **never let the score regress** — if a change drops it, you revert or fix it the same week.

What to do with "we're not showing up anywhere" (the honest answer):

If Search Console shows zero impressions for anything except your own brand name, there's nothing to optimize yet — you have a *content* gap, not a technical problem. The fix isn't more schema. It's the content flywheel from Chapters 9-11 pointed at the pain-point cluster your ICP actually searches (validated by autocomplete + trends, not vibes). Search rewards the pages that answer real queries with real details: your pricing, your setup time, your "does it work with X" answers. That's the same content your prospects need to say yes — you're building it once and it serves two jobs.

PART 4

Finding Your Leads

Automated

Chapter 15: Building a Prospect List Without Spending Hours on It

You need a list of potential clients. Not 10,000 names — 50-100 high-fit prospects is enough to start.

The manual approach (for your first 20):

1. Search Google Maps for businesses in your ICP (e.g., "plumbing company Austin" or "real estate agency Chicago")
2. Visit their website. Note: name, website, phone, email, social links
3. Check if they have a chat widget, WhatsApp button, or answering service
4. Note their pain signal: no chat? Slow response? Outdated site?
5. Find the owner on LinkedIn

The semi-automated approach (for 50-200):

1. Use Google Maps scraping tools (Phantombuster, Outscraper, or custom Playwright scripts) to extract business listings
2. Filter by: industry, size, geography, online presence signals
3. Enrich with LinkedIn data (decision-maker name, title)
4. Score each prospect: fit (1-5) × pain signal (1-5) × accessibility (1-5)

Data structure per prospect:

```
{
  "company": "Austin Plumbing Pros",
  "website": "https://austinplumbingpros.com",
  "industry": "home-services",
  "size": "5-10",
  "location": "Austin, TX",
  "decision_maker": "Mike Johnson",
  "linkedin": "linkedin.com/in/mikejohnson",
  "email": "info@austinplumbingpros.com",
  "phone": "+15125551234",
  "pain_signals": ["no chat widget", "outdated website", "no online booking"],
  "fit_score": 15,
  "status": "to_contact"
}
```

How the kit finds these prospects — a skill, not a script to wire up.

In the starter kit you don't build scrapers or configure a research pipeline by hand. The `prospect-research` skill in `skills/` does the whole job — it knows the targeting rules, the search workflow, and how to slow down and check quality along the way. It runs on top of **Agent Reach** (`skills/agent-reach/`), the free research engine that searches LinkedIn, X, Reddit, Instagram, Facebook and the web from one command line. Once you've logged into the platforms you need, getting 50 high-fit prospects is one prompt away:

```
Analyze my brand-config.json and run the prospect-research skill.
Give me the first 20 prospects that fit my ICP, with the strongest
pain signals. Check each one on LinkedIn, the web and their social
presence, score them with my scoring model, and save the batch to
outreach/prospects.json.
```

The skill searches the platforms your ICP actually uses, extracts the decision-maker, detects the pain signals, scores every prospect, and drops the qualified list straight into the `outreach/prospects.json` file shown above. No n8n workflow, no LLM API keys, no scraping service to pay for. The same skill is what you'll use again — with different targeting — to find micro-influencers in Chapter 27.

> ***How to set this up and run it:*** *See* ***Appendix C: Step-by-Step Setup*** *at the end of this book.*

Chapter 16: Automating Research with n8n + LLM APIs

You can automate prospect research with n8n workflows that combine web scraping and AI analysis.

The workflow:

```
Input (industry + location)
    → Google Maps API / scraping → Business list
    → For each business:
        → Fetch website → AI analyzes pain signals
        → LinkedIn search → Find decision-maker
        → Score fit (AI-driven)
    → Output: Enriched prospect JSON
```

What the AI analyzes:

- Does the website have a chat widget? (pain signal: no automation)
- Is the site mobile-responsive? (pain signal: outdated)
- Are there customer reviews mentioning slow response? (pain signal)
- What's the business model? (fit check)
- Who's the likely decision-maker? (accessibility)

Tools needed:

Tool	Purpose	Cost
n8n	Workflow orchestration	Free (self-hosted)
Google Maps API	Business listing data	Free tier ($200/month credit)
Agent Reach	Social + web research engine (LinkedIn, X, Reddit, Instagram, Facebook search from one CLI)	Free (open-source)
Firecrawl / Playwright	Website scraping	Free (local)
OpenAI / Gemini API	AI analysis	~$0.01/prospect
PhantomCloud / Outscraper	LinkedIn enrichment	Pay-per-use

Cost per enriched prospect: ~$0.05-0.15 (API costs).

Time saved: From 20 minutes manual research per prospect → 2 minutes automated. For 100 prospects: from 33 hours → 3.3 hours.

The research engine in your kit. The prospecting in this book doesn't use scrapers you have to build. It uses **Agent Reach** (github.com/Panniantong/agent-reach) — a free, open-source router that searches LinkedIn, X/Twitter, Reddit, Instagram, Facebook, the web and more from one command line. Your AI agent reads its skill (`skills/agent-reach/`), calls `agent-reach doctor --json` to see which platforms are ready, then searches decision-makers and pain signals on the platforms your ICP actually uses. Install once, login once per platform, and every search is a single command.

Ready-made n8n workflows in the kit. The starter kit ships with three ready-made n8n workflows in `n8n-workflows/` . Import them into your n8n instance, connect your credentials (Google Sheets, Google Maps, and either an LLM API of your choice or DeepSeek), and they run without building anything. Two belong to the research stage of this chapter; the third belongs to the email stage (Chapter 18).

Workflow 1 — `Gather businesses data` . This is the raw-data collector. Given one or more Google Maps search queries (stored in a spreadsheet, one per row), it pulls every matching business and writes a clean row per prospect to your Google Sheet: business name, address, website, phone, rating, review count, industry, and a status column that tracks whether the prospect has been processed yet. Run it once per search term, then let it sit in the "Pending" pile for the audit workflow to consume.

Workflow 2 — `Multi Pages Digital Presence Audit` . This is the pain-signal detector. It takes the businesses collected by Workflow 1, crawls several pages of each website (not just the homepage), and produces a diagnostic row per business: an overall interaction score, an automation-fit score, which contact methods exist (email, phone, WhatsApp link, calendar booking, live chat), whether the site already uses AI tools, the tech-stack modernity level, the friction points it detected, and — best of all — a ready-to-send outreach summary draft you can drop straight into an email. This is your proof for the "personalize the first line" rule: the friction points it finds are the exact hook you quote.

> *The third workflow, `Local Business Outreach` , handles the email sequence and follow-ups — it's explained in Chapter 18 where email automation lives.*

> ***How to set this up and run it:*** *See* ***Appendix C: Step-by-Step Setup*** *at the end of this book.*

Chapter 17: Qualifying and Scoring Automatically

Not all prospects are worth the same effort. Use a scoring system to prioritize.

The scoring model:

Criterion	Points	How to assess
ICP fit (industry, size, geography)	1-5	Automated from firmographic data
Pain signal strength	1-5	AI analysis of website, reviews, social posts
Decision-maker accessibility	1-5	LinkedIn connection, email availability, company size
Budget indicator	1-5	Current tools paying for, company size, online ads
Timing	1-5	Recent pain signal (new bad review, recent job posting for support)

