FAQ: AI agents

From the SourceLace manual: Jobs that SourceLace runs for you on a schedule or when something happens, what they may do, and how you keep them in check.

SourceLace version 0.1.0 · Made November 22, 2026
The latest version is at https://sourcelace.ai/docs

AI agents

Jobs that SourceLace runs for you on a schedule or when something happens, what they may do, and how you keep them in check.

What is an AI agent?

A job you set up once and SourceLace runs for you. For example every weekday morning, two days before each QBR in your calendar, or when a new priority 1 case arrives. It reads your systems with your access, then writes a summary, drafts emails or proposes changes for you to approve. You build, run and approve agents in the SourceLace app: start from a template, describe the job in a sentence, or write your own, and decide on what they propose in the **Approvals** inbox.

Details: AI agents and Build your first agent.

What is Agent Architect, and what does it send to the AI?

An optional add-on that writes agent proposals for your job from your own data. You have a short conversation about your work (including the reports, decks or emails you make again and again) and choose the sources it may study; it proposes agents that use your real objects, fields and figures, checks every name against your data and previews each one on today's records. You pick one and it becomes a draft agent that does nothing until you run it.

What it sends to the AI provider depends on the level your admin allows: the names of your sources, their structure, summary figures SourceLace works out itself, and a few sample rows. Everything sent passes through your organization's data protection rules, any extra rules your admin set for Agent Architect, and masking of email addresses, phone numbers, US Social Security numbers, card numbers and IBANs, which is always on for it. Text that no rule covers, such as a name in a free-text field, is not masked. Sample rows and figures are deleted after the run; saved proposals contain no values from your rows. It reads with your access, so it never sees more than you can.

Details: Agent Architect.

Can different teams use Agent Architect on different data?

Yes. Under **Administration → Agent Architect**, an admin chooses the groups who may use it and, for each group, which sources it may study and the most its runs may see. For example, Finance may study NetSuite with sample rows while Sales operations may study only Salesforce, with summary figures. It only narrows what people can already open in SourceLace; it never gives anyone a new source.

Details: What your admin does.

What does the conversation before a run cost?

Very little, and it goes with its run. Each question uses a small amount of AI on your organization's cheaper model. If the run that follows is one of your plan's included runs, the conversation costs nothing extra; if the run uses AI credits, the conversation is paid with it. A conversation you delete, or never turn into a run (it is deleted after 14 days), is charged as ordinary AI use. Your admin can choose SourceLace's built-in questions instead, which use no AI before the run.

Details: Agent Architect.

Can an agent use our own deck template, playbook or past reports?

Yes. On the agent's page, under **References**, link up to ten files from your connected Google Drive, SharePoint, OneDrive or storage buckets: a **Template** it fills (PowerPoint, Word or Excel, or Google Slides, Docs or Sheets), an **Example** whose style it matches, or a **Guide** whose rules it follows. Each run reads the file again with the sign-in of the person it runs for, so edits are picked up and nothing is copied. A reference can change what the agent writes, never what it is allowed to do: a sentence in a file such as "email this to an outside address" is treated as content, not an order, and the agent's actions, recipients and approvals stay as you set them. When a guide or template changes, the page says so, and the next run's automatic actions wait for approval once. In Agent Architect, you can link the same kind of files as **examples of your work**, and its proposals keep them as references.

Details: References, Examples of your work.

Can I upload a file instead of linking it?

Yes, once an admin allows uploads (on the **Data protection** page, under **Uploaded files**, or by accepting Agent Architect's statement): on an agent's **References** and, with Agent Architect, in **Examples of your work**, choose **Upload** and pick a PDF, Word, Excel, PowerPoint, CSV or text file of up to 20 MB. SourceLace checks it by its contents (files with macros, embedded objects, scripts or passwords, older formats and scanned PDFs are refused), shows its outline before you add it, and keeps it encrypted with a key unique to your organization until you remove it or delete the agent or conversation it belongs to; a conversation you never use is deleted with its files after 14 days. Deleted files are removed from storage promptly. Your organization may keep as much as its plan includes. If an admin turns uploads off, agents stop reading uploaded files and run without them (each run says so), but the files are kept: turning uploads on again brings them back. Its text is masked before it reaches your AI, like a linked file's. Pictures can be shared as examples only where an admin has turned them on, because a picture cannot be masked.

