

// phase 2 deep dive

Market Research



How to turn your target list into a deep intelligence file, and mine it for the real problems worth building your Signal Project around.

baswe.ai // become a software engineer

This is the in-depth guide to phase two of the Signal Project. Work through it once you have a target and a rough company list from phase one. By the end you will have turned that list into a deep intelligence file and a set of real problems worth building around.

Why market research is where the project is won

Phase one told you who you are building for. Phase two is where you learn enough about them to build something they actually care about. This is the phase that separates a generic portfolio project from one that makes a hiring manager stop and pay attention, because everything distinctive about your Signal Project comes from what you uncover here. The specific problem you solve, the language you use to describe it, the way your project fits their world, and the story you tell when you launch all trace back to the research you do now.

It is also the least glamorous phase, which is exactly why most people do it badly or skip it. They grab a few company names, glance at the homepages, and start building, then wonder why their project reads like every other project. The engineers who get hired treat research as the real work, because it is.

The goal of market research is not a list of companies. It is a list of problems those companies actually have. Companies are just where you go to find the problems.

What you are actually doing in this phase

You have two jobs in phase two, and you do them at the same time.

Job one: build an intelligence file. Take your target list and turn each company into a rich, structured profile: what they do, how big they are, what stage they are at, what they are hiring for, what they are shipping, what they are saying publicly, and what their engineers are working on. One company per row, one source of truth you can return to.

Job two: hunt for signals. As you research, you are watching for proto-problems, the repeated themes, pain points, gaps, and manual processes that hint at something worth building. Every signal gets logged. Across your whole list, the strongest recurring signals become the raw material for phase three, where you choose the one problem your project will solve.

You also start one quiet background task in this phase: building a network of loose connections at your target companies, without messaging anyone yet. More on that below.

The output of all this is not a finished project. It is a populated research sheet, a log of candidate problems, and a growing network, all of which feed directly into the next phases.

Set up your research sheet

Do this in the market research template that comes with this lesson, or rebuild it yourself in Google Sheets or Excel. The template has four tabs: a start page, a company research sheet, a network tracker, and a signal log. The company research sheet is the heart of it. Each row is one company, and the columns fall into three groups.

Identity, so you know what you are looking at. Company, website, a one-line description of what they do, industry or vertical, stage or funding, headcount, and location or remote policy.

Signal sources, the five core columns plus a few extensions. The original five are company name, company LinkedIn, relevant open roles, company news, and open source repositories. On top of those, capture their engineering blog and their observed tech stack. These are the places real signals live.

Your synthesis, where the thinking happens. Problem ideas and opportunities you spot, key people at the company, why it fits you, a priority rating, and a status so you know what is left to do.

You do not need every cell filled for every company. You need the identity and a priority for all of them, and deep detail on the ones you are most serious about. The template has dropdowns for priority and status so you can keep it organized as you go.

How to research each company

Work one company at a time on your shortlist, and move through these sources in order. Budget roughly fifteen to thirty minutes per company on a deep pass.

▪ Start with the basics

Open the company website, their LinkedIn page, and their Crunchbase profile. In a few minutes you can capture what they do in one line, their industry, their stage and most recent funding, their rough headcount, and where they are based. This orients everything else.

▪ Read the LinkedIn company page and jobs

On their LinkedIn page, note headcount and whether it is growing, and skim their recent posts to see what they are promoting. Then open the jobs tab and search for the roles relevant to you. Capture how many engineering roles are open and paste the links to the relevant ones. Volume matters: a company hiring ten engineers is in a very different moment than one hiring one.

▪ Mine the job descriptions for signals

This is the single richest source in the whole phase, and most people skip it. A job description is a company telling you, in public, exactly what they need. Read the ones relevant to you closely and pull out the tech stack they use, the team they are building, and above all the problems they mention. Phrases like scale our platform,

build out our AI capabilities, improve reliability, reduce latency, or launch a new product line are direct statements of what hurts and what matters right now. Log anything that sounds like a problem in your signal log.

▪ **Gather news and signals**

Search the company name alongside terms like raises, funding, layoffs, launches, partnership, and AI, and search the founder or CEO name alongside interview and podcast. Use Google News, Crunchbase, their own blog and press page, and executive interviews on YouTube and podcasts. If the company is public, their shareholder letters and earnings calls are goldmines. Read each signal for what it reveals: a funding round tells you they are spending, and on what; layoffs and their stated reason tell you where they are cutting or refocusing; a launch tells you their current priority; and an executive interview often tells you the roadmap and where they think the next frontier is. That last one is exactly the kind of forward-looking problem a standout project can get ahead of.

▪ **Explore their open source and engineering content**

Find the company's GitHub or GitLab organization and look at their active repositories and, importantly, their open issues, which are literally a public list of problems they are working on. Then find their engineering blog. Between the two, you learn their stack, the challenges they care about, and often specific problems described in their own words. This is some of the best signal you will find anywhere.

▪ **Infer the tech stack**

Pull their stack together from the job posts, their engineering blog, and tools like StackShare or BuiltWith. Knowing what they build with tells you what to build your project with so it feels native to them, and whether your existing skills already line up.