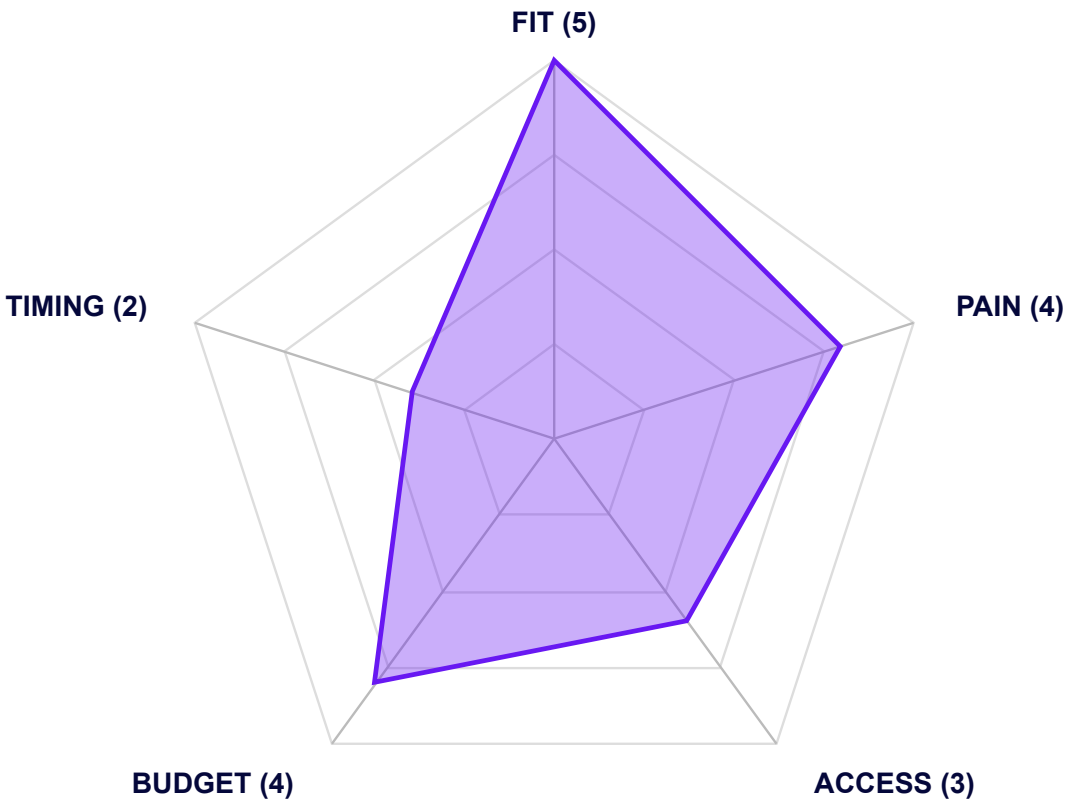


Fig 17.1 — Example radar plot for a Warm lead (total 18/25).

Score interpretation:

- **20-25:** Hot lead. Personal outreach today.
- **15-19:** Warm lead. Automated sequence, then personal follow-up.
- **10-14:** Cool lead. Nurture with content, revisit monthly.
- **Below 10:** Not now. Archive.

The outreach dashboard handles this. The `outreach-dashboard` skill in the starter kit applies your scoring rules automatically. It reads your criteria from `brand-config.json`, scores every prospect as it enters the system, and sorts them into Hot/Warm/Cool buckets. You only spend time on the ones worth calling. No spreadsheets, no manual sorting — the dashboard does it and shows you a prioritized list.

How to open the dashboard locally. The dashboard is a small Node.js server that lives in the kit's `outreach/` folder (`server.mjs`) and stores all your data in local JSON files (`prospects.json` + `statuses.json`) — no external CRM, no cloud, nothing leaves your machine. To open it:

```
cd outreach
node server.mjs     # Node built-ins only — no install, no dependencies
```

Then open **http://localhost:4173** in your browser. Seed the dashboard with your `prospects.json`, and as you work through prospects you mark each one Contacted/Rejected right on the page — the dashboard keeps the state for you across sessions.

PART 5

Outreach

Chapter 18: Email — Scripts + Automating Sequences and Personalization

Email outreach is still the most effective way to reach business owners — if you do it right.

The rules:

1. **Personalize the first line.** Not "Hi [First Name]" — reference something specific. "I saw your Google reviews and noticed several customers mentioned slow response times."
2. **Keep it short.** 4-6 sentences maximum. Business owners don't read long emails.
3. **One CTA.** Don't ask for a call AND a demo AND to check your website. One action. "Want to try it free for a month?"
4. **Follow up.** 80% of deals close after the 2nd-5th follow-up. Most people stop after 1.

The email sequence:

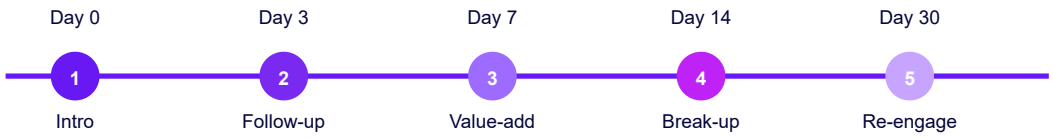


Fig 18.1 — The five-email sequence. 80% of deals close after follow-up #2.

Email	Timing	Purpose
1	Day 0	Introduction + offer
2	Day 3	Follow-up: "Did you see my message?"
3	Day 7	Value-add: share a relevant case study or result
4	Day 14	Break-up: "Is this not a priority right now?"
5 (optional)	Day 30	Re-engage with new offer or content

Automating with n8n:

```
Prospect enters sequence (score ≥ 15)
    → Day 0: Send email 1 (personalized template)
    → Wait 3 days
    → Check: did they reply? If yes → notify you, stop sequence
    → If no → Send email 2
    → Wait 4 days
    → Check reply → ...
    → Continue through sequence
    → Log outcome to tracking sheet
```

The n8n workflow included in the kit (`n8n-workflows/Local Business Outreach.json`) runs this sequence for you — with a human-in-the-loop guardrail. Connect your credentials (Google Sheets, OpenAI or another LLM API, SMTP, and Telegram) and it:

1. **Runs on a schedule** (default 11:00) and pulls the next pending rows from your Google Sheet.
2. **Audits each business's website** and generates enhanced analytics in the background — business size, digital maturity, AI integration level, social presence, estimated monthly lost revenue, strengths and weaknesses.
3. **Drafts a persuasive, personalized email per prospect** with the LLM, and assembles it into an HTML email.
4. **Sends the draft to you on Telegram for approval before anything goes out.** You tap Approve → it sends via SMTP and marks the row as sent (tracking contact date and email count). You tap Reject → it asks for the reason, logs the note, and moves on. This is the "never send something robotic" rule baked into the workflow — you keep final say on every message.
5. **Follows up in sequence** (first email → timed follow-ups based on email count) and stops the sequence when the prospect replies. It also re-engages rejected leads later with a fresh angle from its Angle & Strategy Engine.

Nothing sends without your approval, so deliverability stays clean and every email stays personal. No additional tool required.

Personalization tokens:

Token	Source	Example
[first_name]	LinkedIn / website	Mike
[company]	Google Maps / website	Austin Plumbing Pros
[pain_signal]	AI analysis	"missed customer calls on evenings"
[specific_detail]	Manual or AI	"your Google reviews mention slow replies"

Tools for email automation:

Tool	Cost	Best for
n8n + SMTP	Free (self-hosted)	Full control, custom sequences
Resend / SendGrid	Free tier (100/day)	Reliable delivery
Instantly.ai	$30/mo	Done-for-you cold email
Smartlead	$39/mo	Multi-inbox rotation

Deliverability tips:

- Use a custom domain (not gmail.com)
- Warm up the domain for 2 weeks before cold outreach
- Send 20-30 emails/day per inbox max
- Never use tracking pixels in cold emails (spam trigger)
- Include an unsubscribe link (legal requirement in EU)

Chapter 19: Calls — What's Automatable vs. What Stays Human

Calls are the highest-converting channel. But most of the process can be automated.

What's automatable:

Task	Tool	How
Booking	Calendly / Cal.com	Self-serve link in your email
Reminders	n8n + SMS/email	24h and 1h before the call
Prep	AI agent	Summarize prospect data before the call
Notes	AI transcription	Record + transcribe (with permission)
Follow-up	n8n + email template	Send within 1 hour of the call
CRM update	n8n	Log call outcome automatically

What stays human:

- The actual conversation
- Reading between the lines
- Building rapport
- Handling objections (see Chapter 22)
- Asking for the close

The pre-call prep template:

Before every call, have this ready (AI can auto-generate from prospect data):

```
Prospect: [Name], [Company]
Industry: [Vertical]
Size: [Employees]
Pain signals: [What they struggle with]
Current tools: [What they're paying for now]
Likely objection: [Price / trust / timing]
Your angle: [How your offer solves their specific problem]
Goal of call: [Get commitment for free month setup]
```

The post-call follow-up email (send within 1 hour):

```
Subject: Great chatting — here's what I'll set up for you

Hi [Name],

Thanks for the time today. Here's what I understood:
- [Pain 1]
- [Pain 2]

I'll have your [AI assistant / tool / system] ready by [date].
I'll send you the login details and a quick walkthrough.

In the meantime, [any information you need from them — website URL,
brand assets, FAQ document].

Talk soon,
[Your name]
```

Chapter 20: Direct Contact and DMs — Partial Automation Without Sounding Robotic

DMs work on platforms where business owners actually hang out: LinkedIn, Instagram, Facebook Groups, WhatsApp Business, Twitter/X.