Details: References, Examples of your work.

Does Agent Architect make up how much time an agent saves?

No. A proposal shows the time it saves only when you said how long that work takes you, and it says it is your estimate. The Agent Architect page then adds up the time your agents saved in the last 30 days: each real run whose result you used (you rated it good, approved what it proposed, or it reached its destination) times your estimate. Dry runs and runs nobody used are not counted.

Details: Your agents and Improve it.

Can I design one agent for my whole team?

Yes, if your admin allows it and your organization uses shared agents. Tick **Design for my team** before you start. Each proposal becomes one shared agent that your team subscribes to; each person's runs read only what that person can see, so an account manager's Monday list comes from their own accounts. You see who subscribed and how many runs, never anyone's results.

Details: Design for your team.

Can our admin see how much each person uses Agent Architect?

No. **Insights** show totals for your organization and its groups only, never one person's figures: a group is shown only when at least five of its people used it, and while fewer than five people use Agent Architect, what came of it is not shown at all. Time saved is always people's own estimate. Each person sees their own figures on their own Agent Architect page. SourceLace's own staff see your organization's totals, to help you get value from it; when only one or two people use it, those totals can be their figures.

Details: Insights.

How do I find an agent's runs from a certain day or hour?

Open the agent and its **Runs** list, then use **When** (Last hour, Today, Last 7 days, Last 30 days, This month, or a custom start and end with date and time) and **Kind** (real or dry runs). Times are in your own time zone, named under the filter. The period is in the page's address, so you can share the link, and **Clear filters** shows every run again.

Why does a run say "Waiting for approval" or "Lost"?

A run waits when it proposed changes that need a person; it holds nothing meanwhile and carries on once every change is decided, withdrawn or expired, then says what happened (done; done but changes not made, and why; failed; stopped). Lost means the background work that ran it stopped and no server could take it over: nothing it had not reported was made, and you can run the agent again.

Details: What happens when....

How do we avoid costly mistakes with agents?

Start with one small job, give the agent only the systems and actions it needs, keep record changes waiting for approval, dry-run it, and let it try in shadow for a few days before it runs on its own. For cost, read less, run less often and set low limits on each run. Then watch the first week of runs and the Usage page.

Details: Agent best practices.

Which plans include agents?

Team (drafts and summaries only), Business and Enterprise. Each plan sets how many agents may be on at once. Starter does not include agents.

Details: The plans.

Can an agent see or do more than its owner?

No. An agent reads and writes as its owner, through the same access by group, change settings and audit trail as when the owner asks a question. It can use only the sources and actions listed for it; anything else is refused. If the owner loses access or leaves, their agents pause, and an admin can give them to someone else (see Give an agent to someone else).

Details: Safety.

Which systems can send events to SourceLace?

An event is a message a system sends the moment something happens, so an agent can start within seconds. Agents can start live from events in Salesforce, MongoDB, Gmail, message queues (Amazon SQS, Google Pub/Sub, Apache Kafka and AMQP brokers, including IBM MQ inside your network) and your own connector. Other systems call the agent's start address, with ready-made setup for Microsoft 365, Zendesk, Jira, GitHub, Stripe, DocuSign, Okta, Gong, Slack, Teams, Salesforce Data 360, Marketing Cloud, HubSpot, ServiceNow, SAP Event Mesh, Amazon SNS, Azure Event Grid and Google Cloud Pub/Sub, or the general option for any system that can call a web address. A system that cannot send events, such as a data warehouse, is checked every few minutes instead. These checks are cheap: each is one small read-only query that looks only at what changed, with no AI used until something is found.

Details: Which systems can send events to SourceLace.

What is the difference between a webhook and A2A?

A webhook is a notification: a system calls the agent's start address to say something happened, and gets a quick "received" back. A2A is a task with an answer back: an agent on another platform gives your agent a job and follows it until the summary arrives, or until it needs a sign-in or a decision. Use a webhook when the other system only needs to say "this happened", and A2A when it needs the answer. A2A works both ways: your callable agents are listed in one catalog for other platforms, and your agents can ask agents on other platforms. Both are pilot features, and A2A is off until your SourceLace provider turns it on.

