



Building Agents with MindsDB

Today's Speaker



Jorge Torres
CEO
MindsDB

Agenda

1

From Generic to Custom Enterprise AI

2

Agents Reference Architecture


3

Demo

4

Q&A





FROM: Generic AI

Great off the shelf AI that doesn't know your data and solves generic problems.

```
response = AI( request )
```



TO: Custom-Enterprise-AI

AI capable of continuously learning from your data, tailored to solve your specific enterprise challenges.

```
response = AI( request + Your Data )
```

Developers love us, because they can build production AI systems in a few lines of code

Enterprises love MindsDB, because it turns their existing workforce into an AI army.

MindsDB ecosystem

One of the leading Open Source Applied AI projects

GitHub

21k

Stars

Docker

300k+

deployments

Open Code

750

Contributors

\$55M in Funding

Mayfield

BENCHMARK

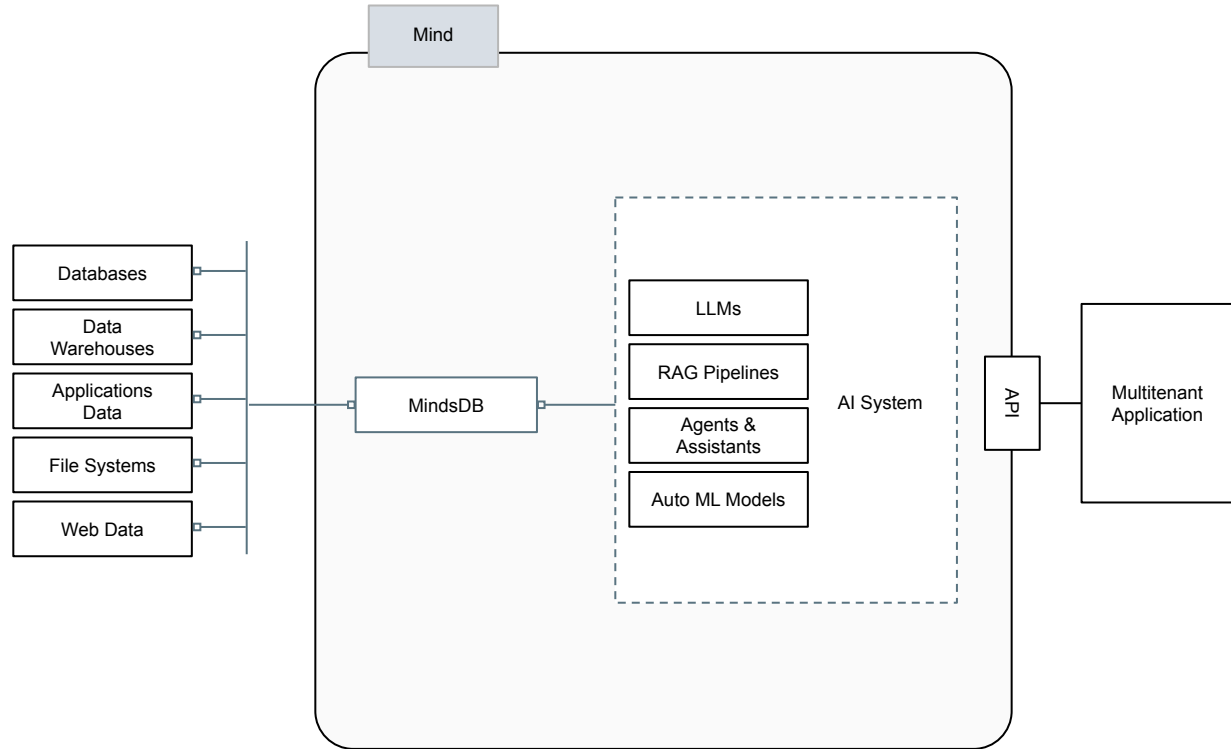


Industry Leading Tech Partners



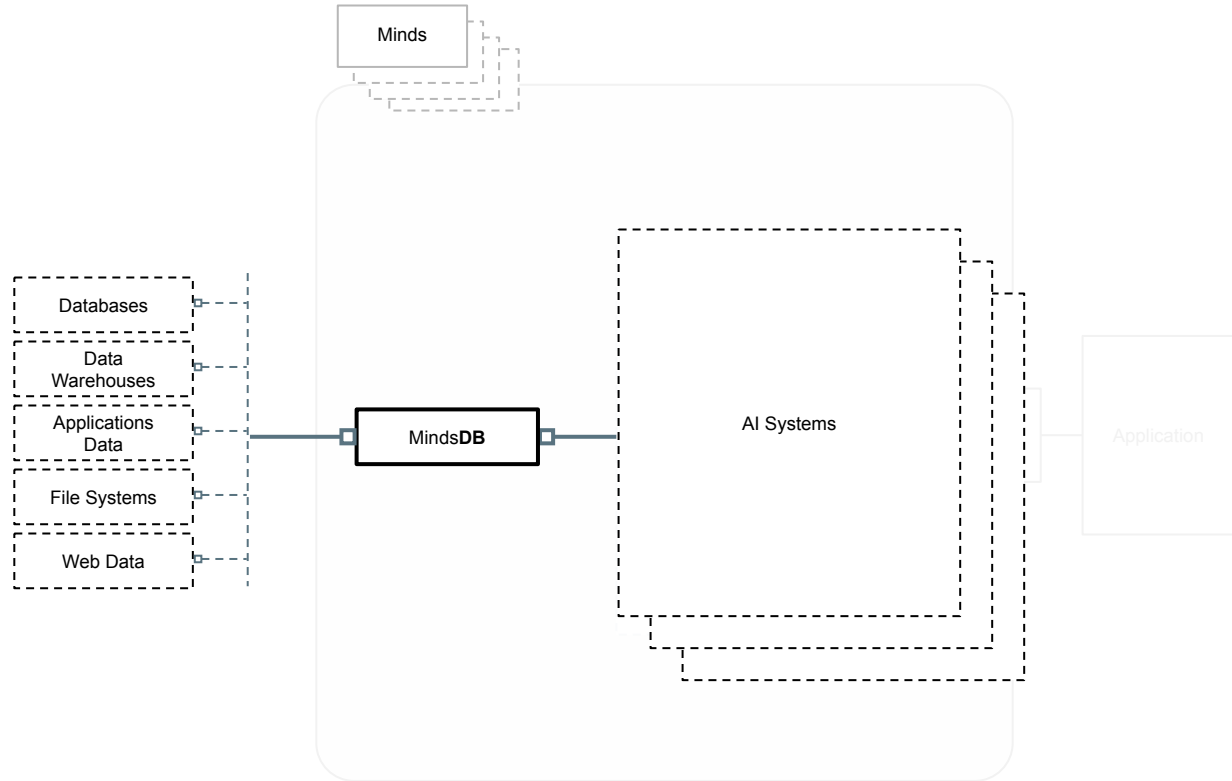
Minds

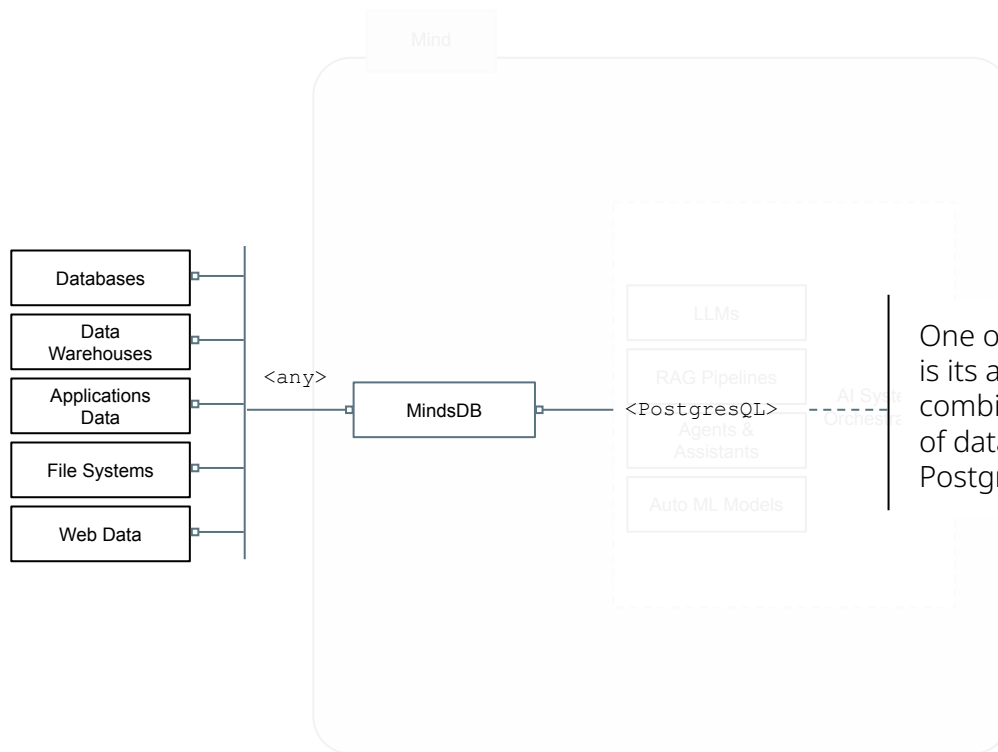
Minds are AI systems built for solving application challenges that require both **AI** and **data**.



