

Lead generation

ar002 · 21 June 2026

Abstract

This is the lead-generation notebook for the design-partner campaign (July goal #1, ar006: **200 qualified leads**). It defines who we target, the channels we source them from, and the pipeline that turns a raw company name into a verified, send-ready contact. The mechanics are automated by the **prospecting CLI** (`src/clis/prospect_cli`), which finds and verifies contacts via Hunter and records them to `src/artifacts/leads.jsonl`. A lead only becomes a **Qualified Lead** when its email is *verified deliverable* **and** the contact is an *ICP-role fit*, anything else is held for review or rejected, so we never bounce on the warmed domain. The live results are at the bottom; run `... stats` for the current count.

Methods

Who we target (ICP)

The decisive filter is **who owns the agent's quality**, and the practical filter is **who Hunter can reach**:

- **Builders' customers, not managed.** A team that *built* its agent (in-house or on Botpress / Voiceflow / Rasa / ...) owns its quality, and its testing, which nobody's doing. That's the regression pain testpath solves. With a **managed / turnkey** vendor (Ada, Intercom Fin, Decagon) the vendor owns quality, so the customer feels nothing. **Builders' customers are targets; managed customers are noise.** The teardown (ar001) proves the gap: every builder ships testing, none handle stochasticity.
- **Mid-market (Series A–C, ≈50–500 people).** Big enough to be indexed in Hunter *and* to have budget; small enough to own the agent and feel the pain. Seed-stage companies tend to be invisible to Hunter (no contacts found → rejected), and enterprises already run internal eval teams, both deprioritized.

Source channels (ranked by yield for this ICP)

1. **Builder case studies**, Voiceflow /customer-stories + /pathways, Cognigy, Kore.ai, Yellow.ai. Companies *proud* they built a support agent; the hook writes itself. Catch: front pages skew enterprise; Botpress 403s automated fetch (manual pass). G2 / Capterra reviews are the back door, reviewers name their employer.
2. **Widget / tech-stack detection**, BuiltWith / Wappalyzer “sites using Voiceflow/Botpress”, or detecting a builder's chat widget directly. High precision for “owns a builder agent”, enumerable at scale. Catch: stale detection; filter out managed widgets.
3. **Targeted job posts**, Greenhouse / Lever / Ashby / LinkedIn for *Conversational AI, Support Automation, AI Support Engineer, CX Platform Engineer* (not generic “Applied AI”). Sharp hook. HN “Who is hiring” monthlies are a free vein, but lean seed → Hunter-reject heavy.
4. **Engineering blogs / conference talks**, the DIY angle (where Gusto, DoorDash came from): technical teams who demonstrably own *and* care about agent quality. Lower volume, more reading per lead.
5. **Channel-partner agencies**, the implementation shops that deploy a builder repeatedly (Helpline Hero, Parkfield, Streamline). One relationship reaches a whole base of SMB clients, a partnership play, after direct outreach proves the pitch.

The pipeline (prospect CLI)

Each candidate runs through the same steps, recorded one row per contact to `leads.jsonl` (*company · domain · person · role · email · hook · source · status*):

1. **add**, stage a candidate found via the channels above (company, domain, source, and the **hook**, the specific reason we're reaching out). Lands as **manual**. Dedups against the active list *and* the reject store.
2. **enrich**, Hunter **domain-search** lists contacts; we rank them by ICP role (support / CX / success / conversational-AI / applied-AI / founder / CTO), then **email-verifier** checks the best one.
3. **Promote / reject**, verified **deliverable + ICP-role fit** → **ready** (a Qualified Lead). Deliverable but **off-ICP** → **verify** (human review, never the send list). **Undeliverable** or **no contact found** → moved out to `rejected.jsonl`.
4. **Never re-process**, the reject store means a dead-end domain is never re-sourced or re-enriched (no wasted Hunter quota). **add** refuses a previously-rejected company unless `--force`.
5. **Never send an unverified guess**, bounces wreck the warmed sending domain, so nothing reaches the send list without a deliverable verdict.

Status model: `manual` → `verify` → `ready` (QL); rejects exit to `rejected.jsonl`. Quality over volume, a vetted contact with a real hook beats a 1,000-row blast.

```
uv run python src/clis/prospect_cli/cli.py add --company X --domain x.com \  
    --source builder_case_study --connection "uses Voiceflow, hiring an AI eng"  
uv run python src/clis/prospect_cli/cli.py enrich --all --yes  
uv run python src/clis/prospect_cli/cli.py stats
```

Results

Qualified-lead pipeline (live)

The operational list, the CLI run end to end with **verified** emails, recorded to `src/artifacts/leads.jsonl`. Each row carries its source, a connection link (the blog post, case study, or job ad we open the email with), outreach stage (0 not sent → 3 final follow-up), and last contact. It renders **only in local development** (`leads.jsonl` is gitignored, contact data is never published; production shows a placeholder).

The live qualified-leads table renders only in the operational app (contact data, never published).