

# Security questions

From the SourceLace manual: Direct answers to the questions in vendor security reviews and questionnaires (such as SIG Lite or CAIQ), grouped the way those questionnaires are. Each answer describes what SourceLace does today; where something is not in place yet, it says so. The full detail is on Sec

SourceLace version 0.1.0 · Made October 6, 2026

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

# Security questions

Direct answers to the questions in vendor security reviews and questionnaires (such as SIG Lite or CAIQ), grouped the way those questionnaires are. Each answer describes what SourceLace does today; where something is not in place yet, it says so. The full detail is on [Security and privacy](#). For anything not answered here, or to arrange a review call, email [security@sourcelace.com](mailto:security@sourcelace.com).

## Company and certifications

### **Does SourceLace hold SOC 2, ISO 27001 or any other certification?**

No. SourceLace does not yet hold any security certification or attestation: no SOC 2 report, no ISO 27001 certificate, and no HIPAA attestation. Please do not record SourceLace as holding one in your vendor register.

### **Do you sign a data processing agreement (DPA) or a Business Associate Agreement (BAA)?**

Not yet. A GDPR data processing agreement is not yet available, and there is no HIPAA Business Associate Agreement programme (see [Healthcare and HIPAA](#)).

### **Do you have written security policies, staff background checks or security training?**

Not yet formalised as written policies we can share. Ask [security@sourcelace.com](mailto:security@sourcelace.com) for company details a questionnaire needs, such as the contracting entity and address.

## Hosting and data location

### **Where does SourceLace run, and where is our data?**

SourceLace's cloud runs on **Render**, in its **Oregon (United States)** region: the application, its database and, when enabled, the short-term store for query results. Data is hosted in the United States only today; there is no EU or other regional hosting yet.

### **Will the hosting provider change?**

It may. If it does, this page will change, and the subprocessor list on [Security and privacy](#) will be updated before any new provider handles your data.

### **Can we run SourceLace ourselves?**

Yes, in two ways:

- **Self-hosting:** your IT team runs the same SourceLace image in your own cloud account or data center, with your own database. Self-hosting is in Preview; ask [support@sourcelace.com](mailto:support@sourcelace.com). See [Self-hosting](#) and [Self-hosted installs](#) below.
- **Runners:** people keep using SourceLace's cloud, but chosen AI agents run inside your network, with your own AI key, and their results stay there. See [Runners](#).

## Is customer data separated between organizations?

Yes. Everything SourceLace keeps for your organization is tied to it, every request is checked against the signed-in person's organization, and each organization's secrets, chats and agent data are encrypted with that organization's own key (see [Encryption at rest](#)). An email address belongs to one organization only.

## Encryption

### Encryption in transit

SourceLace is served over HTTPS only. Connections from SourceLace to your systems use TLS with certificate checks; for databases SourceLace also checks the certificate was issued for the host name your admin entered, and TLS cannot be turned off for a database other than one on the same computer as SourceLace. A runner connects to SourceLace over HTTPS only.

### Encryption at rest

Each organization has its own random data key. These are encrypted with it, using Fernet (AES-128 in CBC mode with an HMAC-SHA256 integrity check):

- each person's sign-ins to your systems (access and refresh tokens, database passwords);
- your admins' secrets (single sign-on, app registrations, custom connector secrets, Amazon S3 keys) and your own AI key;
- saved chats and the files made in them, and skills;
- AI agents' settings, run summaries, triggers and changes waiting for approval;
- query results, while they are held in the short-term store.

Other records, such as the names and roles of your people, your source settings, usage counts and the audit trail (which holds metadata only), are kept in SourceLace's database without your organization's key; their protection at rest is the hosting provider's.

Tokens SourceLace issues (to AI apps, runners and on-premise agents) and one-time sign-in codes are stored only as SHA-256 hashes, so a copy of the database cannot be used to sign in.

### Who holds the keys?

See [Keys: who brings each one, and where it is kept](#). In short: your organization's data key is itself stored only in encrypted form, under a master key that is a secret on SourceLace's servers and never in the database. On a self-hosted install, you hold the master key.

### Can keys be rotated? Can we bring our own key?

The master key can be rotated: a new one is added and the old one kept only for reading. Rotating an individual organization's data key is not yet available. On SourceLace's cloud, a customer-managed key (BYOK or a key in your own KMS) is not yet available; Self-hosting is the way to hold the master key yourself today. You do bring and control your own AI key, if you choose to use one.