The rules:

1. **Never lead with a pitch.** Lead with value or a genuine observation.
2. **Don't mass-DM.** 5-10 personalized messages per day max.
3. **Engage first.** Comment on their posts, share their content, be visible before you message.
4. **Keep it to 2-3 messages max.** If they don't reply, move on.

What can be automated:

Task	Tool	Notes
Finding profiles	Agent Reach (LinkedIn/X/Reddit/IG/FB searches)	Extract profiles matching ICP from one CLI
Initial engagement	Manual	Like/comment on 3-5 posts before DMing
Message templates	AI agent	Personalized per prospect
Follow-up timing	n8n / CRM reminder	"DM [Name] in 3 days if no reply"
Connection requests	Manual (with personalized note)	Never automate — platforms flag it

What stays human:

- The actual conversation
- Tone judgment (knowing when to be formal, casual, funny)
- Deciding when to stop (not everyone is a prospect)

LinkedIn DM templates:

Message 1 (connection request note):

> *Hi [Name] — I work with [industry] businesses on [pain area]. Saw your profile and thought we might have useful conversations. Happy to connect!*

Message 2 (after connection, 1-2 days later):

> *Thanks for connecting, [Name]. Quick question — how do you currently handle [their specific pain]? I've been helping similar businesses [specific result] and might have something useful for you.*

Message 3 (if they engage):

> *Interesting! I built something that does exactly that — [one-line positioning]. I'm offering a free month to a few businesses in [their industry]. Want me to set it up for you? Takes about 20 min.*

The key: Only send message 3 if they engaged with message 2. If they didn't reply, don't push.

The outreach dashboard replaces manual DM tracking. The kit's `outreach-dashboard` skill runs locally as a Node.js server. It collects prospects via a simple web form (name, company, email, phone, social links), auto-scores them based on your criteria in `brand-config.json`, sends AI-generated personalized outreach messages, and stores everything in local JSON files — no external CRM needed. This is the tool that replaces expensive outreach platforms. It costs nothing to run, stores your data locally, and the AI writes every message based on what it knows about your product and your prospect.

PART 6

Converting and Closing

Chapter 21: Structuring the Sales Conversation

When someone agrees to a call or a demo, you're not selling. You're diagnosing.

The framework:



Fig 21.1 — The five-step sales conversation. Listen more than 60% of the time.

1. **Open (2 min):** "Tell me about your business and how you handle [their pain area] today."
2. **Dig (5 min):** "How often does [pain] happen? What does it cost you? What have you tried?"
3. **Present (3 min):** "Here's how I can help — and here's the offer."
4. **Handle objections (5 min):** Address concerns honestly.
5. **Close (2 min):** "Want me to set it up? No card, no commitment. Let's start."

The golden rule: Listen more than you talk. If you're talking more than 40% of the time, you're pitching too hard.

Questions that uncover pain:

- "How many [leads/calls/messages] do you get per week that slip through?"
- "What happens when someone messages you at 9 PM?"
- "How much time do you spend answering the same questions?"
- "What's the most frustrating part of your customer communication?"
- "If you could fix one thing about how you handle [pain], what would it be?"

Chapter 22: Handling Objections

Every objection has the same root: fear. Fear of wasting money, fear of looking stupid, fear of change. Your job is to address the fear, not just the stated objection.

Common objections and responses:

Objection	Response
"It's too expensive"	"The first month is free. If it doesn't save you [time/leads/money], you walk away. If it does, it pays for itself."
"I don't trust AI"	"It doesn't replace you — it handles the repetitive part so you can focus on the conversations that matter. And you can review everything it does."
"I already have a solution"	"Great — does it handle [specific gap]? I'll set mine up alongside yours for a month. If yours is better, no harm done."
"I don't have time"	"That's exactly why this exists. I do the setup. You just use it. 20 minutes to start, then it runs itself."
"Let me think about it"	"Totally fair. What specifically would help you decide? I can address that now."
"What if it makes mistakes?"	"It learns from your content and only answers what you've approved. You can review and correct anything."

The meta-rule: Never argue. Acknowledge, reframe, offer proof.

Chapter 23: Landing Your First Client with No Track Record

This is the chicken-and-egg problem: you need clients to get testimonials, but you need testimonials to get clients.

The solutions:

1. **Your own use case.** Are you using your own tool? Show it. "I built this for myself first. Here's how it works for my business."
2. **Beta testers.** Frame your first 3-5 clients as "beta partners." They get extra attention, you get feedback and a case study.
3. **The free month as proof.** "I'm so confident this works that I'll set it up for free. If it doesn't deliver results in 30 days, you owe me nothing."
4. **Screenshots and demos.** Even without client results, show the tool in action. Use your own data, test accounts, sandbox environments.
5. **The reverse testimonial.** If you can't get a testimonial before the first client, get a testimonial from the first client immediately. "Can I write up how this went and share it? I'll send you the draft first."
6. **Borrowed credibility.** "I've built [X] number of these systems" or "I helped [type of business] solve [problem]" — even if it was pro-bono or for friends/family.

Chapter 24: Closing Without Sounding "Salesy"

The best close doesn't feel like a close. It feels like a natural next step.

The assumptive close:

> *"I'll have this ready for you by tomorrow. What's the best email to send the login details to?"*

The choice close:

> *"Do you want to start with the WhatsApp integration or the website chat first?"*

The summary close:

> *"So just to recap: free setup, free month, no card, you test it with your real customers. If it works, we continue at [price]/month. Sound good?"*

The deadline close (only if genuine):

> *"I'm onboarding 5 businesses this month and I can only handle 2 more setups personally. Want me to reserve your spot?"*

What never works:

- High-pressure tactics ("This offer expires tonight!")
- False scarcity ("Only 3 spots left!" when you have unlimited capacity)
- Guilt trips ("I spent hours on this...")
- Over-explaining (the more you talk, the more reasons they find to say no)

The ultimate close:

> *"No pressure. Try it for a month. If it's not for you, it's not for you."*

PART 7

After the First Sale

Chapter 25: Turning a Happy Client into a Referral Engine

Your first clients are your best marketing channel. But they won't refer you unless you ask.

When to ask for a referral:

1. **After a success moment.** "Your agent just handled 12 inquiries while you were sleeping? That's awesome. Know anyone else who'd benefit from this?"
2. **After a positive review.** "Would you mind if I shared this as a case study? I'll keep it anonymous if you prefer."
3. **After a billing cycle.** "You've been using this for a month now and it's working well. I'm looking for 2-3 similar businesses to onboard. Anyone come to mind?"

How to make referrals easy:

- Give them a one-page summary they can forward
- Offer a referral incentive (1 free month for each referral)
- Create a simple referral link or email template they can copy-paste
- Make the case study so good they want to share it

The referral email template (send to client):

```
Subject: Quick favor?

Hi [Name],

Glad the tool is working well! I'm looking to help a few more
businesses like yours.

If you know anyone who struggles with [pain], I'd love an intro.
Same deal: free setup, free month, no commitment.

I'll send them a personal note — nothing generic.

Thanks!
[Your name]
```

Chapter 26: Going from 1 to 10 Clients

Once you have 1-3 paying clients, the playbook changes. You're no longer selling — you're scaling.

