

Build your first agent

From the SourceLace manual: A step-by-step guide for people with no technical background. You will set up an agent from a template, try it safely with a dry run, turn it on and check its work. Then you will build one from scratch and learn how to write instructions that work. For the full reference,

SourceLace version 0.1.0 · Made October 7, 2026

The latest version is at <https://sourcelace.onrender.com/docs>

Build your first agent

A step-by-step guide for people with no technical background. You will set up an agent from a template, try it safely with a dry run, turn it on and check its work. Then you will build one from scratch and learn how to write instructions that work. For the full reference, see [AI agents](#).

What an agent is, in one minute

An agent is a job you describe once. SourceLace then does it for you, on a schedule or when something happens.

- It reads your systems **as you**, with your own access. It can never see more than you can.
- It can only do the actions you list: draft emails, make a document, propose record changes or post in Slack or Teams.
- Emails are always **drafts** in your own mailbox. SourceLace never sends email.
- Record changes wait for your approval, unless an admin allowed otherwise.
- Every agent starts as a **draft**. A draft can only do dry runs, which change nothing.

You find agents under **Agents** in the menu. If you see "You can't create agents", ask an admin: your plan or your organization's settings decide who can.

Start from a template

A template is a ready-made agent for a common job. It is the fastest and safest way to start.

1. Go to **Agents** and press **Start from a template** (or open the **Templates** tab).
2. Search or pick a category. Each card says what the template does, who it is for, when it runs and what it needs.
3. Look at the pill at the bottom of the card. **Your systems fit** means you can use it now. **Needs a system you don't have yet** means an admin must add a source first.
4. Press **Use template**. A short wizard with four steps opens.

Worked example: Stale opportunity nudges

Say you are a sales manager at Corvanta. Deals in Salesforce go quiet and nobody notices. You want a short note to each rep every weekday, listing their quiet deals and asking for a next step.

Step 1. Systems. The wizard asks which of your systems to use for each need:

- **CRM:** choose *Corvanta Salesforce*.
- **Mail (for drafts):** choose your Gmail or Outlook. This is where the drafts appear.

If a system shows (*connect it first*), sign in to it on **Data Sources**, then come back.

Step 2. Questions. The template asks a few things. Keep the defaults for your first try:

Question	Default	What it means
Days without activity	21	A deal counts as quiet after this many days with no activity.

Question	Default	What it means
Stages to include	all open stages	Or name stages, separated by commas.
Minimum amount	10000	Smaller deals are left out.
Most drafts per run	15	At most this many reps get a draft each run.
Tone	Friendly	Or Direct.

Step 3. When and how. The schedule is *Weekdays* at 09:00. The time zone is filled in from your browser; change it if needed. Under **Actions** you see "Draft a nudge to an opportunity owner" with *Drafts only, nothing is sent*.

Step 4. Name and results. Give it a name, such as "Corvanta stale deals". Under **Where results go**, keep *A new project named after the agent*. Each run's summary becomes a chat in that project. Press **Create agent**.

You land on the agent's page with the message "Created as a draft. Try a dry run, then turn it on." Now go to [Try it with a dry run](#).

Build one from scratch

Use **Agents → New agent** when no template fits. The editor has five tabs; the  next to each field says what it does, with an example. **Create draft** stays at the bottom of the screen and saves every tab at once.

The job. Give a **Name**, a short **Description** and the **Instructions** (see [Write instructions that work](#)). **Skill to follow (optional)** lets the agent follow one of your organization's skills, such as a house style for QBR decks.

Systems and actions: Systems it may read. Tick only what the job needs. The agent cannot read anything else. Tick your mailbox, calendar, or Slack or Teams here too if an action or trigger uses them; otherwise saving fails with "add it to the agent's sources first".

Systems and actions: What it may do. Add actions with the buttons + **Change records**, + **Draft emails**, + **Make a document (in the results)** and + **Post in Slack or Teams**. See [Choose systems and actions](#).

When it runs. Pick **Only when I press Run**, **On a schedule**, **Before calendar events** or **When something happens**. See [Decide when it runs](#).

Results. Pick a **Results project** so summaries become chats you can ask follow-up questions in.

Limits and alerts. Keep the default limits at first (they are marked recommended). **When a run fails** can also draft an email to you in your own mailbox.

Press **Create draft**. Every later save is **Save as a new version**, and each run records which version it used.

Write instructions that work

Write instructions as you would brief a new colleague on their first day. They are smart, but they do not know your company, your fields or what "good" looks like.

Checklist for good instructions

- **What to read:** which system and which records ("open opportunities in Salesforce").
- **Which ones count:** the filter, with numbers ("no activity for 21 days or more", "amount 10,000 or more").
- **What to compare or work out:** totals, changes since last week, what is missing.
- **What to produce:** drafts, a document or a summary, and what each one contains.
- **How much:** a limit ("at most 15 drafts", "the 10 largest").
- **What NOT to do:** the things a new colleague might get wrong.
- **How to finish:** a short summary with the key numbers.
- **Short numbered steps.** One step, one thing.

