

# Data protection

From the SourceLace manual: Choose what the AI may see of the values SourceLace reads from your systems: hide fields, mask values such as email addresses and card numbers, and still let the right people see and use the real values. Admins set this on the Data protection page.

SourceLace version 0.1.0 · Made October 6, 2026

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

# Data protection

Choose what the AI may see of the values SourceLace reads from your systems: hide fields, mask values such as email addresses and card numbers, and still let the right people see and use the real values. Admins set this on the **Data protection** page.

## How it works

Your rules apply to everything that reaches an AI through SourceLace: your own AI app (over MCP), the assistant in the SourceLace app, and AI agents. They add to what SourceLace already does: each person uses their own login, access by group narrows what they reach, results are never copied to a database, and every call is audited.

- **Hidden** fields never reach the AI. They are left out of every result, and a query that names one is refused.
- **Masked** values reach the AI as placeholders, such as [EMAIL\_K3F9QX]. The same value gets the same placeholder throughout one chat, agent or AI app, so the AI can still say "the same person" and work with it. Another chat gets different placeholders.
- **The person who asked** sees the real values in the SourceLace app, and when the AI proposes a change, a draft or a message with a placeholder, SourceLace puts the real value back before the preview they confirm and before anything is written. The AI never sees it.

Everything is off until an admin turns it on.

## Field rules

A rule names a source, an object (or \* for every object) and a field, and says what to do with it:

| Rule         | What the AI gets                                                                                            |
|--------------|-------------------------------------------------------------------------------------------------------------|
| <b>Hide</b>  | Nothing: the field is removed from every result, and queries that select, filter or sort on it are refused. |
| <b>Mask</b>  | A placeholder instead of the value, such as [FIELD_Salary_7QW2NA].                                          |
| <b>Allow</b> | The value as it is, even when a detector would match it (for example a shared support address).             |

Names are not case-sensitive. A rule can list **groups whose members see the real value**. For them a hidden field is masked instead of removed, and a masked value is shown to them in SourceLace and can be used in changes. Without groups, everyone who can read the record sees the real value in SourceLace. The AI only ever gets the placeholder.

## Detectors and your own patterns

Detectors find values inside any text, such as notes, email bodies, chat messages and file text, and mask them. Turn on the ones you need:

- **Email addresses**
- **Phone numbers** (written with separators or a leading +, such as +44 20 7946 0958 or (415) 555-0134)
- **US Social Security numbers** (such as 123-45-6789)

- **Payment card numbers** (only numbers that pass the card checksum)
- **IBANs** (only numbers that pass the IBAN checksum)

Add **your own patterns** for identifiers such as employee numbers, as regular expressions (for example EMP-\d{6}). SourceLace refuses patterns that are invalid, match empty text, or could run for too long. If a pattern ever takes too long on a value, SourceLace masks the whole value rather than let it through.

**Test a sample** on the page: paste one row as JSON to see exactly what the AI would get. The sample is not stored.

## Real values for the right people

- **In the SourceLace app** answers, tables, charts and change cards show real values to the person who asked, if they may see them. Saved chats keep the placeholders together with what they stand for, encrypted with the chat and deleted with it, so you see real values when you reopen a chat.
- **Changes, email drafts and chat messages** proposed with placeholders get the real values server-side, only for people who may see them. A placeholder that is unknown, has expired, or belongs to a value the person may not see is refused and nothing is changed.
- **In your AI app** (over MCP) the AI and the previews it shows you have placeholders only. Open the SourceLace app to see real values.
- **AI agents** work with placeholders too; their drafts and changes get the real values, and the agent's results chat shows them to its owner.

What placeholders stand for is held like a query result, for your organization's result time limit (see [Data retention controls](#)), then deleted. It is never logged or written to the audit trail.

## Queries

SourceLace refuses a query that names a hidden field, or that renames, wraps or computes on a masked one (for example salary AS pay or UPPER(email)), for SOQL, SQL, WQL, OData, Oracle Fusion and the JSON query formats. For the other query languages (Zendesk search, JQL, Marketing Cloud, Workday REST and custom connectors), results are still protected, but a filter on a hidden field is not refused, so it can still narrow which rows come back.

## What is not covered

- What can be **inferred**: filtering or sorting on a masked field can reveal something about its values. Hide fields that must not be reasoned about at all.
- Values the AI **already knows** or can work out from other fields, and anything a person types into a question.
- Data inside **images, scans and attachments** SourceLace does not read as text, and values written in formats the detectors do not recognise.
- **Field names and record ids** are never masked.
- **Documents** the assistant makes (Word, PowerPoint, PDF) show placeholders.

## Audit

Every rule change is in the audit trail. Each call's entry records how many fields were hidden and values masked, never the values themselves.

A SourceLace Runner applies the same rules.