The flywheel:

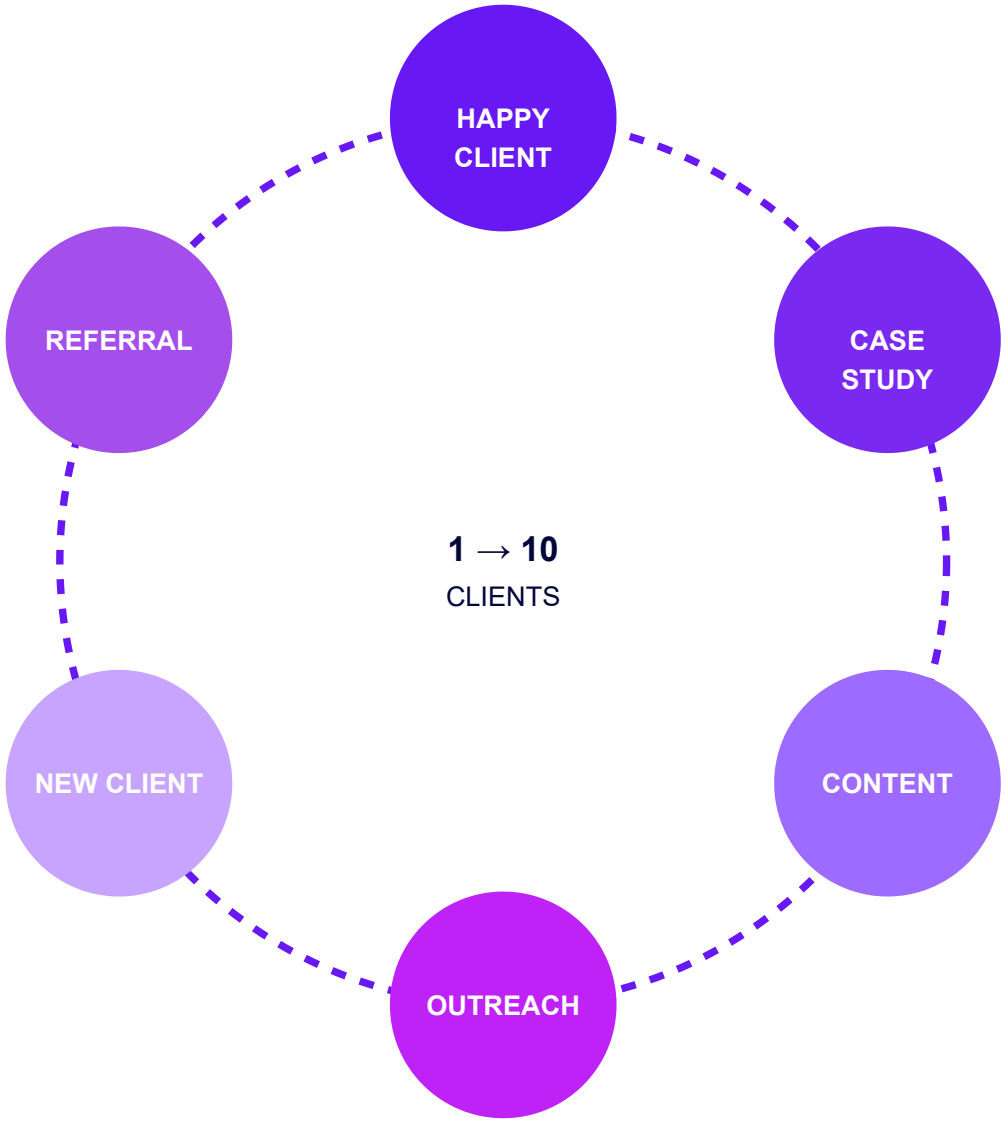


Fig 26.1 — The client acquisition flywheel that compounds from 1 to 10.

```
Happy client → Case study/testimonial → Content about the result →
Outreach with proof → New client → Happy client → ...
```

What to track:

Metric	Target
Outreach sent per week	50-100
Response rate	5-10%
Demo/Call rate	20-30% of responses
Close rate	30-50% of demos
Client retention	>80% after month 1
Referral rate	>20% of clients refer at least 1 person

The hiring trigger:

When you're spending more than 20 hours/week on outreach and delivery combined, it's time to either:

1. Increase prices (to fund help)
2. Automate more (use the systems in this book)
3. Hire a virtual assistant for repetitive tasks

PART 8

Scaling

Chapter 27: Partnering with Micro-Influencers

Micro-influencers (1,000–50,000 followers) in your target industry are the fastest way to reach warm audiences.

Why it works:

- Their audience trusts them
- They need content (you provide value)
- It's cheaper than ads
- It's more credible than cold outreach

How to find them:

1. Search Instagram/TikTok/LinkedIn for your industry keywords
2. Find people posting about business tips, industry advice, or tools
3. Check engagement rate (likes/comments ÷ followers) — aim for >3%
4. Verify they're actually in the industry (not generic "growth hackers")

Finding them with Agent Reach. The kit's research engine — Agent Reach — does the searching for you. Your AI agent runs commands like `opencli instagram search "real estate tips"` (Instagram creators), `mcporter call linkedin.search_people keywords="real estate content creator"` (LinkedIn), or the X/Reddit equivalents, reads their recent posts, computes the engagement rate, and drops matching creators straight into `prospects.json` with their links and engagement scores. The same command language your prospect research uses — no new tool to learn.

The outreach approach:

1. **Engage first.** Comment meaningfully on 5-10 of their posts over 2 weeks.
2. **DM with value.** Not "promote my product" — "I built something that might genuinely help your audience. Can I show you?"
3. **Offer free access.** Give them your tool for free in exchange for honest feedback.
4. **Let them decide.** If they like it, they'll post about it naturally. Don't force it.
5. **Provide content assets.** Make it easy: pre-made graphics, sample captions, key talking points.

The deal:

- Free lifetime access to your tool (in exchange for regular mentions)

- Or: free month + affiliate commission (10-20% of referred clients)
- Or: free setup for their business (if they're an ICP fit too)

Same pipeline, different target. The outreach dashboard in the kit isn't limited to cold prospects. Use the `prospect-research` skill to search Instagram, TikTok, and LinkedIn for creators in your niche. Add them to the outreach dashboard with their social profile links. The AI writes a personalized collaboration pitch based on their content style — not a generic template, but something that references their recent posts, their audience, and why your tool fits their content. Score them on engagement rate, audience fit, and content quality. The same Hot/Warm/Cool buckets apply. The same follow-up sequences apply. Different target, identical pipeline.

Chapter 28: Breaking Into Reddit the Right Way

Reddit is the most underutilized channel for B2B SaaS. It's also the easiest to get banned from if you do it wrong.

The rules (non-negotiable):

1. **Never post about your product directly.** Reddit hates self-promotion.
2. **Be a community member first.** Spend 2-4 weeks commenting helpfully before posting.
3. **Tell stories, not pitches.** "I helped a real estate agency stop missing leads" — not "Check out my AI tool."
4. **Use the right subreddits.** Find communities where your ICP hangs out.
5. **Never use multiple accounts.** Reddit detects and bans vote manipulation.

The strategy:

Phase 1: Be helpful (weeks 1-2)

- Join subreddits where your ICP hangs out (e.g., r/smallbusiness, r/realestate, r/entrepreneur, r/SaaS, r/Juststart)
- Answer questions. Share genuine advice. Build karma.
- Never mention your product.

Phase 2: Share value (weeks 3-4)

- Post a case study: "I helped a business [specific result]. Here's what I learned."
- Share a framework: "5 things to look for when choosing [tool category]."
- Ask for feedback: "I'm building something for [pain]. Would this be useful to you?"

