FAQ: Governance and admin controls

From the SourceLace manual: The controls your organization's SourceLace admins have: roles, who can get in, access by group, stop switches, limits and setup reviews.

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

Governance and admin controls

The controls your organization's SourceLace admins have: roles, who can get in, access by group, stop switches, limits and setup reviews.

What roles are there?

Admin and **member**. Members connect sources with their own logins, ask questions, keep their own chats, skills and projects, and share skills with named people and their own groups. Admins also manage people, email domains, sources, groups, single sign-on and settings, share skills with the whole organization, and read the audit trail, usage and invoices. Every organization always keeps at least one admin. An admin can also name **group managers**: people who add and remove the people of a group, and share skills with it, without being admins. See Group managers.

Details: Roles.

How do we control who can sign in?

Only people your admins invite, or people with a verified address at one of your organization's email domains, can sign in. Public domains such as gmail.com are refused. With automatic provisioning, your identity provider decides who is in.

Details: People and sign-in.

Can I limit a source to some people?

Yes, with access by group (Business and Enterprise). Each group gets **no access**, **read** or **read & write** per source, and per object if you like, such as "read all of Salesforce, but not Opportunity". Someone in several groups gets the most any of them gives. Groups can be kept in step with your identity provider (they then show as synced and their people are locked), or made by hand, where you can add people by pasting emails or uploading a CSV.

Details: Access by group, SCIM provisioning and Group sync.

Can access by group give someone more than their own access?

No. It only narrows access, never widens it. People still sign in to each system with their own account, so they see a record only if both their groups and the system allow it. Once access by group is on, someone in no group sees no sources at all, admins included: their **My connections** page says so. To fix it, give one of their groups access in **Groups**, or turn access by group off.

Details: Access by group.

Can people add their own accounts, and can we control it?

Yes, once an admin turns it on; it is off for every organization until then. A person can then add their own extra account of a system your organization already uses (a second Gmail or Outlook mailbox, a personal Google Drive, SharePoint and OneDrive, Google Sheets) under **My accounts** on their My connections page, without an admin. Only that person sees and uses it; it is not one of your sources and does not count against your plan's connectors. Only systems with a personal sign-in on the Standard connector class can be added this way, never Business or Enterprise systems. Your access by group and data protection rules for that system apply to it, and a person never gets more on their own account than on your connection of the same system. Under **Personal accounts** on Connections, admins switch it on or off for the organization or for a group, set how many accounts each person may add (3 by default, 10 at most), choose the systems, decide whether changes are allowed through them, see how many exist per system and who owns each one, and can revoke any. Removing or deactivating a person removes their personal accounts too.

Details: Personal accounts.

Can I limit which rows a group sees?

Yes, with read rules (part of access by group). On a group's **Sources and objects** tab, add a read rule for an object: pick a field, an operator and a value or list, such as Region **is** EMEA and Stage **is not** Closed, and optionally the most rows one answer may return. A preview says in plain English what the group will see. The rules apply in the chat, agents, AI apps and Search everything. Someone in several groups gets the most permissive result of the groups that may read that object. They are most useful for warehouses and databases reached through one shared account; for systems people sign in to themselves, the system's own permissions still apply first. Only admins change them.

Details: Read rules.

How do we remove someone?

Remove them on the **People** page. Their SourceLace sign-in and their AI app connections stop working at once, and their sign-ins to your systems are deleted. If your organization admits everyone at an email domain, they can come back as a member if they sign in again (signing in to everything afresh: nothing they had before starts working again), so also remove them in your identity provider or remove the domain. With automatic provisioning, deactivating them in your identity provider does all of this for you.

Details: People and sign-in and When someone leaves.

Can we act quickly in an incident?

Yes. Admins can, at once, remove a person, remove a source, turn changes off, pause or stop AI agents (one switch stops them all), and revoke a SourceLace Runner. People can also disconnect any source themselves.

Details: Can you stop access quickly in an incident?.

Who decides what the AI may see?

Your admins, on the **Data protection** page. They can hide fields, mask values, and choose groups whose members still see the real values in the SourceLace app. Everything is off until an admin turns it on.

Details: Data protection.

Where did Organization settings go?

The old **Settings** page and the **Organization settings** item in the account menu are now separate pages under **Administration**, one for each topic: **AI model**, **Sign-in** (single sign-on and provisioning), **Sharing** and **Chat defaults** (instructions, starter questions, how long chats are kept and how answers look). Old links and bookmarks still open the right page. The account menu keeps only what is about you: **Your settings**, **Your instructions** and **Sign out**.

