


# Connect. Unify. Respond.

Any data, anywhere with  
human level intelligence.





Most widely-adopted AI  
query engine in the world



400K Deployments




30K GitHub Stars



200+ Enterprise Datasources

Zero-ETL Support for all major  
large-scale enterprise datasources



Industry-leading tech partners



Google



Microsoft

ORACLE



Backed by top investors



BENCHMARK

Mayfield

Y Combinator



**Unify Data**



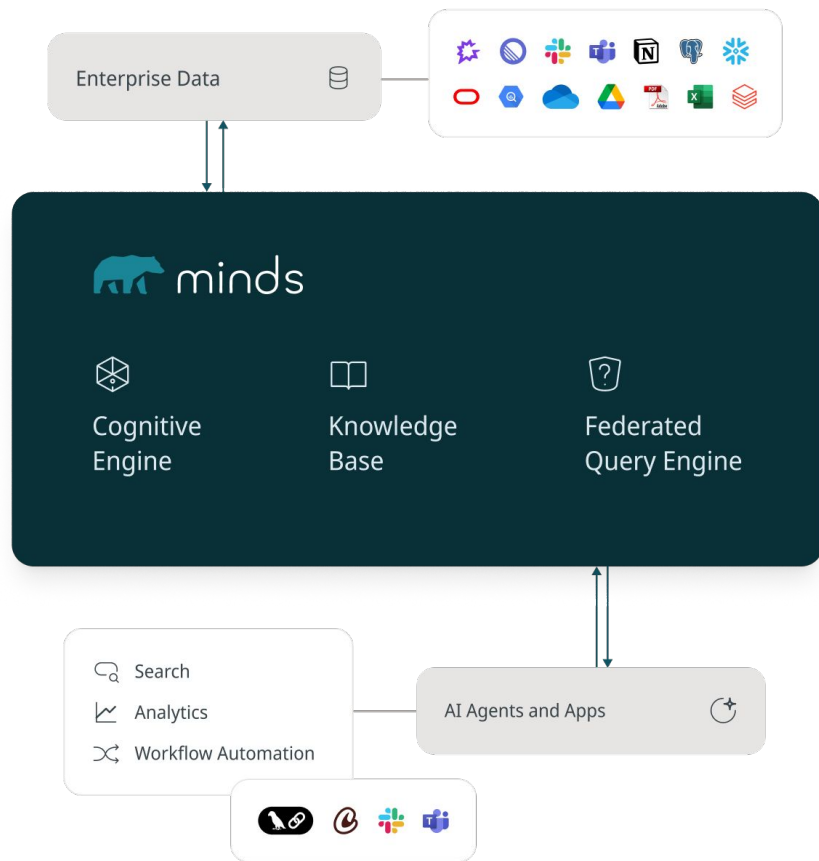

**Accelerate Insights**



**Empower Teams**



# MindsDB Simplifies and Derisks AI Implementations





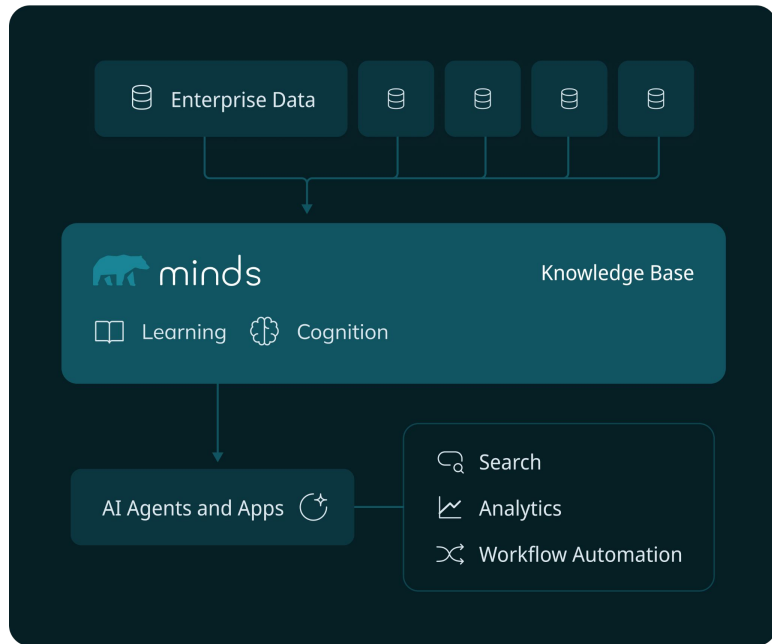
# Why MindsDB

Eliminate the complexity of fragmented enterprise data

## Minds

MindsDB empowers AI and humans to make data-driven decisions instantly.

