
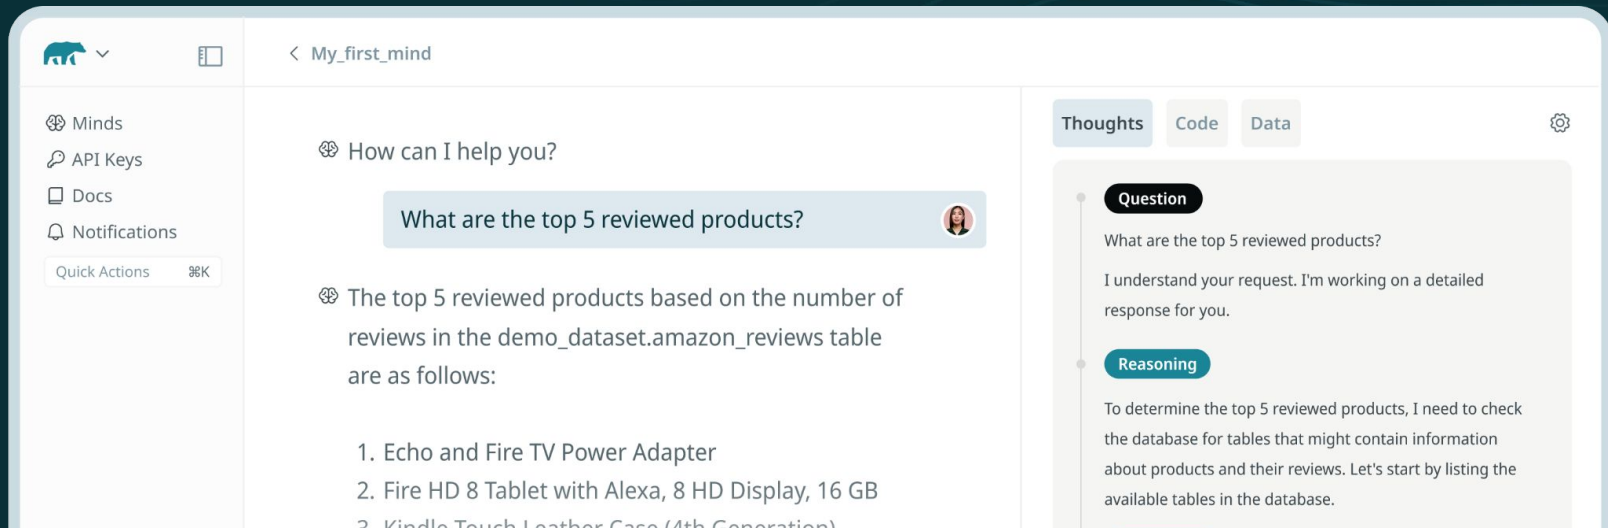


Connect. Unify. Respond.

Any data, anywhere with
human level intelligence.



Enterprise AI Analytics with Minds





Performance

Benchmarks

Easy:

What was the order volume for NVDA on January 2, 2024?


Medium:

What was the average order volume for all stocks on April 1, 2024?

Hard:

Which trading day in Q4 2024 had the highest trading volume and how does that compare to NVDA's performance in Q1-3 of that same year?

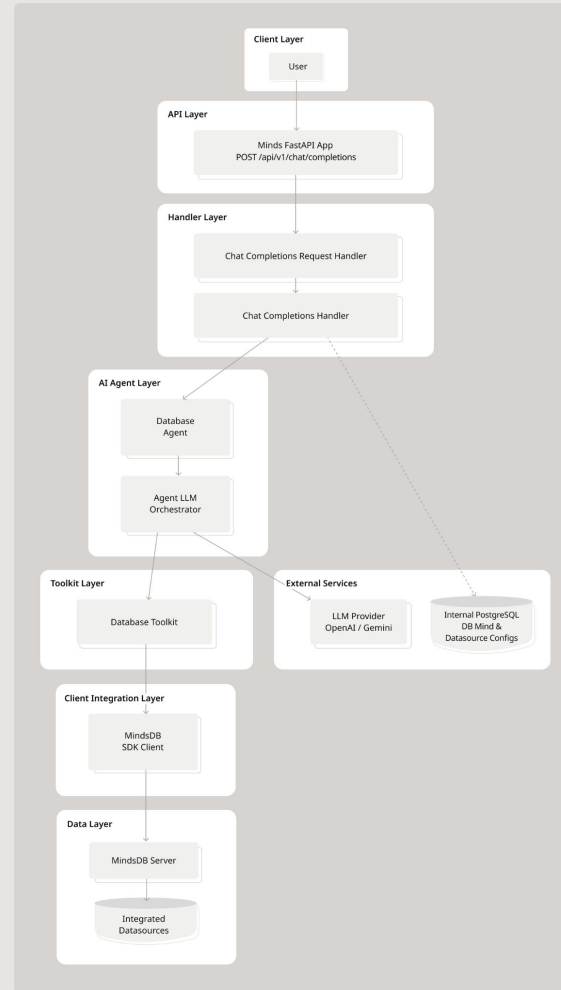
OOTB Response Accuracy across 109 questions



