BUILD-ALONG · 45 MINUTES · ANY JURISDICTION

Build a horizon scanning agent in Wordsmith.

Horizon scanning is the most common agent our customers build, in every market we work in, regulated or not. This guide walks through building one from scratch: a scheduled agent that monitors the regulations that matter to your business and delivers a scan and a memo to the legal team, on a cadence you set.

Work through it in Wordsmith as you read.

BUILT BY

Sam Ruback
Legal Engineer, Wordsmith AI

Elly Meenan
Legal Engineer, Wordsmith AI

1 Describe · 2 Build · 3 Inspect · 4 Test · 5 Finalize

Before you build

Three principles carry the whole process. The middle steps of the build (inspect, test, prune) feel unnatural to most lawyers, because legal work rewards perfection and clarity while agent building rewards iteration. Get comfortable with that and the rest follows.

- **Spot agent-ready work.** Look for recurring tasks where the business engages legal: processes important enough to be worth capturing end to end. Horizon scanning qualifies almost everywhere.
- **Destination, not journey.** Describe what you are trying to get to, not the steps to get there. Wordsmith fills the gaps, or prompts you to fill them, and you iterate toward the outcome.
- **Less prompting, more vibe-smithing.** Don't wait to craft the perfect prompt. Dictate your instruction, typos and all. Value sooner beats polish. Wordsmith does the heavy lifting: it suggests connectors and knowledge bases, and asks questions where you were vague. You stay in the part you're good at: saying how the work should run.

FIRST, CHECK YOUR COMPANY PROFILE

Company profiles hold entity-level context: corporate information, risk profile, what the business buys and sells, what matters to it. With the profile turned on, everything you tell the agent builder is read against it, so your agent starts targeted rather than generic. Set one up for your entity before you build.

