

Automated Job Search Monitoring System

A step-by-step guide to building a structured, AI-assisted job search pipeline from scratch

Table of Contents

1. What This System Does
2. Prerequisites and Tools You'll Need
3. Building Your Target Company List
4. Building Your Job Log
5. Setting Up Gmail Job Alerts
6. The 5-Dimension Job Scoring Framework
7. Running the Morning Sweep
8. Automating the Sweep with Claude Cowork
9. Tracking Your Pipeline
10. Continuous Improvement

1 What This System Does

This system automates the most time-consuming part of a professional job search: consistently monitoring job boards across dozens of target companies, deduplicating new postings, scoring them against your criteria, and logging them — without spending hours manually checking each company's careers page.

At its core, the system combines three things:

- **A structured target company list** — every company you care about, with direct links to their job boards
- **A job log spreadsheet** — a single source of truth for every role you've seen, scored, or are pursuing

- **An AI sweep agent** — an automated workflow that checks your Gmail alerts and each company's ATS (applicant tracking system), scores new roles, and logs everything

The result is a morning sweep that takes minutes instead of hours, surfaces only roles that match your criteria, and keeps a persistent record of your entire search. The sweep runs on demand or on a schedule, producing a structured report each time.

2 Prerequisites and Tools You'll Need

Required Accounts and Access

- **Gmail account** — used for job alert emails from LinkedIn, Indeed, Built In, and similar job boards
- **Claude Cowork** (or Claude with computer use / browser access) — the AI agent that performs the sweep
- **Google Chrome** — the browser Claude accesses to navigate ATS job boards
- **Microsoft Excel or Google Sheets** — for your job log and company list

Files You'll Create

- `Target_Companies.csv` — your master list of companies to monitor
- `job_log_MASTER.xlsx` — your master job tracking spreadsheet
- A **Sweep Instructions document** — the prompt file Claude reads to run the sweep (described in Section 7)

Optional but Recommended

- A Gmail label or filter for job alert emails to keep them organized
- A shared cloud folder (OneDrive, Google Drive, Dropbox) so your files are accessible from anywhere

3 Building Your Target Company List

Your target company list is the backbone of the sweep. The more accurate and complete it is, the better your results. Create a CSV file (`Target_Companies.csv`) with the following columns:

Column Name	Description	Example Value
Company	Company name (canonical, used for deduplication)	[Company Name]
Tier	1 = proactively target; 2 = monitor only if inbound	1 or 2
Industry	Industry category for scoring purposes	[Your Industry]
ATS Type	Which ATS platform they use	Greenhouse, Lever, Ashby, Workday, iCIMS, Custom
ATS URL	Direct link to their job board or careers page (see below)	https://boards.greenhouse.io/[token]/jobs
Notes	Any flags — bot detection, requires login, board is inactive, etc.	Bot detection active — check manually
Last Checked	Date the ATS was last successfully checked (updated each sweep)	04-May
Status	Active, Inactive, Unknown	Active

How to Find ATS URLs for Each Company

The most reliable way to find a company's ATS board URL:

1. Go to the company's website and look for a "Careers" or "Jobs" link
2. Click through to their job listings — the URL will reveal the ATS platform
3. Use the direct API or listing URL, not just the homepage

Common ATS URL patterns to know:

- **Greenhouse:** `https://boards.greenhouse.io/[token]` or the API at `https://boards-api.greenhouse.io/v1/boards/[token]/jobs`
- **Lever:** `https://jobs.lever.co/[company-slug]`
- **Ashby:** `https://jobs.ashbyhq.com/[token]`
- **Workday:** `https://[company].wd[N].myworkdayjobs.com/[board-name]`
- **iCIMS:** `https://careers-[company].icims.com`
- **Rippling:** `https://ats.rippling.com/[company]/jobs`
- **Custom/Careers page:** Use the direct search results URL with a role keyword filter if possible

Pro tip: For Greenhouse, use the API endpoint (`boards-api.greenhouse.io`) rather than the visual board. It returns clean JSON that Claude can read instantly without JavaScript rendering. This is the most reliable ATS for automated sweeps.