Hunt for signals as you go

This is the mindset that makes the whole phase pay off. You are not just filling in cells, you are collecting proto-problems. The moment you notice something that smells like a problem, log it in the signal log with the company, the evidence, a one-line note on why it matters, and whether you have seen it at other companies too.

The kinds of things to log:

- ▶ A hiring theme that repeats, like several companies all building AI features onto an existing product.
- ▶ A pain point stated directly in a job description or engineering blog.
- ▶ An open issue or gap in a company's open source work.
- ▶ A topic an executive keeps returning to in interviews.
- ▶ A visible gap or rough edge in the product itself, as a user.
- ▶ An old, manual, or paper-based process in the industry that no one has bothered to fix.

Pay special attention to signals that show up across several of your target companies. A problem that multiple companies in your industry share is the strongest possible foundation for a Signal Project, because solving it is relevant to all of them at once. These recurring signals are exactly what you will sort through in phase three.

Remember the logistics lesson: an obscure but real problem, like a document that is still handled on paper, often beats a flashy one, because it proves you understand the industry from the inside. Watch for the unglamorous problems, not just the obvious ones.

Start your silent network

While you research, quietly build a network at your target companies. For each company you are serious about, note the departments and teams that are hiring from their job posts, then go to the people section of their LinkedIn and search employees in those teams: engineering leaders, hiring managers, and engineers on the relevant team. Send connection requests to the relevant ones, and send nothing else. No pitch, no message, just the connection.

The reason is timing. In later phases you will reach out to real people, first to ask for their expertise while you build, and later to ask about roles. Those conversations land completely differently when you are already a connection who reached out weeks ago, rather than a stranger showing up the day you need something. You are planting seeds now so they are grown when you need them.

Keep it clean and log it. Aim for a handful of relevant people per company you care about, record each in the network tracker with their name, title, team, and connection status, and resist the urge to message. Do not spam requests to people who are irrelevant, and do not pitch. The whole advantage comes from patience.

Tools to use

- ▶ **LinkedIn:** company pages, the jobs tab, and the people tab, for research and for building your network.
- ▶ **Crunchbase:** stage, funding history, and recent company news in one place.
- ▶ **Google News:** recent coverage; search the company and the founder by name.
- ▶ **Company blog and press page:** the primary source for launches, priorities, and announcements.
- ▶ **GitHub or GitLab:** their open source organization, active repos, and open issues.
- ▶ **Engineering blogs:** the clearest window into their stack and the problems they care about.
- ▶ **StackShare and BuiltWith:** to identify the tools and technologies a company uses.
- ▶ **levels.fyi:** compensation bands, and a way to discover teams and orgs.
- ▶ **Glassdoor:** culture signals and interview intelligence.

- ▶ **hiring.cafe and Wellfound:** aggregated roles across many companies, useful for finding openings fast.
- ▶ **YouTube and podcasts:** founder and executive interviews, where roadmaps and future bets often slip out.

How to work through it efficiently

You will not deeply research a hundred companies, and you should not try. Use two passes.

First pass, quick triage. Go down your whole list and spend just a couple of minutes each capturing the basics and setting a priority. This tells you which companies are worth your real time.

Second pass, deep research. Take your top fifteen to thirty companies and go deep using the steps above, roughly fifteen to thirty minutes each. This is where the real signals come from.

Update the status column as you go so you always know what is left. The aim is depth on your best targets, not thin coverage of everything. It is better to deeply understand ten companies and their shared problems than to skim a hundred and understand none.

Common mistakes to avoid

- ▶ **Treating this as list-building.** The point is finding problems, not collecting logos. If you finish with a tidy list and no problem ideas, you did not do research, you did data entry.
- ▶ **Going shallow on everything.** Skimming a hundred homepages teaches you nothing. Go deep on your best targets.
- ▶ **Skipping the job descriptions.** They are the richest signal source you have. Read them closely.
- ▶ **Not logging signals as you find them.** You will not remember them by phase three. Write them down the moment you see them.
- ▶ **Messaging your connections too early.** The silent network only works if it is silent. Connect now, reach out later.
- ▶ **Reading only the homepage.** The real signal is in the jobs, the news, the repos, and the engineering blog, not the marketing site.
- ▶ **Researching companies you do not care about.** You will do it half-heartedly and miss the good signals. Keep the list to companies you are actually interested in.

What you should have before moving on

By the end of phase two, you should have:

- ▶ A research sheet with your shortlist of ten to thirty companies researched in depth, and the rest triaged with a priority.
- ▶ For each shortlisted company: what they do, their stage and stack, their open roles, recent news and signals, their open source and engineering content, and key people.
- ▶ A signal log full of candidate problems, with the recurring ones flagged.
- ▶ A network tracker with a growing set of relevant connections at your target companies.
- ▶ A clear sense of which five to ten companies are your strongest targets, and what problems keep coming up.

That signal log and that sense of recurring problems are exactly what phase three needs. Do not move on until you can name a few real problems your target companies share.

→ **With your research and signals in hand, move on to Phase 3: Problem Selection.**