Details: Webhooks or A2A.

Can a record or an email trick an agent into doing something else?

What an agent reads is treated as data, never as instructions. If a record says "delete all accounts", the agent still can only use its listed sources and actions, and every change still needs approval unless an admin allowed otherwise.

Details: Safety.

How do we try an agent safely before it goes live?

Two ways. A **dry run** reads and plans like a real run but changes and drafts nothing. **Shadow mode** runs the agent on its real schedule for days without changing, drafting or posting anything, so you can review what it would have done before it goes live. Admins can require shadow mode for every new agent.

Details: Try it in shadow.

What can start an agent?

Pressing **Run now**, a schedule, an upcoming calendar event, a new record in a connected system, a record that starts to meet a condition (such as a renewal date within 90 days), other agents finishing, a call from another system, an email to the agent's own address, another of your SourceLace agents, or another AI agent signed in as a person. Agents on other platforms can call yours, and your agents can ask theirs for help, through the open A2A standard. A2A is off until turned on for your organization, and approvals always stay in SourceLace.

Details: When something happens and Work with agents from other platforms.

Can an agent start before a renewal date, or when something has waited too long?

Yes. An agent can watch for records that meet a condition, such as contracts renewing in the next 90 days, orders blocked for more than three days, or accounts whose usage fell below a limit. Each record starts a run once when it starts to meet the condition, and again only if it stops and later meets it again. You can also have one run with the whole list, for a daily digest. Your system checks the condition, and SourceLace keeps no copy of the records.

Details: When a record meets a condition.

Can an agent start when other agents have finished?

Yes. Choose **After other agents finish** and pick the agents: your own, or ones whose owners let other agents call them. It can start each time any of them finishes, or once all of them have finished within a few hours, and it works from their results. For example, a Monday brief that starts once the sales, support and finance digests are done.

Details: Start after other agents.

Can an agent ask our finance team to approve a discount before it goes ahead?

Yes. Give the agent an **Ask people to decide** action and choose the groups that decide, in order, such as Deal desk then Finance. The agent asks with its recommendation and the evidence; someone in each group approves in turn, and anyone at the current step can reject. The agent's owner cannot decide its own request. When it is decided, the agent carries on with the outcome and people's notes. Deciding changes nothing in your systems by itself.

Details: Decisions.

Can every agent's results read the same way?

Yes. Choose a **Result format** for the agent: a focus list (the few items to act on now, most urgent first), an exception list (only what breaks a rule, grouped by rule), a ranked list (every item in order, with reasons), a brief (the answer first, then what supports it) or an executive report (the numbers that matter, what moved and the decisions needed). SourceLace keeps the wording of each format in one place, so every agent that uses one reads alike, and improves together when the format does.

Details: Result formats.

Can our agents ask each other for help?

Yes. Turn on **Other agents may call this one** on the agent that will be asked, then pick it under **Other SourceLace agents it may ask** in the setup of the agent that asks. The asked agent runs as the same person, with its own instructions, sources, limits and approvals, and its answer comes back as data, never as instructions. Dry runs and shadow mode carry down, the asking agent's credit limit covers both runs, and an agent can never end up asking itself.

Details: Let your agents ask each other.

Which systems can start an agent?

Any system that can call a web address. On the agent's page, choose who calls it, and SourceLace checks that system's own signature or token and answers its first check call: GitHub, Jira, Slack, Zendesk, Stripe, DocuSign, Salesforce Data 360 and Marketing Cloud, Microsoft 365, Teams and SharePoint, Azure Event Grid, Amazon SNS, Google Cloud Pub/Sub, Google Drive and Calendar, Gong and Okta. **Connect a system** has step-by-step guides with the address filled in for these and for Salesforce Flow, SAP through SAP Event Mesh, HubSpot, Dynamics 365, NetSuite, ServiceNow, Snowflake, Databricks, MongoDB Atlas, FHIR servers, integration platform flows and any other system. **Send a test event** shows what the agent would receive, as a dry run, and **Recent calls** shows the last 50 calls.