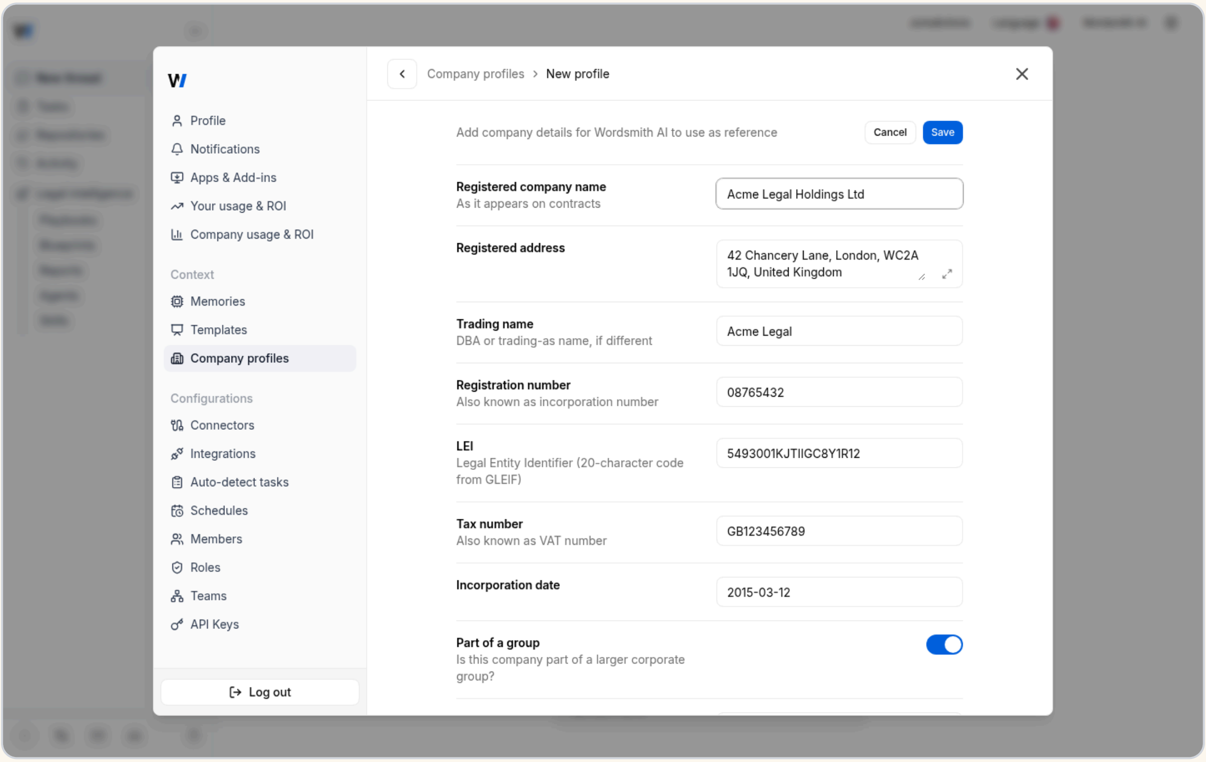

The company profile under Settings → Context: registered entity details, addresses and identifiers Wordsmith reads as reference.

1 Describe the agent in plain words

Open a chat and describe the agent you want. Dictation is faster than typing, and the transcript doesn't need to be clean; a coherent instruction beats a polished one. Name your jurisdictions, and explicitly tell it to ask you questions before it builds.

DICTATE OR TYPE — PROMPT 1

"I'd like you to create a horizon scanning agent for [your business]. It should keep us abreast of [the regulatory topics that matter to you] in [your jurisdictions], so we can develop our products and operations in line with them. Ask enough questions about the business, our product line and our operations to build a targeted agent, so I can stay on top of these issues, manage risk effectively, advise the business as needed in each market, and understand how things develop over time."

2 Answer the questions, then build

Wordsmith responds with clarifying questions before it builds anything. Short answers are fine; it already has your company profile. Expect questions like these:

WORDSMITH WILL ASK	HOW TO ANSWER
Which product lines are in scope?	Name only the categories that carry regulatory exposure.
Which features drive that exposure?	The capabilities regulators actually care about: autonomy, data collection, safety-critical functions.
What is your market footprint?	Where you make things versus where you sell them; often different regimes apply.
Which sectors do you serve?	Sharpens the scan in each market.
Which regulatory domains should it prioritize?	Anchor to your original instruction. Leave marginal domains out; scope discipline beats coverage.

Then the build decisions

Next come configuration choices. Our recommendations for a horizon scanning agent:

- **Output style and audience**: a combination of a tabular scan and a brief but detailed memo for the GC. Legal-team framing, not board framing.
- **Cadence**: match the pace of your market. Monthly, first business day at 9am, suits most; weekly is usually too often.
- **Attach the horizon scanning methodology skill?** Yes, it gives the research structure and steps to follow.
- **Create a knowledge base?** Yes, even if you have no repository yet. Wordsmith creates one and stores every scan in it going forward.
- **Seed standing memories?** Yes, they hold what you've told it (your product, output format, priority domains) so the agent behaves consistently run after run.
- **Dedicated intake email?** Usually no. Horizon scanning runs on a schedule; it rarely needs to be triggered ad hoc.

Review the draft, then create

Wordsmith drafts the agent's standing instructions: scope, jurisdictions, prioritization, output format, things to avoid. Read them. Small changes are a one-line ask in the chat:

DICTATE OR TYPE PROMPT 2

"Two small changes: create the brief memo to the GC as a .docx so I can share it around the business, and on each scheduled scan, email it to [you@yourcompany.com]."

Confirm your timezone, then say *create it*. Wordsmith writes itself a to-do list and works through it; you can watch each step tick off, switch tabs, and come back; the list doesn't go away.

3 Inspect what it built

Don't run it yet. Open the agent and read through each panel to check the logic makes sense before you trust the output.

PANEL	WHAT TO CHECK
Objective	Does it match the destination you described?
Intake	How the agent gets triggered. For now, the schedule only; you'll add sources next.
Instructions	The meat of the agent: the regulations and frameworks it monitors. If a regulator or regime you care about is missing, ask in chat to include it.
Skills	The methodology skill should be ticked. Untick it only if you want your written instructions followed literally, with no imposed structure.
Knowledge	The memories it seeded, plus any repositories you connect; this is where scans get compared against your policies and contracts later.
Connectors	Grant the agent the connectors it needs, e.g. your email connector so the memo actually sends. See the note below.
Activity	Once it runs: how it's been engaged, by what method, and what the responses look like.

CONNECTORS LIVE AT TWO LEVELS

Connectors you set up at account level don't automatically flow to your agents. Each agent is granted its connectors separately: deliberate friction, so an agent used by colleagues never exposes tools they shouldn't reach.