Tiers: How to Prioritize

Use a two-tier system to manage scope:

- **Tier 1:** Companies you would proactively apply to — you're actively monitoring these and will reach out even without an inbound signal. These should be your top 50–100 companies.
- **Tier 2:** Companies that are acceptable if they come to you (recruiter outreach, referral, LinkedIn alert) but you won't proactively source from them. These stay in the list but get auto-skipped unless inbound.

Start with **50–75 Tier 1 companies**. A list of 150+ is manageable once the sweep is automated, but a smaller, higher-quality list is better when starting out.

How to Build Your Initial List

1. Start with companies you already know and want to work for
2. Look at LinkedIn "Companies" search filtered by your industry and size range
3. Check job boards (LinkedIn, Indeed, Built In) for which companies are actively hiring for your role type
4. Look at your network — where do people you respect work?
5. Research industry reports, "best places to work" lists, and trade publications in your field

6. Add competitors and adjacent companies to any you already have

4 **Building Your Job Log**

Your job log (`job_log_MASTER.xlsx`) is a single spreadsheet that captures every role you've ever seen, scored, or pursued. It's the system of record for your entire search.

Required Columns

Column	Description	Example
Date Logged	When this row was added	29-Apr-2026
Company	Company name — must match Target Companies.csv exactly for deduplication	[Company Name]
Role Title	Exact title from the job posting	[Director, Marketing Strategy]
Score	Your 5-dimension composite score (see Section 6)	7/10
Verdict	Quick-read action label	APPLY STRONG, MAYBE MONITOR, SKIP
Why	Brief rationale — which dimensions passed/failed and why	D1 passes — upstream strategy scope. D2 Tier 1. D3 strong match.
Status	Current pipeline stage	New, Outreach Sent, Applied, Screening, Interview, Offer, Rejected, Closed
Date Posted	When the role was posted (if visible)	29-Apr
Source	How you found this role	LinkedIn Alert, ATS Sweep, Recruiter Outreach, Referral
ATS URL / Job Link	Direct link to the job posting	https://jobs.lever.co/[company]/[job-id]
Notes	Free-form notes — contacts, interview notes, comp range, flags	Hiring manager is [Name]. Comp range \$X-\$Y. Phone screen scheduled [date].
Last Action	Date of the most recent activity on this role	04-May-2026

Column	Description	Example
Date		
Next Action	What needs to happen next	Follow up if no response by [date]

Deduplication Rule

A role is a duplicate if the **Company + Role Title combination** already exists in the log. The sweep agent checks for this before logging any new entry. Roles that are already logged get their status confirmed as "still live" but are not re-added.

Score Tier Color Coding (Optional)

Apply conditional formatting in Excel or Sheets to make the Score column visually scannable:

- **GREEN (7–10):** Apply — meets all primary criteria
- **YELLOW (5–6):** Monitor — worth a closer look, needs full JD review
- **RED (1–4):** Watch/Skip — logged for awareness but not worth pursuing

5 Setting Up Gmail Job Alerts

Gmail job alerts are often the highest signal-to-noise source in the sweep. They surface roles from your target companies without requiring any ATS navigation.

Set Up Alerts From These Sources

- **LinkedIn Job Alerts** — create saved searches for your target role titles + industries + locations. LinkedIn will email you new matches daily or weekly.
- **Indeed Job Alerts** — same approach: saved search with email notifications.
- **Built In** — if you're in tech, Built In sends curated alerts for tech roles by city and function.
- **Hiring Cafe** — an aggregator that emails curated matches. Worth setting up if it covers your field.
- **Google Alerts** — set alerts for "[Your Target Role Title]" "[Your Industry]" to catch press coverage that mentions hiring.

Useful Gmail Search Queries for the Sweep

The sweep agent uses Gmail search to find relevant emails. The core queries:

- `from:jobalerts-noreply@linkedin.com newer_than:4d`
- `from:donotreply@jobalert.indeed.com newer_than:4d`
- `from:support@builtin.com newer_than:4d`
- `from:[alert-sender@hiringsite.com] newer_than:4d`

Adjust the `newer_than` value based on your sweep frequency. For a Monday morning sweep covering the weekend, `newer_than:4d` catches Friday through Monday.