- Query any data, anywhere
- Secure and scalable to any infra
- No data engineering required





# The MindsDB Stack

A Seamless Opportunity for Maximum AI Value

## Federated Query Engine (Connect)

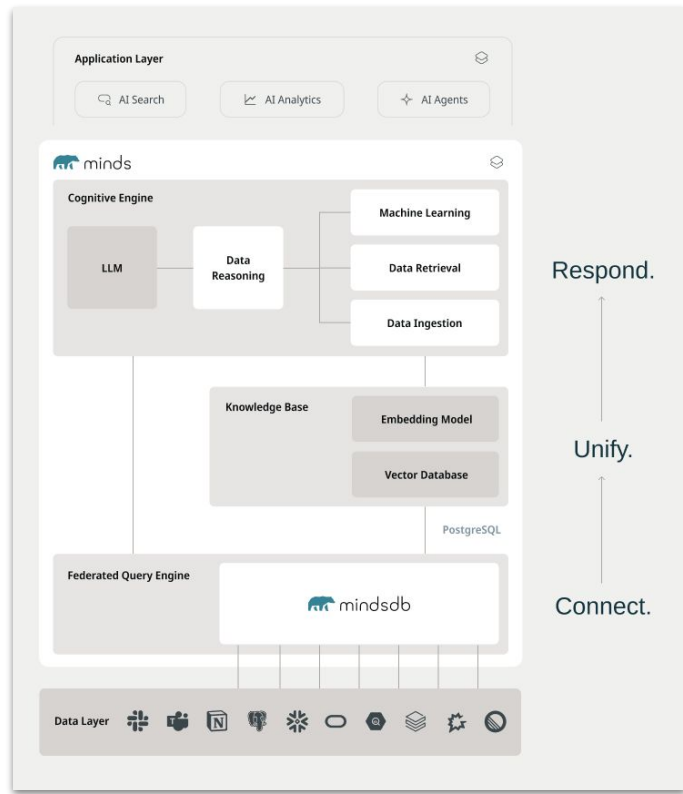
- Retrieves accurate, up-to-date knowledge
- Translates from SQL to any query language

## Knowledge Base (Unify)

- Makes sense of structured & unstructured data
- Avoids cost and complexity of moving data

## Cognitive Engine (Respond)

- Supports leading proprietary and OSS LLMs
- Orchestrates w/reasoning and planning to surface the best knowledge





# MindsDB OSS and Minds Enterprise

## Feature comparison

	MindsDB Open Source	Minds Enterprise
	Explore, prototype, build locally	<b>Scale securely &amp; reliably in production</b>
<b>Connect</b>	Any data, manual integration	Any data, <b>managed &amp; supported integrations*</b>
<b>Unify</b>	Knowledge Base w/manual updates	Knowledge Base w/ <b>dynamic updates</b>
<b>Respond</b>	BYO-LLM, basic UI	<b>Managed LLMs, biz rules &amp; personalized UI</b>
<b>Manage</b>	Basic setup, config, logs	<b>RBAC, Limits, Privacy, Observability</b>
<b>Deploy</b>	Single Mind, container, server	<b>Unlimited Minds, clusters, HA/failover</b>
<b>Scale</b>	Local / individual projects	<b>Unlimited scale</b>

\* Managed refers to Enterprise-level support, maintenance, SLAs, etc. It also includes potential optimizations we may add that improve how Minds work at scale or in multi-tenant deployments

# MindsDB Agents





# MindsDB Agents - Features

## Data-Driven Intelligence

Unlike regular LLMs, agents can access and analyze your actual connected data.

## Unified Access

Combines both structured data (tables) and unstructured data (knowledge bases) in one interface.

## Context-Aware

Understands your data context through custom prompt templates.



# MindsDB Agents

## Databases

- Databases tables from connected data sources (MySQL, PostgreSQL, MongoDB, etc.)
- Data warehouses (Snowflake, BigQuery, Redshift)
- Applications (Salesforce, HubSpot, Slack)