Within a connector, each tool can be set to **always allow**, **require approval**, or **disable**. A sensible default: allow the reads, gate the writes; searching is harmless, sending email on your behalf deserves an approval step.

Add intake sources

With no source instruction at all, Wordsmith researches from your prompt plus its per-jurisdiction research guidelines. If specific sources matter to you (a regulator's site, official legislation portals), name them. And where a source publishes an RSS feed, connect it as an intake: every time the feed updates, the agent listens.

DICTATE OR TYPE PROMPT 3

"Amend the agent to include relevant sources for staying on top of the regulations identified in the instructions, in [your jurisdictions], via RSS feed."

Feeds carry everything a source publishes, so as your intake grows, narrow it; otherwise the output gets less relevant, not more:

DICTATE OR TYPE PROMPT 4

"When reviewing the RSS feeds, focus specifically on items containing the keyword [your product or topic]."

CHOOSING AN INTAKE CHANNEL

Schedule: the default for horizon scanning. The agent runs on its cadence whether or not anything pinged it.

RSS feed: for sources that expose one. Updates flow in as published; add a keyword filter to keep them relevant.

Email, Teams or Slack: when someone in the business should trigger the agent, e.g. the product team flagging a new higher-risk feature for review. Less useful for a scan that belongs to legal.

No feed? Use the newsletter workaround: sign up to the source with your own address, then set an auto-forward rule so its emails route to the agent's intake address.

It's not either-or. If the sources are well codified and the cadence is right, the schedule alone achieves most of your aims; add RSS where you want changes tracked the moment they land. There's no reason you can't have both.

4 Test it

Say *test run the agent*. It writes a fresh to-do list and works through the scan while you watch, or while you do something else and come back. Then inspect the run:

- **Sources**: the agent saves static copies of the pages it drew from, so you can open exactly what it read.
- **Citations**: good when they point to good sources and let you build confidence quickly. A hundred granular citations on a monthly scan is noise, not rigor. Tell the agent to match the citation depth to the task.
- **Outputs**: the tabular scan in the chat, the memo as a .docx, and both stored in the agent's knowledge base. No template supplied? Give it one and re-run.
- **The schedule**: under your profile's schedules you'll see the recurring run, with its cadence and timezone. Pause it there if you need to stop it for a period, and restart when ready.

A scan is only useful if it turns into work. When an item needs follow-up, assign it without leaving the chat:

DICTATE OR TYPE PROMPT 5

"Create a task for [colleague] to research and write a note on [the development], one week in advance of [the date it takes effect]."

The task appears in your colleague's account, timed to the deadline. Then loop back: tweak, nudge and prune the agent (back through build and inspect) until the output feels genuinely useful. That iteration is the method, not a failure of it.

5 Finalize and hand over

Lock it down and make it someone's job. An agent falls down quietly when the business changes around it: the knowledge base goes stale, a connector changes, nobody notices until the scan is wrong. Decide who manages the agent, who QAs its outputs, and who watches for drift.

One more thing: the people doing the work are the best placed to suggest the next agents to build, and they may not be the people you want in admin roles. Create space for them to build and propose, then ship the good ones.

BEFORE YOU RELY ON IT

- [ ] Company profile reflects the entity
- [ ] Cadence and timezone confirmed
- [ ] RSS intake filtered by keyword
- [ ] Test run inspected; citation depth sensible
- [ ] Instructions name your regulations and regulators
- [ ] Memo delivered as .docx, emailed to the right people
- [ ] Connectors granted at agent level; writes gated
- [ ] An owner named for QA and upkeep

Agents, skills, connectors: who does what

The overlap confuses everyone at first. The shorthand: skills say how, agents say when, connectors say where.

AGENT

Another member of the team: a digital persona that draws on documents, history, sources and people, makes decisions, and works out how the whole job gets done end to end.

SKILL

A recipe: a neatly packaged mini-workflow for one part of the job, sitting beneath the agent. Skills shine for formatting too: "put the output into a color-coded spreadsheet" lives at skill level, not in every chat.

CONNECTOR

Where the agent reaches: the tools it reads from and writes to, typically via MCP. Native connectors are built in; anything with its own MCP server can be added as a custom connector.