Phase 3: Soft launch (weeks 5+)

- "I built a tool that does [X]. Been testing it with a few businesses. Here's what happened."
- Include real results, screenshots, honest limitations
- Let the comments drive the conversation

Best subreddits for B2B SaaS:

Subreddit	Audience
r/smallbusiness	SMB owners
r/Entrepreneur	Founders
r/SaaS	SaaS builders
r/juststart	Solo entrepreneurs
r/realestate	Real estate professionals
r/healthcare	Healthcare professionals
r/freelance	Freelancers and agencies

The post format that works:

```
Title: "I helped a [industry] business [specific result] in 30 days — here's what I
learned"

Body:
- The problem (what they were struggling with)
- What I built (general approach, not a sales pitch)
- The results (numbers, timeline)
- What I'd do differently
- Lessons for others in similar situations

[Optional: "If anyone wants to know more, happy to share details
in the comments"]
```

BONUS

Toolkit & Checklists

Chapter 29: Toolkit and Recommended Stack

The complete stack (all free or low-cost):

Layer	Tool	Purpose	Cost
Content creation			
Images	web-image (HTML→PNG)	Post/carousel image rendering	Free
Carousels	HTML templates + brand tokens	Multi-platform carousel export	Free
Video	HyperFrames + GSAP	Motion-graphics video rendering	Free
Voice-over	ElevenLabs	AI narration	~$5-20/mo
Stock photos	Openverse, Picsum, LoremFlickr	Free images, no API key	Free
Scheduling & distribution			
Instagram/Facebook	Meta Business Suite	Post scheduling	Free
LinkedIn	Native scheduling	Post scheduling	Free
X/Twitter	X Pro / TweetDeck	Tweet scheduling	Free
TikTok	Native scheduling	Video scheduling	Free
Multi-platform	n8n + APIs	Automated cross-posting	Free (self-hosted)
Prospecting & outreach			
Lead research	Agent Reach (LinkedIn, X, Reddit, Instagram, Facebook search)	Social + web prospecting from one CLI	Free (open-source)
Lead research	Google Maps API + Playwright	Business listing extraction	Free (with limits)
Website analysis	Playwright + AI	Pain signal detection	~$0.01/site
Email outreach	n8n + Resend/SendGrid	Automated sequences	Free tier
CRM / tracking	outreach-dashboard (local JSON)	Prospect management	Free

Layer	Tool	Purpose	Cost
Analytics & optimization			
Post analytics	Native platform insights	Engagement tracking	Free
SEO (ground truth)	Google Search Console	Queries, pages, positions, CTR, last 28 days	Free
Demand validation	Google Trends + Google Autocomplete	Real user phrasing + growth signals, no API key	Free
AI visibility (GEO)	`llms.txt` + JSON-LD schema	Being found and cited by ChatGPT/Perplexity/Gemini	Free
SEO orchestration	`seo-factory` skill	The 30-minute weekly visibility cadence (Chapter 14)	Free
Conversion tracking	UTM parameters + Sheets	Link attribution	Free

The AI agent stack (for your own use):

Component	Purpose
opencode	AI coding assistant (runs locally)
Agent Reach	Research engine — searches LinkedIn, X, Reddit, Instagram, Facebook, web (github.com/Panniantong/agent-reach)
n8n	Workflow automation (self-hosted)
Playwright	Browser automation
ffmpeg	Media processing
Starter kit skills	Content creation, search/AI visibility (seo-factory), outreach, analytics, craft layer (encoding motion-design taste)

Total monthly cost for the entire stack: less than $1 (the system is built on free, open-source, self-hosted tools — the only variable cost is AI voiceover and email volume, which for your first month stays under a dollar total).

Getting Started

> *The full step-by-step walkthrough — installing opencode, unzipping the kit, filling your config, and the exact prompts to run — lives at the end of this book in **Appendix C: Step-by-Step Setup**.*

The skill factories: how they're built, and why they work for ANY product

The three content skills in the kit — `posts-factory`, `carousel-factory`, `video-factory` — are not templates for one product. They are **generic rule-contracts** (this is documented in the kit as `PORTING.md`). They were derived from a real brand-specific skill system that produced daily content for one SaaS brand, and the product-layer was stripped out. A fourth skill, `seo-factory`, was ported the same way — from that same brand's search and AI-visibility command center (see Chapter 14). Here's what makes them product-agnostic:

1. **Configuration is external.** Every skill reads `brand-config.json` — your product name, colors, fonts, logo path, domain, CTA, language. The skill files themselves contain ZERO hardcoded brands, colors, domains, or languages. Replace the config, and the same skills produce *your* brand.
2. **Secrets are external.** API keys (ElevenLabs for voiceover), product test-account credentials, and the product URL live in `.env`. The scripts reference `PRODUCT_EMAIL`, `PRODUCT_PASSWORD`, `PRODUCT_BASE` — never a real account.
3. **Templates are tokenized.** The HTML templates ship with placeholders (`{{SLUG}}`, `{{KICKER}}`, `{{TITRE}}`) and a placeholder logo you swap for your own. The first "setup" prompt you ran re-templates everything with your brand tokens.
4. **Platform rules are absolute, not opinions.** These hard limits survive for every product: X/Twitter max 4 images per post; LinkedIn carousels export as a native PDF document; TikTok and Reels are 9:16 native; the video "safe zone" protects the center square that Instagram and LinkedIn crop into. A brand swap never relaxes these.
5. **Every build passes a QC gate.** Posts/carousels run a text-overflow checker, videos run a QC + loudness gate (-14 LUFS), and every caption set must pass a placeholder check (zero `{{…}}` remaining). A gate that fails blocks delivery — that's by design.
6. **Anonymization is mandatory.** Any screenshot of a logged-in dashboard (real customer data) must be masked (`--mask-text` / `--hide`) and visually re-checked before publishing. Public pages need none.

Porting to a new format. The kit's `PORTING.md` contains the exact recipe used to build the factories, so you can add a new format (a weekly newsletter PDF, a new social network's native size, podcast covers) by copying the pattern: extract invariants → tokenize the template → move

config/secrets out → ship the QC gate → verify with a leak-check. The rule that keeps the whole system clean:

- If a skill file contains your product name, colors, or domain, the port failed.

Chapter 30: Final Checklist of Mistakes to Avoid

Before you start:

- [ ] You have a clear ICP (not "everyone")
- [ ] You can explain your positioning in one sentence
- [ ] Your offer includes a free trial or zero-risk guarantee
- [ ] Your digital presence looks professional (even if minimal)

Content:

- [ ] You're posting consistently (3x/week minimum on primary platform)
- [ ] You're rotating visual layouts (never the same post twice)
- [ ] You're adapting content to each platform (not copy-pasting)
- [ ] You're using the right hashtags (rotating 2-3 sets)

Search visibility:

- [ ] Your landing page answers a real query it can rank for (title + description + JSON-LD schema)
- [ ] You submitted your sitemap to Google Search Console and get impressions
- [ ] You keep a weekly `content/seo/TRACKING.md` score — and it never regresses
- [ ] You validated every optimizer's suggestion against Search Console + Trends + Autocomplete data (free signals only)

Outreach:

- [ ] You're personalizing every first message
- [ ] You're following up (at least 3 times)
- [ ] You're not pitching in DMs — you're starting conversations
- [ ] You're tracking all prospects in a system (not in your head)

Sales:

- [ ] You're listening more than talking on calls
- [ ] You're handling objections with empathy, not arguments
- [ ] You're offering the free month as your primary close
- [ ] You're following up within 1 hour of every call

After the sale:

- [ ] You're onboarding clients properly (not just handing over login)
- [ ] You're checking in weekly during the first month
- [ ] You're asking for testimonials and referrals
- [ ] You're documenting case studies for marketing

Scaling:

- [ ] You're reaching out to micro-influencers in your niche
- [ ] You're building a presence on Reddit (helpful first, promotional never)
- [ ] You're automating repetitive tasks as you grow
- [ ] You're tracking metrics (outreach → response → demo → close)

APPENDICES

Reference

Appendix A: The 30-Day Launch Plan

Day	Action
1-2	Define ICP and positioning (use icp-discovery and offer-builder skills)
3-4	Build minimum digital presence (landing page + JSON-LD schema + `llms.txt` + submit sitemap to Search Console via seo-factory) + 5 posts
5-7	Create first batch of content (10 posts, 2 carousels)
8-10	Build prospect list (50 prospects via prospect-research skill)
11-12	Set up email outreach (import n8n workflow or use outreach-dashboard)
13-14	Send first outreach batch (20 emails) + run your first weekly seo-factory snapshot
15-17	Follow up + engage on social (comment on prospects' posts)
18-20	Send second outreach batch (20 emails) + ship your first SEO fix (best ROI page from GSC data)
21-23	Book and conduct first calls
24-26	Onboard first client(s)
27-28	Ask for testimonial, create case study
29-30	Double down on what worked, cut what didn't + check your visibility score hasn't regressed

