

CONVERSATIONAL DATA EXPLORATION

AI LLMQuery

Ask your database questions in plain English and get answers you can verify. Every query is generated from your own schema, runs read-only against your data, and stays inside your network — no value from your database is ever sent to the AI.

Read-only by design

No account required

Zero data stored by us

Bring your own Claude or GPT key

WHAT YOUR PEOPLE WOULD ASK ON MONDAY

The questions nobody gets around to asking

Not because they don't matter — because asking means a ticket, a queue, and an answer that arrives after the decision. These are real questions, answered from a real schema.

Give me the latest uploaded note for FC-76409.

Which department is creating a bottleneck in the timeline, and what activity is causing it?

Which servicer sends me the worst referral packages, and what does it cost me in days?

Which judges or counties are quietly adding the most time to my cases?

Which Judges have the best outcomes and why?

Do any note documents disagree with the interest rate in our loan system?

Which mortgages disclose a high environmental risk on the property?

Which documents should we re-scan?

Find me something in this data I am probably not looking at that is costing me money.

How many active files have no deed uploaded?

Are any of my vendors costing me more in delay than they save me in fees?

Where am I breaching SLAs that carry financial penalties, and what is the exposure?

Show me files that have a high-likelihood of turning into a long process and give me the breakdown by Servicer, State & Stage

Which loans carry a prepayment premium, and how much is it?

Which title reports show an unresolved curative item?

Do any notes disagree with our loan system?

Which mortgages disclose high environmental risk?

No Vendor Cloud

No account to create and no data stored by us. Your key is never seen, received, or transmitted by Alchemy Loop.

Read-Only, Always

Every connection opens read-only. One statement at a time — multi-statement requests are rejected outright.

Your Data Stays Put

No value from any table is ever transmitted. Not a name, a number or a date. This is how the software is built, not a setting.

Nothing's a Black Box

Every answer shows the query behind it and how it was measured — and a request trace can record exactly what was sent.

HOW IT WORKS

It works from your schema on day one

AI LLMQuery doesn't guess at your data — it reads the tables, views and relationships you already have, so it answers real questions with no modelling exercise first. Add plain-English column notes and business rules as you go, and every answer gets sharper.

1 Ask in plain English

Type a question the way you'd ask a colleague. AI LLMQuery translates it into a read-only query grounded in your own structure and definitions. Keep drilling down and interrogating.

3 Choose your level of intelligence

Plug in the AI provider you already use and switch between Claude or GPT-based models at any time. Your account, your spend, your choice of model.

2 Inspect and verify every answer

Every response states how it was measured and can show the query behind it. Open the results table, chart them, or export to CSV — and correct a measurement you disagree with.

4 Review before anything runs

AI-assisted view creation drafts CREATE VIEW scripts for your review — they are never executed automatically. A person always applies them separately.

● Query safeguards

- ✓ One statement only — multiple statements are rejected outright
- ✓ Must be a read query
- ✓ Write and structural operations are rejected
- ✓ Generated view scripts are never auto-executed — for review only
- ✓ Time limits on both the AI request and the database query
- ✓ Administration sits behind its own password, apart from everyday use

● What about PII?

What's sent with your question

Your question plus the relevant table, view, and column *names* — never data values. **And not always the question.** Before anything is sent, the question is checked against true PII patterns and against the values in the columns you have marked sensitive. When it matches, the question is refused outright — it never reaches the AI, there is no confirm-and-send-anyway, and the blocked attempt is still written to the request trace.

What is never sent

Any value from any table. There is no second request that returns your rows to the AI — the query runs on your machine and the results are displayed there.

Hide sensitive columns

Hide a column and it is removed from the schema the AI is shown, so it cannot write a query against it. Likely candidates are flagged for review, so an SSN column does not depend on somebody spotting it.

Governed answering - not an LLM wrapped around a database

LLMQuery combines advanced retrieval, query decomposition, contextual reasoning and data-layer techniques to produce precise, reliable and auditable answers.

- Runs on your premises — no external or third-party access
- One shared folder can serve every workstation
- Plug in the AI provider you already use — Claude or GPT
- OpenAI requests are sent with the `store:false` flag
- Switch providers or models at any time