Details: Finding your way around Administration.

What is the Health page?

Health (admins, under **Administration**, **Trust and health**) is one look at whether what your organization relies on is working: connections that failed lookups, agents whose runs failed, whether the audit trail checks out and your streaming destinations deliver, seats and use against your plan, and sign-ins that were refused. Problems come first, and each tile links to the page that has the detail. The tiles only read what SourceLace already keeps and show no data from your systems. You can also be emailed, at most once a day, when something new needs attention: add up to five addresses in the Alerts section at the bottom of the page and choose which kinds you want. Nothing is sent until you add an address.

Details: Health and Emails when something needs attention.

Can we control who creates AI agents?

Yes. Under **Agents → Settings**, admins choose who can create agents: everyone, people in some groups, or admins only. They can also require every agent to run on your own SourceLace Runner, require shadow mode first, and set which objects agents may change without approval (Business and Enterprise).

Details: For admins.

Can we set stricter limits than SourceLace allows?

Yes. On the **Limits** page, admins can search every limit, browse them by topic, and choose stricter values, such as shorter sign-in sessions, shorter query times, smaller files, and lower limits for AI agents. You can never choose more than SourceLace allows for your plan.

Details: Limits.

How do we know our setup is sound?

The **Business setup** page gives a readiness score from 0 to 100, with checks for sign-in, access, data protection, sources, AI model, agents, skills and business terms, and audit and retention. Most checks offer a fix, and no fix changes anything until an admin approves it.

Details: Readiness score.

Can we copy a setup from a test organization?

Yes. An admin downloads the setup as one file: groups and what they may read, data protection rules, business terms, skills, agent settings and more. It never holds people, passwords, keys, chats or data from your systems. Bringing a file in puts each difference in a review queue for an admin to approve.

Details: Setup files and Review queue.

Which account does the setup assistant read my system with?

The sign-in you connected to that source. For a source with a personal sign-in that is your own account and your own permissions. A source with a shared account is read with that shared account's permissions. Either way a scan keeps only counts, never values. The assistant page lists your organization's sources and says whether each is connected or **Connect first**; connect a source under **My connections** before you scan it.

Details: Setup assistant.

Will the setup assistant keep suggesting what I rejected?

Not groups, terms, agents or "turn on access by group": when you scan again, it leaves out what your organization rejected before and tells you how many of each kind it skipped. Data protection rules are the exception and are always suggested again, because a field that looks sensitive may need protecting after all. SourceLace remembers a rejection only while the review batch it was in is kept (30 days by default; an admin can change it), so after that the suggestion can come back. To see skipped suggestions again sooner, tick **Include suggestions I rejected** before you scan, or choose **Scan again with them**.

Details: Setup assistant and Review queue.

Can SourceLace change our settings without us?

Your SourceLace contact manages your plan, admins and contract, and every change is in your audit trail. Setup changes they suggest wait in your review queue like any other suggestion. You decide whether they may approve their own suggestions, and you can turn that off at any time.

Details: Help from your SourceLace contact.

Can admins manage SourceLace from chat?

Yes, mostly. Admins can add people, sources and groups by asking their AI app in plain words. Single sign-on is set up in the app only, so a client secret never has to be typed into a chat. Every admin change is recorded in the audit trail.

Details: Admin guide.

Does SourceLace learn from what our people ask?

Only if an admin turns on **Learn from use** (off by default). SourceLace then notices questions several people repeat in the SourceLace app and suggests a skill or agent for an admin to review. Admins never see who asked, and see example questions only if they turn that on too. It stays inside your organization and is never used to train AI models.

Details: Learn from use.

How do I let people attach photos in the chat?

Open **Chat defaults**, the **Chats** tab, **Attachments in chat**. Nobody can attach a photo until an admin reads the statement, ticks the box, types their full name and chooses **Accept**, and then chooses who may attach: everyone, or only the groups you pick. The acceptance is written in your audit log with who and when, and your other admins are emailed. Photos work today only with the AI included with SourceLace and your own Anthropic key; with OpenAI, Azure OpenAI, Amazon Bedrock, Google Gemini or Mistral nobody gets the button, and the box says so. You can withdraw it at any time; that also clears who may attach, so accepting again starts from nobody. Photos are kept encrypted with the chat, for as long as the chat is kept; after a deletion, storage can keep an older copy for up to 30 days.

Details: Attachments in chat.