## Knowledge Bases

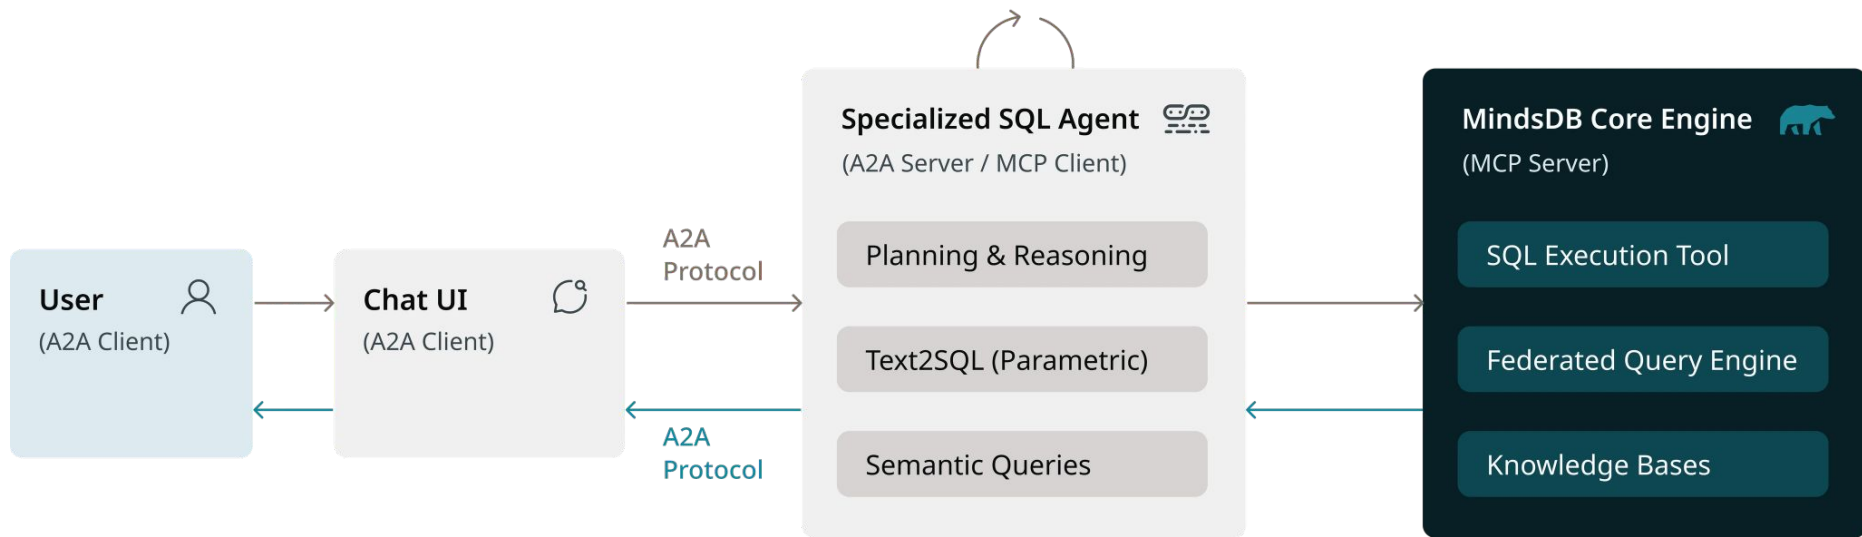
- Databases with unstructured data (comments, reviews, etc.)
- Web content crawled from websites
- Uploaded files (PDFs, documents, spreadsheets, etc.)

## A2A Ready

Follows the Agent2Agent Protocol standard for seamless agent communication.



# MindsDB Agents - Architecture



# Connect a Datasource

Example



# Connect a Data source

MindsDB lets you connect to your favorite databases, data warehouses, data lakes, etc., via the **CREATE DATABASE** command.

The MindsDB SQL API supports creating connections to integrations by passing the connection parameters specific per integration.

```
CREATE DATABASE postgresql_conn
WITH ENGINE = 'postgres',
PARAMETERS = {
  "host": "127.0.0.1",
  "port": 5432,
  "database": "postgres",
  "user": "postgres",
  "schema": "data",
  "password": "password"
};
```

# Create a Knowledge Base

Example



# Create a Knowledge Base

**embedding\_model1:** Specifies the model used to convert your text data into vector representations (embeddings). This allows for semantic searching (finding content based on meaning, not just keywords).

**reranking\_model1:** Defines an optional model to refine the search results. After an initial retrieval using embeddings, this model re-evaluates and re-orders the results for better relevance.

```
CREATE KNOWLEDGE_BASE my_kb
USING
    embedding_model = {
        "provider": "Google",
        "model_name" : "gemini-embedding-001",
        "api_key": "sk-..." -- optional, default from env variable
    },
    reranking_model = {
        "provider": "Google",
        "model_name": 'gemini-2.5-flash',
        "api_key": "..." -- optional, default from env variable
    },
    storage = my_vector_store.storage_table, -- optional, default ChromaDB
    metadata_columns = ['date', 'creator', ...], -- optional
    content_columns = ['review', 'content', ...], -- optional, default content
    id_column = 'id';
```

# Create an Agent

Example





# Create Agent

**Model:** which language model or AI engine to use. It includes parameters like **provider** (e.g., Google), **model\_name** (e.g., gemini, and authentication details like **api\_key**.

**Data:** specifies the data sources connected to your agent or knowledge base, including **tables** from databases and **knowledge\_bases** for semantic search capabilities.