# Identity and access

## How do our people sign in to SourceLace?

With Google, Microsoft, or your own single sign-on: OpenID Connect (Business and Enterprise) or SAML 2.0 (Enterprise, in Preview), with providers such as Okta and Microsoft Entra ID. SourceLace has no passwords of its own and never sees your people's passwords. You can require single sign-on for your email domains. See [Single sign-on](#).

## Do you support MFA?

Through your identity provider. Because SourceLace has no passwords of its own, multi-factor authentication is whatever your Google, Microsoft or single sign-on setup requires. SourceLace has no MFA of its own.

## Who can get in?

Only people your admins invite, or people who sign in with a verified address at one of your organization's allowed email domains (public domains such as gmail.com are refused). Removing someone on the **Team** page ends their sign-in and deletes their sign-ins to every source. See [People and sign-in](#).

## What roles are there?

**Admin** and **member**. Admins manage people, sources, groups, single sign-on, settings and the audit trail; members use the sources they may use. Every organization always keeps at least one admin. See [Roles](#).

## Can access be limited to some people? Do you sync groups?

Yes, on Business and Enterprise. Admins can limit each source to named groups, and with single sign-on, groups are kept in step with your identity provider at each sign-in. See [Access by group](#) and [Group sync](#).

## Does SourceLace use a shared service account to reach our systems?

Not for most systems. Each person connects them with their own login, so the system's own permissions apply to every call and people only see and change what they already could. Some sources use one **shared account** an admin sets up: Amazon S3, lake tables, Amazon Athena, Gong, custom connectors, SAP S/4HANA and SuccessFactors (SAP's demo sandbox only, for now), and, if the admin chooses it, Amazon Redshift and REST or SOAP sources. Each source is labelled "Personal sign-in" or "Shared account" in the app and for AI apps; see [Personal sign-in and shared accounts](#).

## How long do sessions and tokens last?

| What                                                     | Default | Notes                                                        |
|----------------------------------------------------------|---------|--------------------------------------------------------------|
| A session in the SourceLace app                          | 7 days  |                                                              |
| An AI app's access token (MCP, standard OAuth with PKCE) | 1 hour  | Bound to SourceLace's MCP address, so it works nowhere else. |

| What                         | Default  | Notes                                                                                                          |
|------------------------------|----------|----------------------------------------------------------------------------------------------------------------|
| An AI app's refresh token    | 30 days  | Replaced each time it is used; an AI app not used for 30 days signs in again. Revoking either token ends both. |
| A sign-in on a runner's page | 12 hours |                                                                                                                |

Admins can make the first three shorter on the [Limits](#) page. A removed person is refused at their next request, without waiting for a token to expire.

## Can SourceLace change records in our systems?

Only where an admin has turned changes on for named objects of a source. Every change is previewed and sent only after the person confirms it, and previews expire within 10 minutes. Queries are checked to be read-only before they are sent. SourceLace never sends email: mail sources only make drafts. See [Reads, changes and query checks](#).

## Can SourceLace staff see our data?

SourceLace's operators manage your account (plan, admins, contract, single sign-on) on an operator page. Every change they make is recorded in your own audit trail under their name. That page and its export never show data from your systems, sign-ins to them, secrets or AI keys. A written policy on which staff can reach the hosting environment is not yet formalised.

## Data stored, retention and deletion

### What does SourceLace store, and what does it not?

SourceLace does not copy your records. Query results are held for at most 30 minutes (up to 2,000 rows a query), in server memory or encrypted in a short-term store, and never written to a database or a log. Files are read on demand and never copied or indexed. The longer-lived exceptions are saved chats in the SourceLace app and AI agents' run summaries and pending changes, both encrypted with your organization's key. The full list is in [What SourceLace stores](#).

### Can we make retention shorter?

Yes. On the **Data retention** page an admin can shorten how long results are held (1 to 30 minutes), keep fewer rows, set stricter limits for one source, keep results in server memory only, and delete a chat's results as soon as it is closed (on by default). Each result held and each deletion is recorded in your audit trail, with no row values. See [Data retention controls](#).

### How long are chats, agent runs and the audit trail kept?

Saved chats and agent runs: 30 days after last use by default, set by your admin from 1 day up to your plan's longest. The audit trail: your plan's audit history, 90 days to 7 years. See [How long things are kept](#).

## What happens to our data when we leave?

Ask [support@sourcelace.com](mailto:support@sourcelace.com) to close your account. Nobody can sign in from that moment, and 7 days later everything SourceLace stores for your organization is removed for good, including the audit trail and your organization's encryption key. During those 7 days the deletion can be undone. If you want your settings and audit trail first, ask for an export before the deletion. See [Removing access and data](#).