Appendix B: Resource Links

- **Content automation:** Starter kit skills for post creation (`posts-factory`), carousel generation (`carousel-factory`), and video production (`video-factory`) — all in the delivered kit
- **Outreach automation:** n8n workflows (`n8n-workflows/Gather businesses data.json` , `Multi Pages Digital Presence Audit.json` , `Local Business Outreach.json`) and the outreach-dashboard skill for local lead management
- **Research engine:** Agent Reach — free, open-source, searches LinkedIn, X, Reddit, Instagram, Facebook and the web from one CLI (github.com/Panniantong/agent-reach); install with `pipx install agent-reach` . Per-platform setup guide: https://raw.githubusercontent.com/Panniantong/agent-reach/main/docs/install.md
- **Prospecting:** `prospect-research` skill (runs on Agent Reach) for automated business + micro-influencer discovery and scoring
- **Search + AI visibility:** `seo-factory` skill for the 30-minute weekly SEO/GEO cadence — free signals only (Google Search Console, Google Trends, Google Autocomplete, `llms.txt` + JSON-LD schema); LLM-accessible, no paid SEO tool needed
- **Brand system:** `brand-config.json` as single source of truth for all visual and positioning tokens
- **Master prompt:** `AGENTS.md` chains all skills together — you don't orchestrate manually

Appendix C: Step-by-Step Setup

Here's the full setup flow, step by step — install, configure, and run your first prompt. After this, each chapter's prompt is enough.

1. **Install opencode:**

```
curl -fsSL https://opencode.ai/install | bash        # macOS / Linux
winget install opencode.opencode                     # Windows
```

2. **Unzip the starter kit** into a project folder. You'll see this structure:

```
starter-kit/
├── AGENTS.md              # Master prompt — tells opencode what each skill does
├── brand-config.json      # Your product, ICP, colors, fonts, positioning
├── .env.example           # API keys template → copy to .env (ElevenLabs, SMTP, product
creds)
├── skills/
│   ├── icp-discovery/     # Analyzes your product, defines ICP
│   ├── offer-builder/     # Creates offer, pricing, positioning lines
│   ├── posts-factory/     # Single-image posts for all platforms
│   ├── carousel-factory/  # Multi-slide carousels with layout rotation
│   ├── video-factory/     # Motion-graphics videos with AI voiceover
│   ├── seo-factory/       # Search + AI visibility cadence (free signals only)
│   ├── prospect-research/ # Finds and scores prospects (runs on Agent Reach)
│   ├── agent-reach/       # Research engine — searches LinkedIn/X/Reddit/IG/FB/web
│   ├── web-image/         # HTML→PNG renderer used by the factories
│   ├── design-quality/    # Craft references: motion-graphics-designer, cinematic-cam-
era,
│   │                      #   motion-graphics, apple-design, emil-design-eng
│   └── outreach-dashboard/# Local lead collection + outreach server
├── content/
│   ├── posts/             # HTML templates + scripts (snap.mjs, fetch-bg.mjs)
│   ├── carousels/         # Carousel templates (layout packs) + export scripts
│   ├── videos/            # HyperFrames compositions + VO/SFX scripts
│   └── seo/               # TRACKING.md (visibility score), keywords +
autocomplete/trends scripts
├── tools/
│   └── shot.mjs           # HTML-to-PNG renderer entry point
├── outreach/
│   └── server.mjs         # Free lead dashboard (prospects.json + statuses.json)
├── n8n-workflows/         # 3 schemas: Gather businesses data, Multi Pages
│                          #   Digital Presence Audit, Local Business Outreach
└── exports/               # All outputs land here, organized by date
```

3. **Fill** `brand-config.json` with your product details. This is the single source of truth — every skill reads from it:

```
{
  "product_name": "[Your Product Name]",
  "product_description": "[One-paragraph description of what it does]",
  "icp": {
    "industry": "[Your target industry]",
    "company_size": "[Employee range]",
    "geography": "[Where your clients are]",
    "pain": "[The daily pain they feel]",
    "budget": "[What they currently spend]"
  },
  "positioning": "[Your one-sentence positioning]",
  "brand": {
    "bg": "#fdfdff",
    "primary": "#6818f2",
    "secondary": "#be22f7",
    "text": "#04073a",
    "gradient": "linear-gradient(120deg, #6818f2, #be22f7)"
  },
  "logo_path": "assets/img/logo.png"
}
```

Most fields are marked `DERIVED` in the shipped file — you only fill the bare facts you know (name, domain, what it does, who it's for, language), and the setup prompt resolves the rest for you with research.

4. **Install and configure Agent Reach** (one-time):

```
pipx install agent-reach
agent-reach doctor --json    # see which platforms are ready
```

Login once to the platforms you'll prospect on (LinkedIn CLI, and Chrome for Reddit/Instagram/Facebook/X) — Agent Reach reuses those sessions so you never reconfigure.

5. **Fill `.env`** with your API keys:

```
ELEVENLABS_API_KEY=your-key-here
SMTP_HOST=smtp.gmail.com
SMTP_PORT=587
SMTP_USER=you@gmail.com
SMTP_PASS=your-app-password
PRODUCT_EMAIL=you@yourproduct.com     # test account on your own product
PRODUCT_PASSWORD=your-test-password
PRODUCT_BASE=https://yourproduct.com
```

6. **Open the folder in opencode and run your first prompt:**

```
cd /path/to/pursuit-of-clients-starter-kit && opencode
```

The one prompt that sets everything up:

```
Read AGENTS.md and brand-config.json. Set up my entire pursuit-of-clients
machine: validate my ICP and the offer/pricing in brand-config.json,
update the post/carousel/video templates with my brand tokens and
domain, generate my first batch of content (1 post, 1 carousel, 1 video),
then tell me the first 20 prospects to research for outreach
starting with my best channel.
```

opencode reads the master prompt, applies your config, and runs the `icp-discovery` → `offer-builder` → content → prospect-research chain. From then on you just ask in plain language.

Then, one prompt per part of the machine

Outreach — a daily prospect pipeline. A single prompt turns the research skill into a recurring job that fills your pipeline every morning:

```
Set up a cron job that runs every day at 9:00. Each run, use the
prospect-research skill to find 20 new high-quality prospects that
match my ICP in brand-config.json, score them with my scoring model,
and append them to outreach/prospects.json so the dashboard picks
them up.
```

From then on you wake up to 20 fresh, scored prospects in your dashboard every morning — without opening the tool.

Social media — a 30-day content plan. A single prompt analyses what's already in your workspace and plans a month of content across post, carousel and video formats:

```
Analyze my whole workspace and brand-config.json. Build a 30-day
content plan for my product: the theme of each day, the format
(post, carousel or video), and the platform mix. Save the plan so I
can follow it day by day.
```

Search visibility — a weekly cadence. Once your sitemap is submitted to Search Console, one prompt runs the whole Chapter 14 loop and keeps your score from regressing:

```
Run the weekly search-visibility cadence for my product using the
seo-factory skill: pull the last 28 days of Search Console data,
validate demand for the keywords I should write about (Trends +
Autocomplete), produce a short optimization plan (title, description,
schema, llms.txt, sitemap), snapshot my GEO score in content/seo/TRACKING.md
and make sure it did not regress.
```

Running each day. Once the plan exists, that's all you ever need to say:

```
Generate day 1.
```

opencode produces that day's post, carousel or video, exports every platform variant, and writes the captions for you — then you schedule and move to the next day. "Generate day 2", "Generate day 3", and so on until your pipeline is running on autopilot.