Important: Gmail job alerts often arrive in the *Promotions* or *Updates* tab, not Primary. Configure your sweep to search **all inbox categories**, not just Primary, or the agent will miss most emails.

6 The 5-Dimension Job Scoring Framework

Every role is scored on five dimensions. Dimensions 1 and 2 are **primary gates** — if a role fails either of them, it doesn't pass regardless of other scores. Dimensions 3–5 add nuance and drive the final score.

Dimension	What It Measures	Pass Criteria (customize for your search)	Weight
D1 — Role Scope	Does this role match the type of work you want to do?	Must be [your target scope, e.g., upstream strategy / positioning / GTM] – not pure execution or management of a single channel	Primary gate
D2 — Industry Fit	Is this the industry / sector you want?	Tier 1: [your preferred industries]. Tier 2: [acceptable-if-inbound industries]. Auto-fail: [industries you won't consider]	Primary gate
D3 — Competitive Position	Are you a competitive candidate for this role?	Strong = your background is a direct match. Moderate = adjacent. Weak = stretch.	High
D4 — Compensation	Does the comp range meet your floor?	Minimum acceptable: \$[X]. Auto-fail if explicitly below floor and no path to it.	Medium
D5 — Structure / Never Again	Any structural red flags?	Auto-fail: [company types, structures, or role characteristics you'll never accept, e.g., pure agency work, no remote, contract-only, etc.]	Veto power

Scoring Guide

Score	Meaning	Action
9–10	Exceptional match — rare	Apply immediately. Prioritize above all else.
7–8	Strong match — most primary criteria met	Apply. Move to outreach or application within the week.
5–6	Moderate match — needs more information	Pull the full job description. Re-score before deciding.
3–4	Weak match — one or more primary criteria fail	Log and watch. Do not invest time unless circumstances change.
1–2	Clear mismatch	Log and close. Do not pursue.

Key principle: Score based on what you can see in the job alert or ATS listing. If a role sounds potentially right but the listing is thin, mark it 5–6 MONITOR and pull the full JD before re-scoring. Don't commit to applying or ignoring based on incomplete information.

7 Running the Morning Sweep

The sweep is a structured workflow the AI agent follows each time it runs. You give Claude a detailed instruction document (saved as a text or markdown file in your job search folder) that tells it exactly what to do.

What the Sweep Instruction Document Should Cover

- Your identity and search context** — your target role, seniority level, industries, and any absolute deal-breakers. This is how Claude knows which roles are relevant.
- Gmail sweep instructions** — which sender addresses to search, how far back to look, which inbox tabs to include.
- ATS sweep instructions** — read Target_Companies.csv, navigate to each ATS URL, search for relevant role keywords, extract new listings.
- Deduplication rule** — check job_log_MASTER.xlsx before logging anything. Skip if already logged.

5. **Scoring instructions** — apply the 5-dimension framework to each new role found.
6. **Logging instructions** — add new rows to `job_log_MASTER.xlsx`; update Last Checked in `Target_Companies.csv`.
7. **Report format** — how to structure the morning sweep report (see below).

Sweep Report Format

At the end of each sweep, Claude generates a structured markdown report saved to your job search folder. The report should include:

- **Sweep stats:** Number of companies checked, Gmail queries run, new roles found, total log rows
- **New roles by score tier:** GREEN (apply), YELLOW (monitor), RED (skip) — with score, verdict, and brief rationale for each
- **Top picks:** The 3–5 roles that need attention first and why
- **Previously logged roles still live:** High-value roles already in your pipeline confirmed still open
- **Access failures:** Any companies that couldn't be checked and why, so you know what to verify manually
- **ATS issues to fix:** Broken URLs or tokens to update before the next sweep

Typical Sweep Sequence (What Claude Does)