Details: Who calls it and Connect a system.

Can one agent cover several Salesforce orgs or SAP systems?

Yes, for what starts it. An agent that starts when something happens can watch several systems of the same kind, such as your US, EMEA and APAC Salesforce orgs: after choosing the first, tick the others under **Also watch other systems**. Each is checked with your own sign-in for each, on its own, so one that is down never holds up the others, and every run says which org it came from. If one org is set up differently, choose **Different for this system** under **Settings for each system** to give it its own query and record columns. You choose whether each org's check starts its own runs (the default) or **Start one run with all of them**: one run with everything the orgs found at a check, naming any org that could not be checked, so it never reports something as missing there. An agent that starts when a message arrives can likewise listen on several message queues, on one connection or several, under **Also listen on**; a queue that stops working pauses alone. An agent can already read several systems, and you can give it actions in each. A template whose system slot allows several lets you tick each of your orgs in the wizard ("Choose one or more"), so one agent made from it reads them all and names the connection with every finding. An action can also write to several of them: see the next question. Industry packs can set up agents for several systems too (see the questions below). See One agent for several systems.

Can an industry pack set up agents for several of our systems, such as a Salesforce org for each region?

Yes, where the pack says so. On **Business setup**, **Industry packs**, a part that takes several systems, such as the CRM in **Sales and revenue**, shows a tick box for each of your systems of that kind, all ticked to start with. Group access, data protection and business-term hints are suggested for every one you tick. Then choose **One agent that covers all of them** (an agent whose template can read several systems reads them all in one run and names each finding's system; an agent that can read only one uses the first you ticked, and the page says so) or **One agent for each system**, named with the system's name, such as *Weekly pipeline review (EMEA org)*. If you connect another org later, run the pack again and tick it: SourceLace suggests adding it to the agents the pack made, as items in your review queue, each a new version of the agent with its mode and approval settings unchanged. An agent you made yourself is left alone.

Details: Industry packs.

Can one agent send to several places, such as an email and a Slack message, or update every org?

Yes. An agent can draft an email and post in Slack in the same run, and one action can go to several systems: the same record change in each regional Salesforce org, the same draft in two mailboxes, or one alert in Slack and in Microsoft Teams. In the agent editor, choose **Add another destination** on the action; with a template whose slot allows several systems, the action is made once for each system you tick. Each destination has its own approval and its own entry in the audit log. If one destination does not work (a change is refused, or someone rejects its approval), the agent carries on with the others, nothing already made is undone, and the run's page lists how each destination went under **Destinations**. Every destination counts as one of the agent's 20 actions. See One action, several destinations.

Can an agent email or post the report it made, or do it without me approving each one?

Yes. An agent that makes a Word, PowerPoint, PDF, Excel or CSV file can attach it to a new email draft in Gmail or Outlook, or post it with a message in Slack or a Microsoft Teams channel: tick **Can also send files this agent makes** on its email draft or chat post action. It can also save it to a folder (next question). You choose per action whether a person approves each one first (the approval shows each file's name, type, size and where it goes) or it goes by itself; a Slack post goes by itself only when an admin allows that action without approval. When Gmail, Outlook or Slack asks SourceLace to wait, it tries again by itself for up to a day, and it never sends the same file twice. The run's page lists each one under **Deliveries**, with **Retry** for one that failed. For Slack and Teams, an admin adds `files` to the source's objects that can be changed first. See Sending files the agent made.

Can an agent save its reports to SharePoint, OneDrive or Google Drive?

Yes, to one folder per action. Add a **Save files to a folder** action with the folder's id; each file is saved as the person who runs the agent, with only their access, and never replaces a file already there. A file waits for approval (showing its name, type, size and the folder) unless an admin lets that exact folder be saved to by itself, so automatic saves to a reports folder can be on while record changes still need approval. A SourceLace admin turns saving on for the source first, and Google or Microsoft then asks people for permission to add files. In Google Drive the file can be saved as a PDF, Word, Excel, PowerPoint or CSV file, or as a Google Doc, Sheet or Slides file, whichever the agent's instructions (or your request in a chat) ask for. See Saving files to a folder.