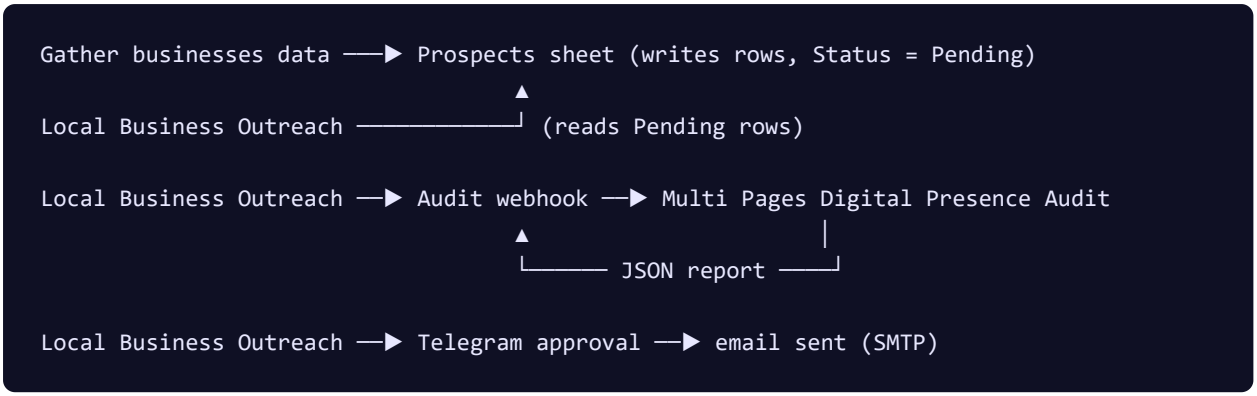
You're in control. The prompts above are starting points, not rules. Once everything is working, tweak them to sound like you — ask for different topics, different platforms, different volumes, different tones. Everything in `AGENTS.md` and `brand-config.json` is editable: change your pricing, your colors, your messaging, your scoring rules, your schedule. Nothing is locked. The machine is built so you can open it up and adjust it whenever your business changes — that's the whole point of it being yours, not a SaaS.

Where outputs land: Everything goes into `exports/<DD-MM-YYYY>/` — posts, carousels, videos, captions. Each platform gets its own subfolder. You schedule from there.

Setting Up the n8n Outreach Workflows

The three ready-made n8n workflows in `n8n-workflows/` automate the whole outreach loop — finding businesses, auditing their websites for pain signals, and sending personalized emails that you approve first. `Gather businesses data` and `Multi Pages Digital Presence Audit` belong to the research stage (Chapter 16); `Local Business Outreach` runs the email sequence (Chapter 18). This guide walks you from a fresh n8n install to a working run.

How the three workflows fit together



- **Gather businesses data** — searches Google Maps (Places API) around a location, dedupes results, and writes new prospects (`Status = Pending`) into a Google Sheet.
- **Multi Pages Digital Presence Audit** — a webhook-triggered workflow. Given a website URL, it crawls multiple pages with a headless-browser service, analyzes them, runs an AI audit, and returns a JSON report (it also logs the result to its own Audit results sheet).
- **Local Business Outreach** — reads Pending rows, calls the Audit workflow's webhook to enrich each business, scores fit and drafts a personalized email with AI, sends it to you on Telegram for approval, then emails it and tracks follow-ups.

Run order: `Gather businesses data` (populate prospects) → `Local Business Outreach` (which calls `Multi Pages Digital Presence Audit` internally). The Audit workflow needs to be **active and reachable** before you run Outreach.

Accounts and services checklist

Needed for	Service	Used by
Storing prospects and results	Google account + Google Sheets	All 3 workflows
Finding businesses	Google Cloud project with the **Places API (New)** enabled + an API key	Gather businesses data
Rendering and analyzing web pages	A Browserless-compatible headless-Chrome service (self-hosted or cloud)	Digital Presence Audit
Writing the audit and email copy	OpenAI API key (billing enabled)	Digital Presence Audit, Local Business Outreach
Approving drafts before sending	A Telegram bot (via @BotFather) + your chat ID	Local Business Outreach
Sending the emails	An SMTP account (Gmail app password, SendGrid, Mailgun, etc.)	Local Business Outreach
Running n8n	Node.js 20.19–24.x, or Docker	All 3 workflows

Part A — Run n8n locally

Pick one:

Option 1 — npx (fastest, nothing to install permanently):

```
npx n8n
```

Option 2 — npm, installed globally (good for daily use):

```
npm install n8n -g
n8n start
```

Requires Node.js 20.19–24.x.

Option 3 — Docker (recommended: easiest to keep your data and update cleanly):

```
docker volume create n8n_data
docker run -it --rm \
  --name n8n \
  -p 5678:5678 \
  -v n8n_data:/home/node/.n8n \
  docker.n8n.io/n8nio/n8n
```

For something you leave running (no `--rm`, restarts automatically), use `docker run -d --restart unless-stopped ...` instead, or a `docker-compose.yml` with the same image/volume.

Whichever option you pick, once it's running open http://localhost:5678 and create your local owner account (email/password — this stays on your machine, nothing is sent anywhere).

About webhooks (important). Two things in this system need to be reachable by an outside caller, not just localhost:

1. The **Telegram bot** (in `Local Business Outreach`) — Telegram needs to deliver messages to your n8n instance from the internet.
2. Cross-workflow calls are fine locally — Outreach calling the Audit workflow's webhook is just n8n calling itself, so `http://localhost:5678/webhook/...` works even with no public exposure.

For the Telegram delivery you have two options:

- **Local testing only:** start n8n with `n8n start --tunnel` (npm/npx) — n8n spins up a temporary public URL that forwards to your machine. Not for production, and not available as a flag on the Docker quick-start above (use the `N8N_TUNNEL` env var or a reverse proxy instead if running Docker).
- **A real setup:** deploy n8n on a small server/VPS with a domain and HTTPS (or use n8n Cloud), and use its real public URL.

Part B — Import the workflows

For each of the three `.json` files:

1. Open n8n → **Workflows** → **Add workflow** → **menu (⋯)** → **Import from File** (or drag the file onto the canvas).
2. Rename if you like, then **Save**.
3. Repeat for all three files.

Part C — Set up credentials

n8n prompts you to select/create a credential the first time you open a node that needs one. Here's what each one is and where to get it:

Credential (name used in the workflows)	Type	Where to get it
Google Sheets account	Google Sheets OAuth2	n8n → Credentials → New → Google Sheets OAuth2 API. You'll need a Google Cloud OAuth **Client ID/Secret** (Cloud Console → APIs & Services → Credentials → OAuth client ID → Web application). Add n8n's shown redirect URL to "Authorized redirect URIs", enable the Google Sheets API on that project, then click Connect in n8n and sign in.
OpenAi account	OpenAI API	platform.openai.com → API keys → Create new key. Add billing/credit to the account or requests will fail. Paste the key into n8n's OpenAI credential.
Telegram account	Telegram API	Message @BotFather on Telegram → `/newbot` → follow the prompts → copy the bot token into n8n's Telegram credential.
SMTP account	SMTP	Any SMTP-capable mailbox or transactional provider (Gmail with an App Password, SendGrid, Mailgun, Amazon SES...). Enter host/port/user/password in n8n's SMTP credential.
Google Places API key — **not** an n8n credential, pasted directly into a node	API key	Google Cloud Console → enable Places API (New) → APIs & Services → Credentials → Create API key. Restrict it to the Places API. Google gives a limited free usage tier; check current Places API pricing before high-volume use.
Browserless / head-less-Chrome token	API key	See Part E, "Multi Pages Digital Presence Audit" below.

Part D — Prepare your Google Sheets

You'll need (at minimum) two spreadsheets — you can add more tabs to the same spreadsheet if you prefer, just repoint each node's `documentId` / `sheetName` accordingly after import.

1. Prospects sheet (shared by all 3 workflows). This is the central sheet: `Gather businesses data` writes new rows here, `Local Business Outreach` reads/updates them, and the audit results get merged in.