1. Read `Target_Companies.csv` to get the full company list and ATS URLs
2. Read `job_log_MASTER.xlsx` to load already-logged roles for deduplication
3. Search Gmail for each configured alert sender — extract role titles and companies from email content
4. For each company in Target Companies: navigate to ATS URL, search for relevant roles, extract new listings
5. For each new role: apply 5-dimension scoring, assign a verdict
6. Log all new roles to `job_log_MASTER.xlsx`
7. Update Last Checked date for each company in `Target_Companies.csv`
8. Generate and save the morning sweep report

Common ATS gotchas: Some ATS platforms (iCIMS, Workday, some custom sites) require JavaScript to render and cannot be read by automated agents. Others use bot detection. For these, note them in your Target Companies list as "check manually" and do a quick manual visit during the sweep. Expect 5–15% of boards to need manual handling at any given time.

8 Automating the Sweep with Claude Cowork

Claude Cowork supports scheduled tasks — you can set the sweep to run automatically on a recurring schedule so it happens whether or not you remember to start it.

Setting Up a Scheduled Sweep

1. Open Claude Cowork and navigate to your job search folder
2. Ask Claude to create a scheduled task: *"Run my job search sweep every Monday morning at [time]"*
3. Point the scheduled task at your sweep instruction file
4. Verify the schedule is set correctly and confirm the first test run

The sweep will run in the background on your schedule. When it's done, the morning sweep report will appear in your job search folder, ready for review when you open your laptop.

Recommended Sweep Frequency

- **Monday morning** — covers all weekend postings from Friday–Monday (highest volume)
- **Wednesday or Thursday** — catches mid-week postings before Friday's slowdown
- **On-demand** — run any time you want a fresh check, especially if you're actively interviewing

Two sweeps per week is a good rhythm for an active search. Once per week is acceptable for a passive search.

9 Tracking Your Pipeline

Your job log doubles as your pipeline tracker. Use the Status column to move roles through stages as you take action. A simple pipeline has these stages:

Stage	Meaning	Trigger
New	Scored and logged, no action taken yet	Automatically assigned at logging
Reviewing JD	Full job description pulled, re-scoring in progress	When you pull the full JD for a YELLOW role
Outreach Sent	You've reached out to someone at the company	After sending a LinkedIn message or email
Applied	Application submitted	After clicking submit
Recruiter Screen	Scheduled or completed initial recruiter/HR screen	After scheduling
HM Interview	Scheduled or completed hiring manager interview	After scheduling
Panel / Final	In final interview stages	After scheduling
Offer	Offer received	When offer arrives
Rejected	No longer in process	When rejected or withdrew
Closed	Role no longer available or you've chosen to skip	When role is removed or you pass

Follow-Up Cadence

Use the Next Action and Last Action Date columns to drive follow-up. A simple rule of thumb:

- After outreach: follow up in 5–7 business days if no response
- After application: follow up in 7–10 business days if no response
- After interview: follow up within 24 hours with a thank-you note
- If ghosted after two follow-ups: move to Closed, keep the company on your radar

10 Continuous Improvement

The system gets better over time as you fix ATS URLs, tune your scoring criteria, and refine your target list. After each sweep, review:

- **Access failures** — fix broken ATS URLs before the next sweep. Each broken URL is a company you're not monitoring.
- **Score calibration** — if you're getting too many YELLOWS that never turn into applications, tighten your D1/D2 criteria. If you're getting too few results, check whether your keyword search is too narrow.
- **Company list health** — add companies you discover during the sweep. Remove companies that have been dormant for 3+ months or that have clearly stopped hiring in your function.
- **Gmail alert coverage** — if a role from a target company showed up somewhere unexpected, add that source to your alert setup.

What makes this system work: Consistency. The value compounds over time — the longer you run it, the more comprehensive your deduplication becomes, the better you get at scoring, and the cleaner your company list gets. A sweep run once a week for four weeks will surface patterns you'd never catch manually.

This guide covers a generic implementation of the automated job search monitoring system. Adapt the scoring framework, company list, and sweep schedule to your specific search context.