

PROJECT BRIEF

AI-Assisted Recruitment Sourcing Agent

Prepared for	Jim
Version & Date	Version 05, 18 February 2026
Status	Draft for discussion

Purpose

This brief proposes building an AI agent to support your recruitment workflow. The agent will do two things: automatically discover relevant job vacancies in the construction sector across the south-west of England, and generate optimised LinkedIn Recruiter search strings so you can quickly find suitable candidates for each vacancy.

The goal is to reduce the time you spend on vacancy discovery and search construction, freeing you to focus on candidate evaluation and relationship-building — the work that wins placements.

Current Situation

Your current process relies on manual searching across job boards and word of mouth to find vacancies. Once you identify a role worth pursuing, you use LinkedIn Recruiter to search for candidates. This works, but has two constraints.

Discovery is limited by your time. You can only check so many sources in a day, and vacancies that appear on boards or company sites you don't regularly monitor may be missed — or found late.

Search construction is repetitive and skill-dependent. Translating a job specification into an effective LinkedIn boolean search takes care. A weak search returns noise; an over-constrained search misses good candidates. Each role needs its own tailored query, and this adds up across multiple live roles.

LinkedIn Terms of Service

Before describing the proposed approach, it is worth explaining why the agent is designed the way it is.

LinkedIn's terms of service prohibit automated access to the platform. This includes logging in via software, scraping search results, automatically viewing profiles, or sending messages through any tool that is not an approved LinkedIn integration. LinkedIn actively detects and enforces against automated usage, and consequences range from temporary restrictions to permanent account bans.

This matters because the most obvious approach — having an AI agent log into your LinkedIn Recruiter account and run searches on your behalf — would put your account at risk. Even well-disguised automation (such as a browser extension operating within your

session) breaches the terms of service. The risk may be low for light usage, but it is real, and losing a LinkedIn Recruiter account would be a serious disruption to your business.

There are legitimate alternatives. LinkedIn Recruiter is specifically designed for the candidate sourcing workflow, and its search interface accepts boolean queries directly. The agent's job is therefore to do the analytical and preparatory work — finding vacancies, analysing job specifications, and constructing search strings — and then hand off to you to execute the searches manually within LinkedIn Recruiter. This keeps your account entirely within LinkedIn's terms of service while still removing the most time-consuming parts of the process.

A third-party data provider (such as Apollo, Proxycurl, or RocketReach) could supply LinkedIn-sourced candidate data via API, eliminating the need to search LinkedIn directly. However, this data is often days or weeks stale, coverage is incomplete, and costs scale with volume. This remains an option if the manual search step proves to be a bottleneck, but the recommended starting point is the approach described above.

How the AI Works

The agent uses an AI language model called Claude, made by a company called Anthropic, to do the analytical work — understanding job specifications, extracting the key characteristics of a role, and constructing effective search strings. It is worth explaining how this works in practice, because it affects both cost and the way the system is set up.

When the agent needs Claude to do something — for example, analyse a job specification and generate search strings — it sends a request containing the job details and a set of instructions. Claude processes the request and returns a response, typically within a few seconds. The agent then takes that response and uses it in the next step of the workflow. This is the same mechanism for every task the agent performs, whether it is categorising a vacancy from a job board, deciding whether a listing matches your criteria, or generating LinkedIn boolean search strings.

Claude can be accessed in two ways, and we will use both at different stages of this project. Claude Desktop is an application that runs on a desktop or laptop computer, included with a Claude Pro subscription. It is the simplest way to interact with Claude and is well suited to trying things out and running tasks manually. For the final hosted solution, the agent will communicate with Claude via Anthropic's API (application programming interface) — a standard, secure connection over the internet that allows software to send requests to Claude and receive responses automatically, without any manual interaction. The API uses pay-as-you-go pricing based on the volume of text processed.

Cost of Using Claude

Claude Pro subscription. Used during the demonstration and pilot phases. This costs approximately £18 per month and includes generous usage of Claude Desktop. There is a usage cap, but it is unlikely to be reached at the volumes involved in this project.

Anthropic API. Used for the final hosted solution. There is no software licence to purchase. You create an API account (free), load a small amount of credit, and pay only for what you use. There is no minimum spend, no monthly subscription, and no contract. The model best suited to this workload is Claude Sonnet, priced at approximately £3 per million tokens of

input and £12 per million tokens of output. In practical terms, processing a single job specification — sending the details to Claude, receiving the search strings back — costs a fraction of a penny. Even processing 50 vacancies a day would cost well under £1 per month. The initial credit minimum is approximately £4. For this workload, that initial credit could last several months.

Approach

Rather than building the full solution in one go, we will take a phased approach that proves the value of each component before committing to the engineering work of a hosted, automated system. There is a clear go/no-go decision at the end of each phase.

Phase 1: Demonstration

Purpose: Prove that the AI produces useful output — that the vacancy discovery finds relevant roles, and that the search strings it generates are good enough to use in LinkedIn Recruiter.