## Does SourceLace learn from our people's questions? What can our admins see?

Only if your admin turns on **Learn from use** (Preview, off by default). SourceLace then looks at the questions people ask in the SourceLace app, inside your organization only, for questions several people repeat, and suggests a skill or agent for your admin to review. Admins see the words those questions share, how many people asked, how many questions and on which days; they never see who asked, and never the questions themselves unless they also turn on example questions, which shows up to three per suggestion with personal details such as email addresses and phone numbers masked. A pattern is shown only when at least 3 different people asked it. People are counted with a one-way scrambled code that is never stored. Suggestions are encrypted with your organization's key, are deleted when the chats they came from are deleted, and are never shared with other organizations, sent to SourceLace or used to train AI models. See [Learn from use](#).

## Can a person ask for their own data to be deleted or exported?

Yes. People can delete their own chats and disconnect any source at any time. For other requests about personal information, email [privacy@sourcelace.com](mailto:privacy@sourcelace.com).

## AI and model providers

### Does SourceLace send our data to an AI model?

- **From your own AI app** (Claude, ChatGPT, Cursor and others): no. SourceLace returns results to your AI app, which sends them to its own provider under your agreement with it.
- **From the assistant in the SourceLace app, and from AI agents:** yes. The conversation (or the agent's instructions) and the results of the tools it used, which can include rows from your systems, go to the AI provider your admin chose.

See [AI and your data](#).

### Which providers, and can we use our own account?

By default, Anthropic, through SourceLace's own Anthropic account. Your admin can instead enter your organization's own Anthropic or OpenAI key, so requests go under your own agreement with that provider and are billed to you. See [AI settings](#).

## Do the providers keep or train on our data?

SourceLace does not train AI models on your data and sends it to no AI provider other than the one your admin chose. Under Anthropic's commercial terms, Anthropic does not train on data sent through its API. Requests to OpenAI are sent with its store option off, though OpenAI's own policies may keep data for a limited time, for example for abuse monitoring. How long each provider keeps requests is set by that provider's terms; a zero-data-retention arrangement on SourceLace's own account is not in place.

## Can we stop some fields from reaching the AI?

Yes, on every plan. On the **Data protection** page your admin can hide a field so it never reaches the AI, or mask values (email addresses, phone numbers, US Social Security numbers, payment card numbers, IBANs and your own patterns) so the AI gets a placeholder. The rules apply to your own AI apps, the assistant, AI agents and runners alike, and people you allow still see the real values in the SourceLace app. Everything is off until you turn it on. The rules do not cover what can be inferred from filters, values typed into a question, or text inside images and scans; see [What is not covered](#) and [Data protection](#).

## Subprocessors

### Who processes our data?

| Company                                           | What for                                                                                         |
|---------------------------------------------------|--------------------------------------------------------------------------------------------------|
| Render                                            | Hosting the application, its database and the short-term store for query results (United States) |
| Anthropic                                         | The assistant in the SourceLace app and AI agents, when people use them                          |
| OpenAI                                            | The same, only if your admin enters your own OpenAI key                                          |
| Google, Microsoft or your single sign-on provider | Signing people in to SourceLace                                                                  |
| Stripe                                            | Billing; receives your billing contact and company name, never data from your systems            |

For SourceLace's own operations, Google Workspace carries SourceLace's company email (including messages you send us) and Cloudflare provides DNS for SourceLace's domains; neither receives data from your systems. The systems you connect receive only the calls each person makes with their own login. What each subprocessor receives is on [Security and privacy](#).

### Will you tell us before adding one?

The subprocessor list is updated before a new provider handles customer data. A contractual notice period is not yet formalised.

## Logging and audit trail

### What is logged?

Every tool call, whether it succeeded, was refused or failed, from the SourceLace app, every AI app and admins; every sign-in, including refused ones; every change to people, sources, groups, single sign-on and settings; every AI agent run, step and approval; and each query result held and deleted. Each entry holds when, who, the tool, the source, the query text, the object and record id, the names of changed fields, how many rows came back, the outcome, the error type and how long it took. See [The audit trail](#).

### What is never logged?

Row values, field values, passwords, keys and tokens. The query text is kept as written, so a value typed into a query (such as a customer name like "Corvanta") appears in the trail. SourceLace's server logs never contain tokens, passwords or query results.

### Is the audit trail tamper-evident?

Yes. Entries are chained by SHA-256 hashes per organization, so an entry that was edited, removed or reordered breaks the chain and can be detected. If an entry cannot be written, the call fails rather than returning data that was not recorded. A button for admins to run the integrity check themselves is not yet in the app.