Column	Written by
Business Name	Gather businesses data
Address, Website, Phone, Rating, Reviews, Industry	Gather businesses data
Status (Pending → draft_generated → email_sent_N / Completed / Missing_Website / email_rejected)	All 3 workflows
Language, Emails_number	Gather businesses data
Email	You (manually, or add a lookup step) — used as the send-to address by the SMTP node and as the match column in the enriched-data update, but isn't populated automatically by these workflows
Business_Size, **Online_Maturity**, AI_Summary, Strengths, Weaknesses, Recommendations, Digital_Maturity, AI_Integration_Level, Social_Presence, Estimated_Monthly_Lost_Revenue, Improvements	Local Business Outreach (AI analysis)
Subject, Body, Subject1, Body1, Subject2, Body2, Subject3, Body3	Local Business Outreach (each follow-up email)
Email_Count, Last_Contact, Last_Revision_Date, Revision_Notes	Local Business Outreach

> ***Spelling matters for `Online_Maturity`.*** *The AI analysis outputs a field `Online_Maturity` and the GSheets node writes it to a Prospects sheet column named exactly `Online_Maturity`. The header must match the node's column name letter-for-letter or the update silently fails — don't "fix" the spelling to anything else, and make sure the row's email address is filled or nothing sends.*

Create the sheet with just the header row (all columns above) and leave it otherwise empty.

2. Outreach tracking sheet (Local Business Outreach only).

Column
Business Name
Status
Last_Contact
Email_Count
Notes

3. Search queries sheet (Gather businesses data only). This is the input list of what to search for — you fill it in yourself.

Column	Example
City	Paris
Latitude	48.8566
Longitude	2.3522
Search Query	hotels in Paris
Status	Pending

Add one row per city/search combo you want the workflow to run. It flips each row to `Completed` after processing.

4. Audit results sheet (Digital Presence Audit only, "SEO_AUDIT"). This one is populated automatically — just create an empty sheet; the workflow appends rows with the audit fields (Timestamp, URL, Automation_Fit_Score, etc.).

Part E — Workflow-specific configuration

1. Gather businesses data

- In **GSheets – Get 1 search query** and **GSheets – Update status**: point `documentId` / `sheetName` at your Search queries sheet.
- In **GSheets – Store Raw GMaps results1**: point at your Prospects sheet.
- In **HTTP – Google Maps Text Search1**: replace the `X-Goog-Api-Key` header value (`"you api key from gcp account (free)"`) with your real Places API key.
- There's an extra node, **GSheets – Store Raw GMaps results** (no "1"), that isn't connected to anything — it's safe to ignore or delete.

- The **CRON – 5x Daily** trigger fires at 8, 11, 12 (the empty `{}` entry), 17, 20, 10, 15, and 22h — edit the times to whatever schedule you want; n8n needs to be running (or n8n Cloud) at those times for it to fire.
- **Activate** the workflow (top-right toggle) once your sheets and API key are set.

2. Multi Pages Digital Presence Audit

This one needs a Puppeteer/headless-Chrome service reachable at `PUPPETEER_SERVICE_URL`, exposing a Browserless-style `/function?token=...` endpoint (the workflow calls `POST {PUPPETEER_SERVICE_URL}/function?token={PUPPETEER_TOKEN}`).

Two ways to get one:

- **Self-host:** `docker run --rm -p 3000:3000 -e "TOKEN=your-token" ghcr.io/browserless/chromium`, then set `PUPPETEER_SERVICE_URL=http://localhost:3000` (or `http://<container-name>:3000` if n8n and this container share a Docker network) and `PUPPETEER_TOKEN=your-token`.
- **Hosted:** sign up for a Browserless-compatible cloud service and use the URL/token they give you.

Set these as environment variables for n8n itself (e.g. `-e PUPPETEER_SERVICE_URL=... -e PUPPETEER_TOKEN=...` on the docker run command, or in a `.env` file with Docker Compose, or exported before `n8n start` for npm).

- In **Save_To_Google_Sheets**: point at your Audit results (SEO_AUDIT) sheet.
- The webhook path is `audit` (POST). Once the workflow is activated, its production URL is `http://localhost:5678/webhook/audit` (or `https://your-domain/webhook/audit` if hosted). Copy this — you'll need it for the Outreach workflow's `AUDIT_WEBHOOK_URL`. While inactive, n8n also gives you a `/webhook-test/audit` URL that only works while you're clicking "Listen for test event" in the editor.

3. Local Business Outreach

- Point every GSheets node at your Prospects and Outreach tracking sheets (same doc/tab per sheet across all the nodes that reference them).
- Set up the Telegram bot credential (Part C) and get your numeric chat ID (easiest way: message your bot once, then call `https://api.telegram.org/bot<TOKEN>/getUpdates` and read `chat.id` from the response). Replace the literal placeholder `telegramChatId` in the two Telegram nodes ("Telegram – send for approval" / "...approval2") with that number, or better, with an expression reading it from an environment variable.
- Set environment variables: `AUDIT_WEBHOOK_URL` (from step 2 above), `BRAND_NAME`, `BRAND_DOMAIN`, `BRAND_SENDER_NAME`, `BRAND_SENDER_TITLE`, `BRAND_SENDER_CONTACT`, `BRAND_FROM_EMAIL`.

- Fill in the bracketed placeholders left in the AI prompt nodes (Draft Persuasive Email, Generate Enhanced business analytics, Generate Email details for rejected emails) and the industry use-case placeholders in Set Industry Context, with your own product/service description — these were intentionally left blank for you to fill in.
- Activate both this workflow and the Digital Presence Audit workflow (order doesn't matter, but both need to be active for a full run).

Part F — Environment variables reference

Variable	Used by	Purpose
PUPPETEER_SERVICE_URL	Digital Presence Audit	Base URL of your headless-Chrome service
PUPPETEER_TOKEN	Digital Presence Audit	Auth token for that service
AUDIT_WEBHOOK_URL	Local Business Outreach	Full webhook URL of the activated Digital Presence Audit workflow
BRAND_NAME, BRAND_DOMAIN	Local Business Outreach	Your product name/site, used in the email template
BRAND_SENDER_NAME, BRAND_SENDER_TITLE, BRAND_SENDER_CONTACT	Local Business Outreach	Signature block in the email template
BRAND_FROM_EMAIL	Local Business Outreach	"From" address for outgoing email

How to set these: with Docker, add `-e VAR=value` flags (or an `env_file` in Compose); with npm/npx, `export VAR=value` before running `n8n start` (or use a `.env` file with a tool like dotenv-cli, since n8n doesn't auto-load `.env` on its own).

> ***Who receives the emails.*** *The SMTP node sends to the prospect's `Email` column in the Prospects sheet (`$json.Email`), not to an environment variable. If that column is empty, nothing goes out — populate it (manually or with a lookup step) before outreach runs.*

Part G — First test run

1. **Activate Multi Pages Digital Presence Audit.** Test it directly: `curl -X POST http://localhost:5678/webhook/audit -H "Content-Type: application/json" -d '{"url":"https://example.com"}'` and confirm you get a JSON response back and a row appears in the SEO_AUDIT sheet.

2. Add one or two rows to your Search queries sheet, then manually execute `Gather businesses data` (or wait for its schedule) and confirm rows land in Prospects with `Status = Pending`.
3. Activate `Local Business Outreach` and manually trigger its "Schedule Trigger- 11am" node (or run it on demand) — watch for a Telegram message asking you to approve a draft.
4. Approve it and confirm the email actually sends (see the note above about the recipient address).

Troubleshooting quick reference

Symptom	Likely cause
Telegram approval message never arrives	n8n isn't publicly reachable (see "About webhooks" in Part A)
Audit workflow times out	Headless-Chrome/Browserless service isn't running or PUPPETEER_TOKEN is wrong
Google Sheets node errors "credential not found"	Re-select/re-create the Google Sheets OAuth2 credential after import
AI node returns garbled JSON / node fails to parse	Normal occasionally with gpt-3.5-turbo; the workflows have fallback error handling, but persistent failures usually mean the prompt placeholders weren't filled in
Places API call returns 403	Places API (New) not enabled on the project, or the key is restricted to the wrong API

> *This book is a living document. The strategies, tools, and templates will evolve as we learn more about what works. The core principle stays the same: build a machine that finds the right people, shows them something they can't refuse, and makes saying yes easier than saying no.*