* Using GPT-4o,

** Was 88.1% before removing ambiguous questions

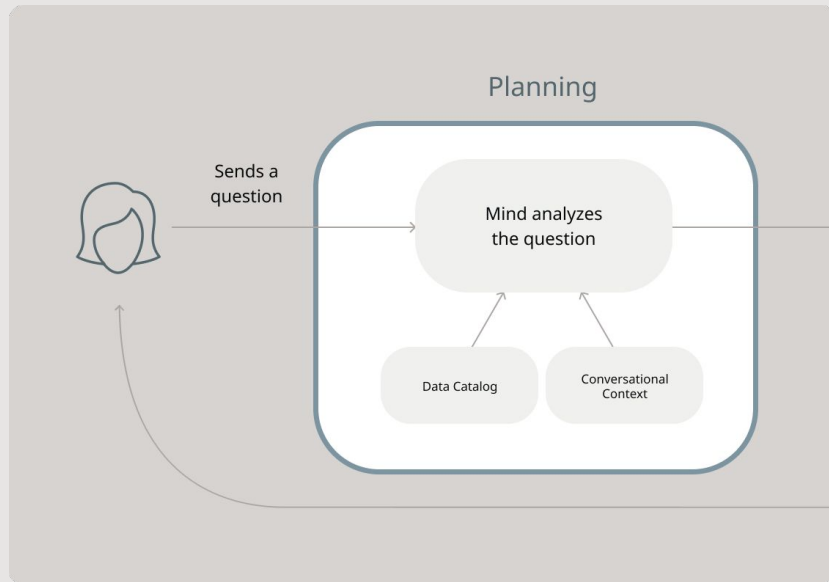
System Architecture

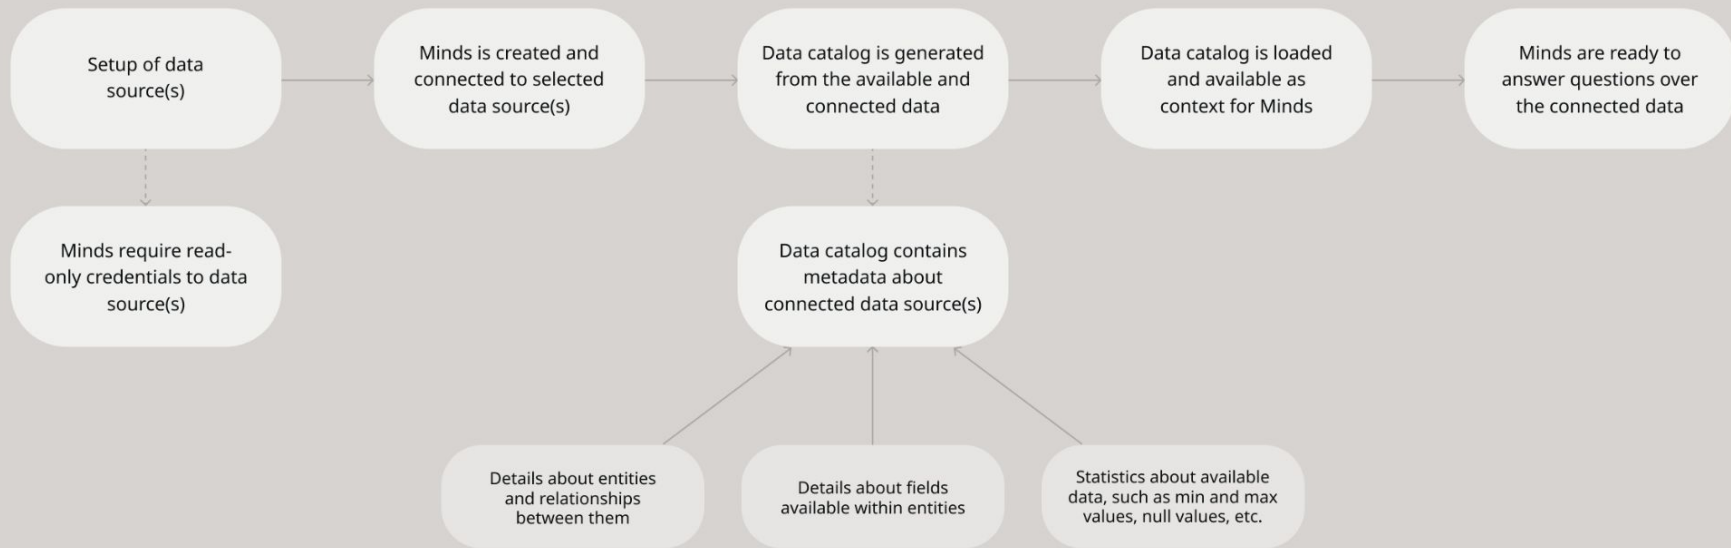




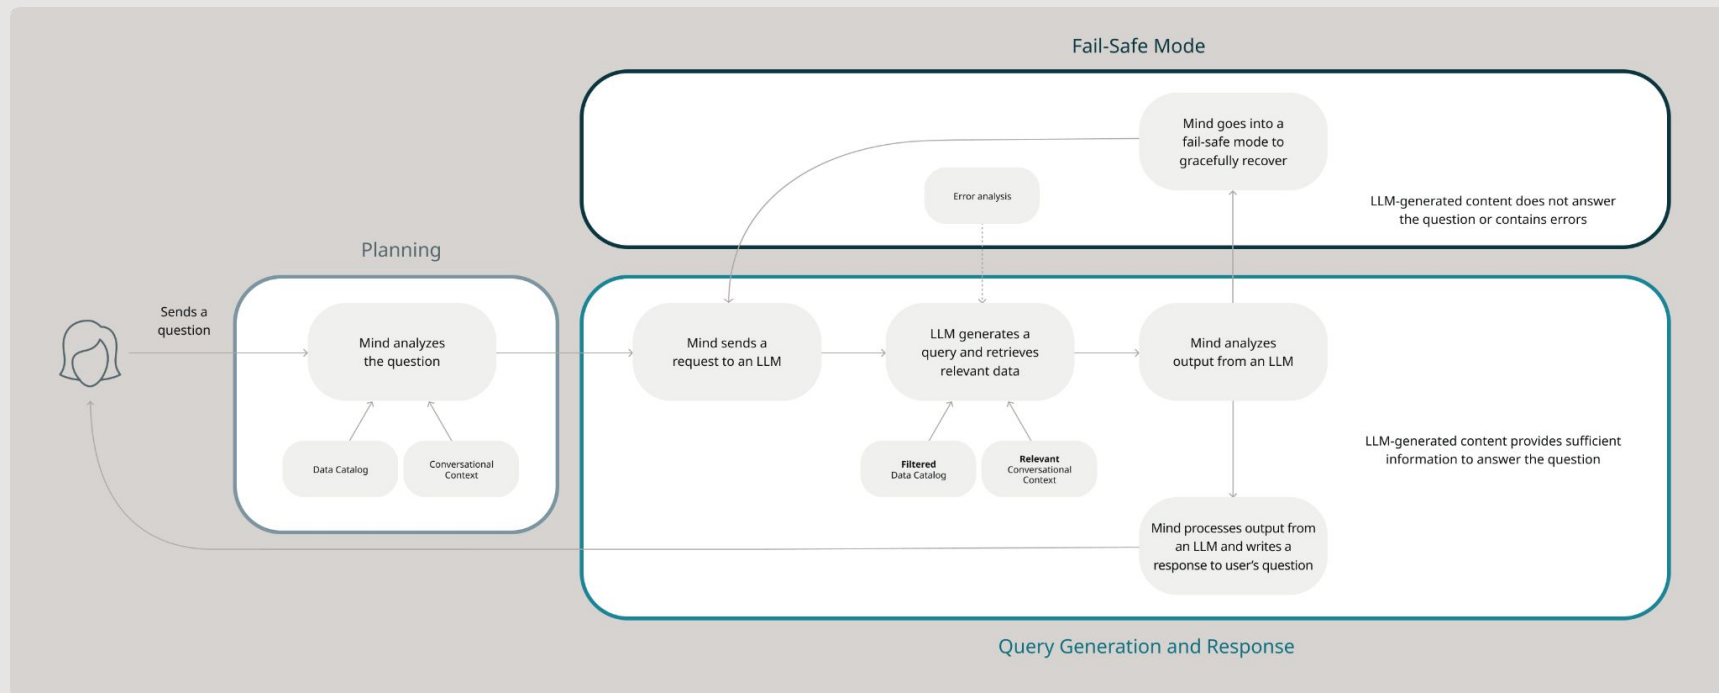
How our Mind Thinks

- Analysis of the user's question against the available data
- Mapping the user's question onto the available data
- Relevant data and context identified





How the Mind Thinks





How do we translate queries from natural language to executable code?

Natural Language → SQL Flow

1. Context Building – Extract user intent and context
2. Planning – Identify datasources/tables
3. SQL Generation – LLM creates query using templates
4. Validation – Sanitize and check SQL
5. Execution – Run query via MindsDB
6. Response – Stream results back to user