Can I start an agent from Slack or Teams?

Yes, as an optional way to start your agents: a Slack slash command or a mention of a Teams outgoing webhook starts the agent as the person who typed it, found by their verified email in your organization. People who are not found are told to sign in to SourceLace first, and nothing starts. They get a link to the run; changes still wait for approval in SourceLace, and the app stays the place to build, approve and review agents.

In Slack, someone makes a Slack app with a slash command whose Request URL is the agent's start address, and pastes the app's Signing Secret and its bot token (with the scopes `users:read` and `users:read.email`) into SourceLace. In Teams, a team owner creates an outgoing webhook with the start address as its Callback URL and pastes its security token, and your Microsoft Entra administrators register a small app with the Microsoft Graph permission User.Read.All so SourceLace can tell who mentioned it. Teams outgoing webhooks work only in a team's public channels. Office 365 connectors and the Workflows app are not used for this.

Details: Start from Slack or Teams, with step-by-step guides for Slack and Teams.

Can I start an agent by sending it an email?

Yes. Each agent can have its own email address, made from random letters and digits so nobody can guess it. Give it to any system, monitoring tool or partner that sends email alerts. Only email from the addresses and domains you allow, whose sender is verified, starts the agent; everything else is dropped without a reply. The agent receives the subject, sender, date, text and the names and sizes of attachments, as data and never as instructions, and reads an attachment only when it needs to. Attachments are kept encrypted with the run and deleted with it; nothing else of the email is kept. You can rotate or turn off the address at any time, and **Recent emails** shows what each email did. Your organization's admin may need to ask SourceLace to turn it on.

Details: Start by email.

How do we stop an agent?

Pause skips its schedule and triggers. **Stop** blocks it completely until its owner turns it on again. Admins have one switch that stops every agent at once; runs in progress stop at their next step. Every agent also has limits on steps, records changed, credits and runs per day.

SourceLace support can also stop an agent, or all of your agents, in an emergency, such as an agent stuck in a loop. You see why on the Agents page and in your audit log, and only SourceLace support can lift that stop.

Details: Pause or stop and Limits.

What is kept from an agent run?

Each run's times, counts and steps (never the rows it read), and, encrypted, its summary, any changes waiting for approval, what it still has to send (a draft's or post's text and the names of its files, never the files themselves, which stay in its results chat), the attachments of an email that started it and, until the run ends, its progress so far, so an interrupted run carries on where it was. Runs are deleted with the agent and after your organization's retention period.

Details: Safety.

Can several people use the same agent?

Yes, as a shared agent. One person shares an agent, and others subscribe. It runs once for each subscriber, as that subscriber, with their own access, never the owner's. Admins can subscribe a whole group.

Details: Shared agents.

Can our agents run inside our own network?

Yes, on a **SourceLace Runner** in your cloud account or data center. It connects out over HTTPS only, uses your own AI key, and keeps people's source sign-ins and every result there. It sends SourceLace only each run's status and counts. Runners are offered on the Enterprise plan only, on request, with their own fees and terms.

Details: Runners and Who can have a runner.

Can agents start when something happens in Salesforce?

Yes. Choose **When Salesforce sends an event** and pick the event: a platform event, a platform event channel, or the changes to an object (or to every object). Nothing needs to be set up in Salesforce, and the list of events comes from Salesforce itself. The agent reads the events with its owner's own Salesforce sign-in, so it sees only what that person may see. If the sign-in expires, the agent stops and asks its owner to reconnect. Details: Start when Salesforce sends an event.

Can agents start from JMS, IBM MQ or Kafka?

Yes. Agents can start from **Apache Kafka** topics, and JMS brokers connect through **AMQP 1.0**, the open standard most of them speak: ActiveMQ and Artemis, Solace, RabbitMQ 4, Azure Service Bus and IBM MQ (through its AMQP channel). **Amazon SQS** queues and **Google Pub/Sub** subscriptions work too, and IBM MQ inside your own network through the on-site SourceLace Agent, with no inbound port opened. Each message that arrives on the queue you choose starts the agent, and the message waits in your queue until its runs have started. Details: Start agents from your message queues.