Example 1: vague goal

Before:

Keep an eye on the pipeline and tell people to update their deals.

After:

1. In Salesforce, find open opportunities with an amount of 10,000 or more and no activity for 21 days or more.
2. Group them by owner. For the 15 owners with the largest total, draft one email to the owner listing their deals (name, stage, amount, close date, days since last activity) and asking for the next step and its date.
3. Do not put one owner's deals in another owner's draft. Do not guess next steps.
4. Finish with a summary: how many quiet deals, their total amount and how many drafts you made.

Why it is better: it says what to read, which records count, what to produce, how many and what not to do.

Example 2: no comparison

Before:

Tell me about support tickets.

After:

1. In ServiceNow, count open incidents by priority and by assignment group.
2. Compare incidents opened in the last 7 days with the 7 days before, by priority, and say what went up or down.
3. List priority 1 and 2 incidents older than their target (4 hours and 1 day), with number, short description, group and age.
4. Make a one-page document with these tables. Do not include incidents from the "Test" group.
5. Finish with three lines: total open, the biggest change and the oldest priority 1.

Why it is better: "tell me about" becomes counts, a comparison, a list with fields and a clear output. Note that each run starts fresh and does not remember earlier runs, so compare within

the data, not with "last time".

Example 3: too much power

Before:

Fix the account data in our CRM.

After:

1. In Salesforce, find accounts of type Customer with an empty Industry or an empty Billing Country.
2. For each, look only at what Salesforce already holds, such as the mailing country of the account's contacts. Only when the answer is clear, propose a change to Industry or Billing Country. Never change any other field.
3. At most 10 accounts per run. If you are not sure, list the account in the summary instead of proposing a change.
4. Finish with a list of accounts changed, accounts skipped and why.

Pair this with a *Change records* action on *Account* that may only **Update records** for the fields *Industry*, *BillingCountry*, with **I approve each change**.

Why it is better: "fix" is open-ended. The new version names the fields, sets a limit, says when to hold back and keeps every change waiting for your approval.

Tips

- Use the names of things as your system shows them: stage names, object names, group names.
- Say who drafts go to and what the subject should be.
- If text in a record or email tells the agent to do something, it is ignored. The agent treats all data as data. You do not need to write rules about that.
- The instructions cannot give the agent more power. If an action or system is not listed, it is refused whatever the instructions say.

Choose systems and actions

Least access. Give the agent only the systems and actions the job needs. A pipeline report needs to read your CRM and make a document. It does not need to change records or draft emails.

The four actions:

Action	What happens	Approval
Draft emails	Creates drafts in your Gmail or Outlook. You review and send them yourself.	Drafts from a template are made at once. Drafts you add in the editor wait in Approvals and are made when you approve.
Make a document (in the results)	Adds a Word, PowerPoint or PDF file to the run's results.	None needed: nothing outside SourceLace changes.

Action	What happens	Approval
Change records	Proposes creating or updating records. You name the System , the Object (such as Opportunity), the Fields it may set and whether it May Update records, Create records or both.	I approve each change (the default), or Without approval (if an admin allows it) .
Post in Slack or Teams	Posts only in the channels you list. At most 10 messages per run.	I approve each message (the default), or without approval if an admin allows it.

Drafts or record changes? Start with drafts, documents and summaries. They change nothing in your systems, so they are safe to try. Add record changes only when the drafts look right for a few weeks. On the Team plan, agents make drafts and documents only.

Approvals. Keep **I approve each change** for anything that changes records. "Without approval" only works if an admin has allowed that object in **Agents → Settings**, on a plan that includes it. If not, saving tells you so.

Decide when it runs

Only when I press Run. Good while you are still testing.

On a schedule. Choose **How often** (*Every hour, Every day, Weekdays, Every week* or *Every month*), the time **At (HH:MM)** and the **Time zone**. For weekly, pick the day. For monthly, pick a day from 1 to 28 or *Last day of the month*. 08:00 stays 08:00 through daylight-saving changes.

Before calendar events. Choose the **Calendar** (your Gmail or Outlook), the words in **Events whose title contains** (such as *QBR*) and **Hours before the event** (48 is two days). The match ignores capital letters. The agent runs once for each matching event, and cancelled events are skipped. The run is told which meeting it is for.

When something happens. The agent checks a system every few minutes and runs once for each new record it finds, such as each new priority 1 case. Think of it as a saved search that SourceLace repeats for you:

- You write the search in the system's own query language, with {since} where the "since when" time goes. For example: `SELECT Id, Subject FROM Case WHERE Priority = 'High' AND CreatedDate > {since}`.
- **Check every (minutes):** 5 at the most often.
- **Most runs per check:** if many records appear at once, the rest wait for the next check.
- Each record starts one run, once. To run again each time a record changes, fill in **Run again when this field changes (optional)**, such as *LastModifiedDate*.
- When you turn the agent on, or save a new version, records that already exist are noted and skipped.

If you do not know the query language, start from the **Sev 1 escalation** or **Closed-won handoff** template. They already have the right search for each system.

