

AI Lead Research and Prospecting Agent Pack

Build one on-demand agent that checks approved public sources for one or two observable signals, preserves raw evidence, classifies candidates visibly, remembers prior events, exports a review pack, and prepares editable drafts without sending outreach.

Course: <https://flowgrammer.ca/courses/ai-lead-research-prospecting-agent>

Complete one section after each lesson. Use permitted real or sanitized evidence when available. Keep fictional fixtures, unknown values, manual work, blocked routes, and human approvals clearly labelled.

1. Define the Offer and Research Decision

A business research brief precise enough to search and narrow enough to test.

Prospecting research starts with a decision, not a scrape. Define who might be relevant and what evidence would justify human review.

Decision, scope, and evidence used

Item or rule	Evidence / source	Unknown or failure action	Owner

[] The brief is specific enough to search. [] Review status is separate from purchase intent. [] Personal data is minimized. [] External actions are outside v0.

Result: pass, repair, manual, blocked, narrower scope, or stop, with the reason

2. Research and Select Observable Signals

A selected-signal register with source, event, date meaning, false positives, coverage, and alternative explanations.

A signal is an observable event connected to the offer by a hypothesis. It is not a hidden-intent detector.

Decision, scope, and evidence used

Item or rule	Evidence / source	Unknown or failure action	Owner

[] Every selected signal has a verified source candidate. [] Event date and first-seen date remain separate. [] Alternative explanations are recorded. [] Unverified signals cannot produce priority status.

Result: pass, repair, manual, blocked, narrower scope, or stop, with the reason

3. Approve Sources, Caps, Cost, and Privacy

A collection contract enforced outside the user interface.

Scope controls belong at the execution boundary. A disabled button or friendly prompt does not enforce a paid or privacy limit.

Decision, scope, and evidence used

Item or rule	Evidence / source	Unknown or failure action	Owner

[] Limits are enforced outside the browser. [] Secrets stay out of artifacts. [] Private or unbounded input fails before collection. [] Paid failures do not loop.

Result: pass, repair, manual, blocked, narrower scope, or stop, with the reason

4. Preserve Raw Evidence and Normalize Records

Raw and normalized fixture layers with traceable run, record, event, entity, and evidence identities.

Save what the source returned before cleaning or interpretation. A crash after collection should leave the evidence usable.

Decision, scope, and evidence used

Item or rule	Evidence / source	Unknown or failure action	Owner

[] Raw evidence exists before analysis. [] Every material fact has a source pointer. [] Unknown values remain explicit. [] Errors do not become negative findings.

Result: pass, repair, manual, blocked, narrower scope, or stop, with the reason

5. Deduplicate Entities, Events, and Repeat Runs

A memory and reconciliation design that suppresses unchanged alerts without discarding legitimate events.

One company can produce several real events, while one event can appear on several pages. Model both identities explicitly.

Decision, scope, and evidence used

Item or rule	Evidence / source	Unknown or failure action	Owner

[] Raw duplicates remain available. [] Counts reconcile. [] First-seen cannot refresh event age. [] Review state survives repeat runs.

Result: pass, repair, manual, blocked, narrower scope, or stop, with the reason

6. Classify Evidence Without Inventing Intent

A versioned rubric for priority review, research needed, and excluded states.

A classification should explain why a record needs attention. It should not pretend to know a buyer's private state.

Decision, scope, and evidence used

Item or rule	Evidence / source	Unknown or failure action	Owner

[] Priority never means purchase intent. [] Every decision is explainable. [] Missing evidence does not become a negative fact. [] Weak evidence creates a question or research state.

Result: pass, repair, manual, blocked, narrower scope, or stop, with the reason

7. Build the Review Report and Grounded Drafts

A styled report with same-run HTML, Markdown, JSON, and CSV evidence plus draft-only messages.

The operator should see the finding, evidence, uncertainty, and next question without reading raw provider data.

Decision, scope, and evidence used

Item or rule	Evidence / source	Unknown or failure action	Owner

[] Every personalized statement traces to evidence or the business brief. [] No guessed contact appears. [] Exports reconcile with the report. [] No send or publish control exists.

Result: pass, repair, manual, blocked, narrower scope, or stop, with the reason

8. Run the Canary, Repeat, Recover, and Hand Off

A verified on-demand Research Agent Pack or an honest structural-only, blocked, manual, or stop result.

A useful live run may return zero candidates. Verification concerns collection, evidence, decisions, state, exports, and safe operation, not a lead quota.

Decision, scope, and evidence used

Item or rule	Evidence / source	Unknown or failure action	Owner

[] Zero results can pass. [] Live and fictional evidence are separated. [] State survives restart. [] Future external actions remain unclaimed.

Result: pass, repair, manual, blocked, narrower scope, or stop, with the reason

9. Capstone and critical gates

Grade the complete AI Lead Research and Prospecting Agent Pack and choose the next operating decision.

Criterion	Missing	Partial	Decision-ready	Score

Criteria: Specific offer and decision; Selected observable signals; Approved collection contract; Raw-before-analysis evidence; Stable entity and event identities; Visible honest classification; Reconciled report and exports; Grounded unsent drafts; Fixture, repeat-run, and canary evidence.

Critical gates: no unsupported claim, no hidden external effect, no invented missing value, human decisions remain owned, and the tested status matches the evidence. Passing standard: at least 16 of 18 points with no missing critical gate.

Final decision, approver, limitations, rollback or stop route, and next review