### Can we see and export it?

Admins read, filter and search the trail on the **Audit Log** page (Business and Enterprise; on every plan every call is still recorded). A self-service download is not yet available: ask [support@sourcelace.com](mailto:support@sourcelace.com) and SourceLace exports your settings and audit trail as a JSON file. Streaming the trail to your SIEM is not yet available.

## Vulnerability management and secure development

### How is SourceLace built and tested?

Every change runs through automated checks before it can be merged: linting, formatting, strict static type checking and the full test suite, which includes tests for security behaviour such as read-only query checks, access rules and that secrets never appear in results. Dependencies are pinned in lockfiles, and builds install exactly what the lockfile lists.

### Do you do penetration tests or run a bug bounty?

Not yet. No independent penetration test report is available, and there is no bug bounty programme. Automated scanning of dependencies for known vulnerabilities is not yet part of the checks either.

### How do we report a vulnerability?

Email [security@sourcelace.com](mailto:security@sourcelace.com).

## How are security fixes delivered?

On SourceLace's cloud, fixes are deployed by SourceLace; you do nothing. On a self-hosted install or a runner, your IT team takes new versions; the **Runners** page shows a runner that is out of date. A formal patching timeline (for example, by severity) is not yet published.

## Incident response and breach notification

### Do you have an incident response plan?

A written incident response plan is not yet formalised. To report a suspected incident, email [security@sourcelace.com](mailto:security@sourcelace.com); for anything else urgent, [support@sourcelace.com](mailto:support@sourcelace.com).

### Will you tell us about a breach?

Yes. If a breach affects your information, SourceLace will notify you and your organization as the law requires. A contractual notification time (such as 72 hours) is not yet formalised.

### Can you stop access quickly in an incident?

Your admins can, at once: remove a person, remove a source, turn changes off, pause or stop AI agents, and revoke a runner (a run in progress on it stops at its next step). People can disconnect any source themselves.

## Business continuity and backups

### How is the database backed up? What are your RTO and RPO?

Not yet documented. Backup frequency, how long backups are kept, restore testing, and recovery time and recovery point objectives for SourceLace's cloud are not yet defined in writing. Two things are true by design: query results are never in the database, so they are never in a backup; and when an organization is deleted its encryption key is deleted too, so anything encrypted for it that remains in a backup can no longer be read.

### Is there an uptime SLA or a status page?

Not yet. No uptime commitment or public status page is published. The app's **Status** page shows, for your organization, whether the assistant is ready and how each source's calls went in the last 24 hours.

### What if SourceLace is unavailable?

Your data stays in your own systems, which SourceLace only reads and changes on request; nothing you rely on lives only in SourceLace except saved chats, skills, projects and agent settings.

## Healthcare and HIPAA

SourceLace has **no HIPAA Business Associate Agreement programme yet** and holds no HIPAA attestation. **Do not use SourceLace with patient data or other protected health information until a BAA is available** and signed with your organization. Ask [security@sourcelace.com](mailto:security@sourcelace.com) to be told when it is.

# Self-hosted installs

## What does a self-hosted install send to SourceLace?

Counts only, once a day over HTTPS, with an online license: SourceLace's version, your plan and seats, how many people have access and are active, how many sources of each kind, how people sign in, and how many questions, tool calls, agent runs and errors (by type). Never records, names or email addresses, queries, chat text, file names, record ids, object or field names, error messages, keys or passwords. The report is signed so SourceLace can tell it is complete, and your admins see its exact contents on the **License** page. With an offline license nothing is sent; an admin emails a monthly file instead, after reading it if they like. See [What SourceLace receives](#).

## Does the license need internet access?

No. The license is a signed file your install checks on its own. When it lapses, nobody is cut off mid-work; see [When a license lapses](#).

## What is in a support bundle?

Only what an admin chooses to download and send: the version, health checks, the license's state, settings by name with secrets shown only as "set" or "not set", 7 days of usage counts and recent log lines with email addresses, tokens and passwords removed. No records, queries, chat text or keys. See [the support bundle](#).

## Who is responsible for what when we host it?

You are: the infrastructure, the database and its backups and encryption at rest, HTTPS, keeping the master key safe, sign-in app registrations, taking new versions, monitoring and logs. SourceLace provides the software, the license and support.

## And a runner?

A runner keeps your AI key, each person's source sign-ins and every run's results on your network. It connects out to SourceLace over HTTPS (you open no inbound port) and sends only each run's status, times, one line per step (tool, source, object, row count, outcome) and health counts; SourceLace refuses any report that tries to include a summary, rows, values, drafts or files. See [Runners](#).