

Asking across systems (Preview)

From the SourceLace manual: Ask one question that needs several systems at once, such as "which accounts with open renewals have overdue SAP invoices?", and get an exact answer: every row is counted, nothing is estimated, and you can see where each number came from.

SourceLace version 0.1.0 · Made October 7, 2026

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

Asking across systems (Preview)

Ask one question that needs several systems at once, such as "which accounts with open renewals have overdue SAP invoices?", and get an exact answer: every row is counted, nothing is estimated, and you can see where each number came from.

Status: Preview: new, for pilots, provided as is. Tell support@sourcelace.com what you ask with it and how the answers look.

What it does

When a question needs more than one system, or more rows than fit in a chat, the AI writes the question once. SourceLace then:

- **Asks each system only for what the question needs.** Each system gets its own filters and fields (for example only open opportunities, or only invoices past due), with your own sign-in, so it returns only rows you are allowed to see.
- **Asks every system at the same time.** Three systems take about as long as the slowest one, not the three one after another.
- **Puts the rows together and works out the answer exactly.** The rows are joined, counted and totalled in a short-lived, encrypted place on SourceLace's server that only your question can read. The AI quotes the numbers it computes; it does not estimate them.
- **Deletes the rows when the result expires,** at the same time as any other result (30 minutes unless your admin chose less), or as soon as you close the chat if your admin set that.

A question to one system works the same way when the AI asks on purpose for more rows than a chat result holds (2,000): SourceLace then keeps every row, up to 250,000 (the number fetched without asking first), so totals and counts cover all of them. Otherwise the result is cut at 2,000 rows and marked incomplete, as before.

What you see

Under the answer's table, **Where these rows came from** lists each system, the object read, how many rows came back and how long it took. Admins also see exactly what SourceLace sent to each system.

The answer says **Complete** when every row was read. If a system had more rows than your organization allows for one question, SourceLace stops and tells you how many each system has, so you can narrow the question. Only when you ask for an answer from the first rows anyway is it marked **Incomplete**, and the AI says so.

You can keep asking about the same answer ("now by region", "which of these are in Germany?") without SourceLace fetching the rows again.

Large questions ask first

If a question would fetch a lot of rows (more than 250,000 unless your organization set another number), SourceLace shows the plan before anything is fetched: each system, about how many rows it would return, and what it costs in credits. In the SourceLace app, click **Run this plan** to go ahead; in your AI app, tell it to go ahead. Or narrow the question instead. A plan can be run for 10 minutes.

What protects your data

- **Your own access.** Every system is read with your own sign-in, and your organization's groups decide which objects you may ask about. They are checked again each time you ask more about an answer.
- **Data protection first.** Fields your organization masks are replaced before the rows are put together, so the real values never reach the place where the answer is worked out. SourceLace can still match and count masked values (for example the same customer in two systems), but refuses questions that would need the real values, such as sorting by a masked field, and says why. The same goes for text in which data protection found and replaced personal data, such as an email address in a notes field. Masked values look different in each conversation, so an answer made in one chat cannot be combined with rows in another: ask again in the chat you are in.
- **Not kept.** The rows are encrypted with a key of their own (or, on a server without encryption set up, held only in its memory, never on disk) and deleted when the result expires. Nothing is copied into a database or cache.
- **Audited.** Your organization's audit trail records the question and, for each system, what was sent and how many rows came back, never the rows themselves.

Limits

How many systems one question may use, how many rows each may return, how long a question may take and how many rows your organization may fetch each month are on the [Limits](#) page.