I will build two lightweight agents: one for vacancy discovery and one for search string generation. These will run locally on my machine using Claude Desktop, with the work instructions and code stored in local folders. I will run the agents manually, using your input on role types, geography, and preferred sources to configure them.

The output — a list of discovered vacancies and a set of LinkedIn search strings for each — will be shared with you as files for review. The focus at this stage is entirely on output quality: are the vacancies relevant? Are the search strings effective? Is the agent finding things you would have missed?

This phase requires no setup on your side and no cost beyond my time. We will review the output together and decide whether to proceed.

Duration: 1–2 weeks.

Decision: Is the output quality good enough to be useful? If yes, proceed to pilot. If not, we adjust or stop.

Phase 2: Pilot

Purpose: Put the agents in your hands so you can run them as part of your real working day and assess whether they save you meaningful time.

I will set up Claude Desktop on your computer along with the agents and their work instructions. You will be able to trigger the agents yourself when you want to — for example, first thing in the morning to check for new vacancies, or when you pick up a new role and need search strings. The output will appear in a local folder on your machine, ready for you to review and act on.

This is deliberately simple. There is no web application, no email notifications, no automation. You run the agent, look at the output, and use what's useful. The point is to test whether the agents fit into your workflow and whether the time saving is real, before we invest in building a more polished system.

The only cost during this phase is a Claude Pro subscription (approximately £18 per month). If you already have one, there is no additional cost.

Duration: 2–4 weeks of regular use.

Decision: Are the agents saving you enough time to justify building a permanent solution? If yes, proceed to the hosted system. If not, you cancel the subscription and stop.

Phase 3: Hosted Solution

Purpose: Turn the proven agents into a robust, automated system that runs without manual intervention.

This phase takes the agent logic that has been tested and refined during the pilot and re-engineers it as a hosted service running on your existing AWS EC2 instance, alongside your current website. The key differences from the pilot are:

Automatic scheduling. The agent runs on its own, once or twice daily, without you needing to trigger it. New vacancies are discovered overnight and ready for you each morning.

Notification by email. You receive a short email each morning telling you how many new vacancies have been found. The email contains a link to a selection page.

Selection page. A simple web page hosted on your server listing today's new vacancies. You scan the list, tick the roles you want to pursue, and press a single "Generate searches" button. One browser tab, one pass, one click.

Search strings by email. Within a few minutes of your selection, you receive an email with the LinkedIn boolean search strings for each chosen role, along with a note explaining the logic behind each one. You paste the strings into LinkedIn Recruiter and proceed as normal.

Deduplication and tracking. The agent maintains a database of vacancies it has already seen, so you only ever see new listings. Over time this also provides a record of the roles you have worked on.

At this stage, the agent communicates with Claude via Anthropic's API rather than through Claude Desktop. This is a pay-as-you-go service with no subscription or licence fee, as described in the "How the AI Works" section above. Notification and search-string emails are sent via Amazon SES, a low-cost email service within the AWS ecosystem you already use.

Because the agent runs on your existing EC2 instance, there is no additional hosting cost.

Duration: 4–6 weeks to build, followed by ongoing refinement.

What the Agent Will Not Do

At every phase, the agent will not log into LinkedIn, view profiles, send messages, or perform any automated actions on the platform. Your LinkedIn Recruiter account remains entirely under your manual control.

The agent will not decide which vacancies to pursue or which candidates to approach. It is a research and preparation tool. You retain full editorial control over every decision that affects your client relationships and reputation.

Where Costs Fall

Phase 1 (Demonstration). My development time only. No cost to you.

Phase 2 (Pilot). Claude Pro subscription (approximately £18/month). My time to set up your machine and support you during the pilot.

Phase 3 (Hosted solution). My development time to build the hosted system. Ongoing running costs are: Anthropic API usage (likely a few pounds per month), Amazon SES email charges (fractions of a penny per email), and any job board API fees if applicable. There is no additional hosting cost because the agent runs on your existing EC2 instance. There are no licence fees, no long-term contracts, and no per-seat charges.

At every stage, costs are transparent and you can stop at any point. The phased approach means you never commit to the next stage of investment without evidence that the previous stage delivered value.

What I Need From You

To get started, I need your input on the following. A 30–45 minute conversation would cover all of it.

Your specialisms within construction. Which types of roles do you typically recruit for? Site managers, quantity surveyors, project managers, trades, estimators, health and safety, design engineers? The more specific, the better the agent's filters.

Geographic scope. We've assumed south-west England. How tightly defined is that? Bristol, Bath, Exeter, Plymouth, Swindon? A set radius, or does it vary by seniority of role?

Job boards and sources you currently use. Which sites do you check today? Are there construction-specific boards that are particularly productive? Are there any you've tried and found worthless?

Key employers. Are there specific construction firms whose vacancies you always want to know about? If so, we can monitor their careers pages directly.

Your LinkedIn search patterns. When you build a search in LinkedIn Recruiter today, what fields and filters do you typically use? Walk me through a couple of recent examples if you can — this helps the agent learn your search style from the start.

Next Step

A 30–45 minute conversation to work through the questions above. That gives me everything needed to start building the Phase 1 demonstration.