**Prompt Template:** contains instructions that guide the AI model's behavior, including tone, response format, and how to interpret the connected data sources.

```
CREATE AGENT my_agent
USING
    model = {
        "provider": "google",
        "model_name" : "gemini-2.5 flash",
        "api_key": "....."
    },
    data = {
        "knowledge_bases": ["mindsdb.sales_kb"],
        "tables": ["postgres_conn.customers"]
    },
    prompt_template='
        mindsdb.sales_kb stores sales analytics data';
```

# Chat with your Data

Example



# Query an Agent

Search By keyword

To query an agent in MindsDB, use the **SELECT FROM AGENT** syntax:

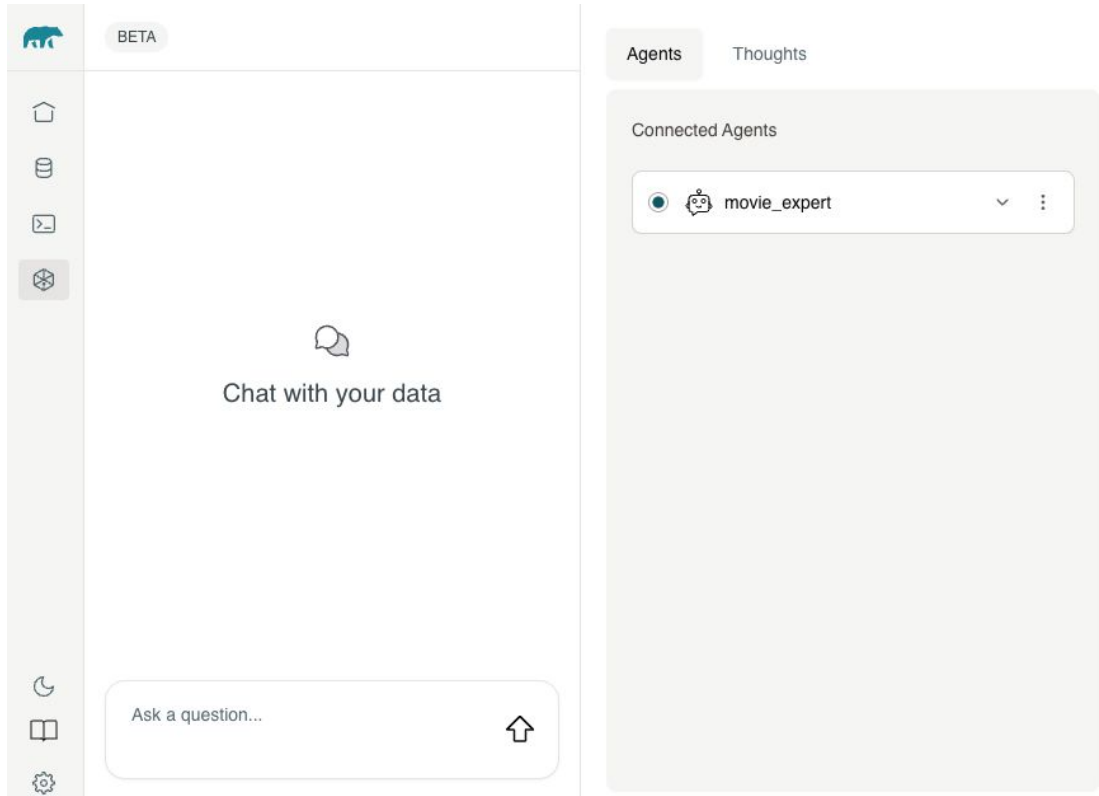
```
SELECT answer
FROM my_agent
WHERE question = 'What is the average number of orders per
customers?';
```



# Chat with Your Data

## Respond

Chat with your data directly on the  
'Respond tab' of the OS.



# A2A - Use cURL requests to talk with your Agents



# Calling A2A Agents via cURL

Once you've created an agent, you can interact with it remotely using a simple **cURL request** that follows the A2A format. This allows you to send tasks, messages, or queries directly to your agent over HTTP.

```
curl -X POST \
  "https://7625f341f0ff.ngrok-free.app/a2a" \
  -H "Content-Type: application/json" \
  -H "Accept: text/event-stream" \
  -H "Cache-Control: no-cache" \
  -H "Connection: keep-alive" \
  -d '{
    "jsonrpc": "2.0",
    "id": "your-request-id",
    "method": "tasks/sendSubscribe",
    "params": {
      "id": "your-task-id",
      "sessionId": "your-session-id",
      "message": {
        "role": "user",
        "parts": [
          {"type": "text", "text": "hi, tell me about a movie about talking dinosaurs"}
        ],
        "....."
```

# Thank You