Try it with a dry run

On the agent's page, press **Dry run**. It starts within a minute. A dry run reads your systems and plans everything like a real run, but it changes nothing and drafts nothing.

Open the run from the **Runs** tab on the agent's page. The run page shows:

- **Result:** *Done*, *Failed* or *Stopped at a limit*, marked *dry run*. A note says "Dry run: nothing was changed or drafted. Below is what it would have done."
- **Summary:** what the agent found and what it would have done. Check the numbers against what you know.
- **Changes:** one card per draft, message or record change it would have made, marked *dry run*. Each shows the fields with their value **Now** and **After**. For the stale deals example, check the To, Subject and Body of each draft.
- **Steps:** each tool it used, **Where** (system and object), how many **Rows** came back, the **Outcome** and the **Time**. *Refused* means it tried something it is not allowed to do. *Failed* gives the reason, such as a wrong field name.
- **Credits:** what the run used. Multiply by the runs per day to see the monthly cost.

What to look for:

- Did it read the right records? Zero rows often means a filter or name is wrong.
- Are the numbers right? Compare with a report you trust.
- Are drafts going to the right people with the right content?
- Any *Refused* or *Failed* steps? Fix the instructions or add the missing system.

Edit and dry-run again until you are happy. It usually takes two or three tries.

Turn it on

Press **Turn on**. The status becomes **On** and **Next run** shows when it will run. Your plan sets how many agents can be on at once; if all are in use, pause one first.

Read runs and results

- The **Agents** page lists your agents with their status, when they run, the last run, the share of runs that succeeded in the last 30 days, credits this month and changes waiting for approval.
- Click an agent to see its setup, its instructions and all its runs.
- When a run has a summary or files, it becomes a chat in the agent's project. Use **Open the results chat** on the run page to ask follow-up questions.
- The run page never shows the data the agent read. That data stays in your systems.

The approvals inbox

Agents → Approvals lists every change waiting for you. The tab shows how many are waiting.

- Each card says what would change (such as "Update Opportunity ..."), which agent and run proposed it, why, and until when it waits.

- The table shows each field's value **Now** and **After**.
- Press **Approve and apply** or **Reject**.
- Approving checks the record again first. If it changed since the proposal, the change is marked *stale* and nothing is written; the next run proposes again.
- Changes not decided within 7 days expire.
- **All** shows past decisions too.

Pause or stop an agent

- **Pause** skips the schedule and triggers until you press **Turn on**. You can still run it by hand.
- **Stop** is for when something is wrong: after asking you to confirm, nothing runs, not even a dry run, until its owner turns it on again.
- **Delete agent** (under **More**, on the **Setup** tab of the agent's page) removes the agent with its runs and waiting approvals.
- Admins can stop every agent at once in **Agents → Settings**. Runs in progress stop at their next step.

Common mistakes and fixes

What you see	Likely cause	Fix
"add it to the agent's sources first" when saving	The mailbox, calendar or chat used by an action or trigger is not ticked.	Tick it under Systems it may read .
"cannot run without approval: an admin must allow ..."	You chose <i>Without approval</i> for an object an admin has not allowed.	Choose I approve each change , or ask an admin.
"agents make drafts only"	Your plan does not include record changes.	Remove the <i>Change records</i> action, or ask about a higher plan.
A run shows 0 rows	A filter is too strict, or a stage or field name is wrong.	Check names in your system and fix the instructions.
Steps marked <i>Refused</i>	The instructions ask for something the actions do not allow.	Add the action, or remove that request from the instructions.
<i>Stopped at a limit</i>	The run hit its most steps, records or credits.	Narrow the instructions first. Raise the limit only if the job really needs it.
The agent is <i>Paused</i> with a reason	You lost access to a system, or need to sign in again.	Reconnect on Data Sources , then Turn on .
Drafts go to the wrong people	The instructions do not say who each draft is for.	Name the field to use, such as "the opportunity owner's email".
An event agent did not run for old records	Existing records are skipped when you turn it on or save it.	This is on purpose, so you are not flooded with runs for old records.
Nothing happens at the scheduled time	The agent is still a draft, or the time zone is wrong.	Check the status is On and the time zone on the agent.

Checklist before you turn it on

- [] The name and description say what it does, so others understand it.

- [] The instructions say what to read, which records count, what to produce, how many and what not to do.
- [] It reads only the systems the job needs.
- [] Its actions are the fewest that do the job. Drafts and documents before record changes.
- [] Record changes list only the fields needed, and wait for approval.
- [] Slack or Teams posts list only the right channels.
- [] The schedule, time zone or calendar words are right.
- [] For event agents, the search finds the right records and uses {since}.
- [] A dry run finished, read the right records and the numbers look right.
- [] The drafts or changes in the dry run are what you would have done yourself.
- [] The credits per run, times runs per day, fit your budget.
- [] The results go to a project you will check.
- [] You know where to look: **Approvals, Needs your attention** and the results chat.

Once the agent works well, you can share it: see [Make templates for your team](#).