



SECURITY ARCHITECTURE

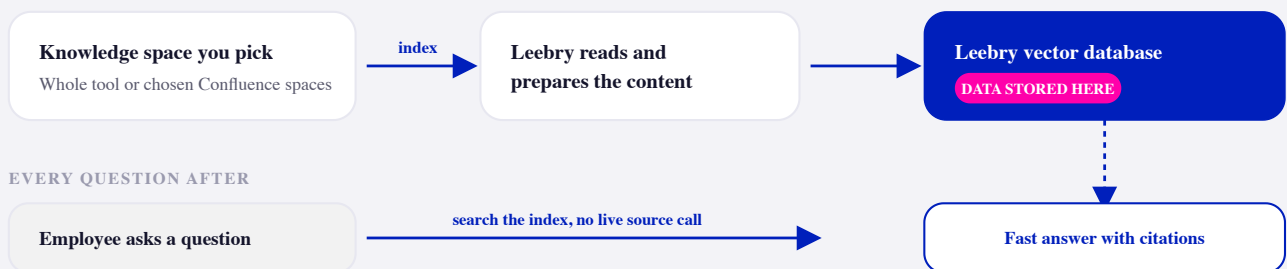
How Leebry Works With Data

Our architectural choices, access modes, and data storage information explained.

Organization-level access (indexed)

You point Leebry at a knowledge space, either a whole tool like Confluence or specific spaces you name. Leebry reads that content, prepares it, and stores it in its vector database. When someone asks a question, Leebry searches that index and answers from it. Because the content is already prepared and searchable, answers come back fast and the quality is high. Use this model for company-wide information that holds no secrets.

SET UP ONCE, INDEXING



Indexed content is prepared once and reused for every question, which is why this model is the fastest.

BEST FOR

Company-wide knowledge that carries no secrets: docs, runbooks, onboarding, policies.

WHAT YOU GAIN

The fastest responses and the highest answer quality, because nothing is fetched mid-conversation.


THE TRADE

A prepared copy lives in Leebry's vector database, so scope it to spaces everyone may read.

Organization-level access (real-time)

This model keeps the broad reach of organization-level access but skips the index. You open up the public or company-wide spaces you want available, and Leeby searches them live at the moment of the question instead of preparing a copy ahead of time. Nothing is indexed and nothing is stored in Leeby's database. Use it when you want public docs answerable through Leeby but would rather not keep a copy of them.

EVERY QUESTION, IN REAL TIME



Leeby reaches the open spaces live and answers from them, without preparing or keeping a copy.

BEST FOR

Public or company-wide docs you want answerable through Leeby without keeping a copy of them.

WHAT YOU GAIN

The broad reach of organization-level access with no stored copy, and answers that always reflect the live version.

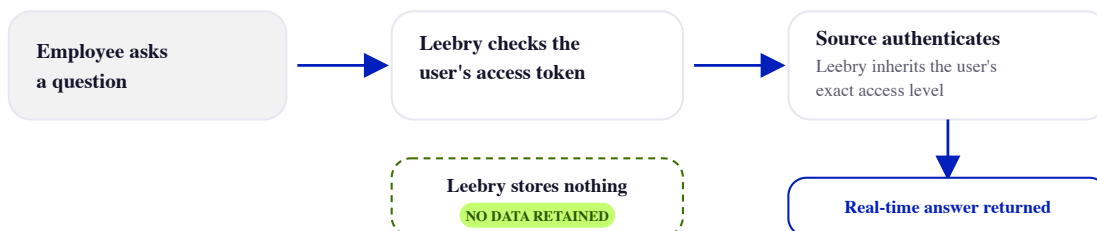
THE TRADE

Real-time lookups are slower than the indexed model, and this covers shared spaces rather than per-user permissioned data.

User-level access (token-based)

Here Leebry indexes nothing. When a user asks a question, Leebry first checks whether that person can reach the source through their own access token. After the user authenticates, Leebry takes on exactly their permission level and can see only what they could see. It queries the source live, returns the answer, and keeps no copy. Because Leebry holds no data of its own, there is nothing to pull from later, by accident or by a deliberate attempt. A Slack request can never surface a salary to someone who was never allowed to open that record.

EVERY QUESTION, IN REAL TIME



Nothing is indexed and nothing is kept, so a query can only ever return what the asking user already has rights to.

BEST FOR

Sensitive, permission-bound information where the wrong viewer is the risk you are guarding against.

WHAT YOU GAIN

No stored copy means nothing to leak, and answers stay scoped to each person's real permissions.

THE TRADE

Each answer waits on a live lookup, so responses run a little slower than the indexed model.

The three models compared

	Org-level, indexed	Org-level, real-time	User-level, token-based
Indexing	Indexed and refreshed	Not indexed, live search	Not indexed, live search
Where data lives	Stored in Leeby's vector DB	Nothing stored	Nothing stored
Access scope	The spaces you index	Public spaces you open up	Only what the asking user can see
Speed	Fastest	Slower, real-time	Slower, real-time
Freshness	As fresh as the last refresh	Always the live version	Always the live version
Use it for	Company-wide knowledge, top speed	Public docs you'd rather not store	Sensitive, permission-bound data

The rule we give customers: **index company-wide knowledge when you want top speed, use real-time organization-level access when you want public docs answerable without a copy in our database, and use user-level access for sensitive, permission-bound data.** Most organizations mix two models: a user-level and one of the org-level.

How your data is handled, end to end

The access model decides what Leebry can reach. These rules cover what happens to your data once it is there, and they hold no matter which model a source uses.

- **Your data never trains a model**

Leebry retrieves, it does not train. Your content is referenced to answer a specific question. It is never used to train or fine-tune a model, never added to a shared knowledge base, and never passed to a third-party provider for training.

- **Encryption standards**

All access data is encrypted using industry-standard protocols, both in transit while being moved and at rest while stored.

- **Stored as vectors, not raw text**

Indexed content is held as mathematical vectors, a numeric form built for retrieval that keeps a step of abstraction away from the original text.

- **Restricted, audited internal access**

MacPaw works under practices aligned with **ISO 27001** and **SOC 2 Type II**. There is no casual team access to your content. Infrastructure access is limited to a few authorized engineers for maintenance, and every instance is logged and monitored.

- **Deletion you control**

Disconnect an integration or close your account and the associated data is flagged for permanent removal. It drops out of results right away, with physical deletion following on the next cleanup cycle.