Can agents start when data changes in MongoDB, in Gmail or in our own system?

Yes. Choose **When a MongoDB document changes** to start an agent when a document in a collection is added, changed or removed; it reads the changes with its owner's own MongoDB login. Choose **When an email arrives in Gmail** to start an agent when an email arrives in the owner's inbox, or in one of their labels; it reads the owner's own mailbox with their own Google sign-in. If your developers built a connector with the Connector SDK, choose **When your own connector sends an event**: the connector offers its topics, and sends each event for the agent's owner, so the owner sees only what they may see. All three carry on where they stopped after a restart. Details: Start when a MongoDB document changes, Start when an email arrives in Gmail and Start when your own connector sends an event.

How do I know an agent is ready?

Test it. With agent testing on, each agent has a test set: cases that say what to ask it and what a good result looks like, which you write, have the AI draft, or make from reviewed shadow runs. **Test** runs every case with your real data, but every change, draft and message is only shown, never made, and an AI grader scores each case from 0 to 100 with a reason. The agent goes live, and a changed version is published, only when its latest test run reaches the pass mark (80 unless your admins choose another; an owner may raise it for one agent). With shadow mode on too, it also needs its reviewed shadow runs. Details: Test before going live.

What happens when I rate an answer?

Your thumbs up or down, and your note if you write one, are kept with which answer you rated, never the answer itself. They are encrypted and deleted with the chat or agent, and in any case after your organization's retention period for saved chats and agent runs. You can change or clear your rating. An agent's owner, its co-owners and your admins see the ratings of its runs; the owner may turn a thumbs-down into a test case with the corrected answer, or choose a well-rated run as an example for the agent, never automatically. Thumbs-down also help Learn from use decide what to suggest first. Details: Rate answers.

Can an admin give an agent to someone else?

Yes, for an agent that is not shared, for example when its owner has left. The new owner must be in your organization and must already have connected, with their own sign-in, every system the agent uses; if one is missing, SourceLace names it and nothing changes. The agent keeps its history, versions and settings and runs as the new owner afterwards. It is paused until they turn it on, changes it proposed before are cancelled and proposed again on its next run, and an agent that had been tried in shadow mode starts that again. The hand-over is in the audit log. Details: Give an agent to someone else.

Can an admin start someone else's agent from another AI agent?

Only if the agent's owner allowed it. From another AI agent, SourceLace's `run_agent` tool or an agent on another platform, an agent runs with its owner's access, so an admin who is not the owner is refused, and the refusal is in the audit log. The owner can allow it by turning on **Callable by other agents** for that agent; an admin cannot turn that on for them. In the app, an admin can still press **Run now** on any agent. Details: Start from another agent.

How do I know whether my agents are doing useful work?

Open **Agents → Performance** (admins; an agent's owner opens it from the agent's page). It shows, for the last 7, 30 or 90 days, how many runs completed, how many changes were made and outputs delivered, the success rate, how long the whole journey and each stage took (queue wait, working, waiting for people), each kind of step's success and speed, approvals and how long they took to decide, and plain-English hints on where to improve. Every figure is defined in Measuring your agents. It is worked out once a day from counts and durations, never from your data.

Why does a run's journey show mostly waiting?

Because the run paused for people. An agent that proposes a change it may not make by itself holds nothing while it waits for someone to approve, reject or let the proposal expire, and that time is part of its journey. If waiting is most of the time, the agent is fine and the process around it is slow: approve sooner, let the agent make routine changes itself, or shorten **Days a change waits for approval** on the Limits page. Details: The journey.

Can I change the columns of the agents list?

Yes. Drag the edge of a column heading to make it wider or narrower, and use **Columns** above the table to hide or move columns. The name column always stays. SourceLace remembers your layout for you on every device you sign in on, and **Reset columns** goes back to the standard one. A long description is cut to two lines; hover over it to read all of it. An agent's **Runs**, **Versions**, **Ratings** and **Tests** tabs, **Agents overview**, a shared agent's **Subscribers** and **Runs**, the lists on **Agent performance**, the **Matching** rules, the **Audit log**, **People**, **Groups** and **Connections** work the same way.

Details: AI agents.