MindsDB

Minds**DB** is the server that **orchestrates** the data workflows between **data-sources** and the Minds' **AI systems**.



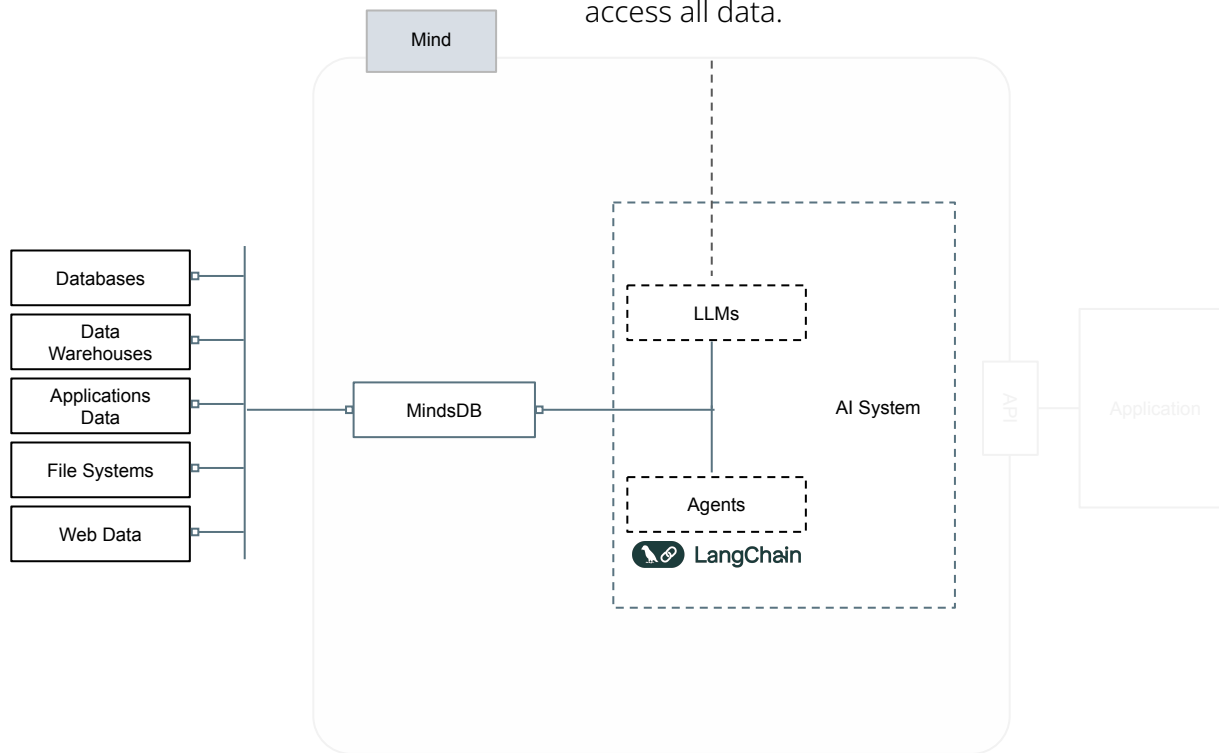


One of MindsDB's advantages is its ability to query and combine data from any kind of datasource using a unified PostgreSQL dialect.

Langchain <> MindsDB Integration

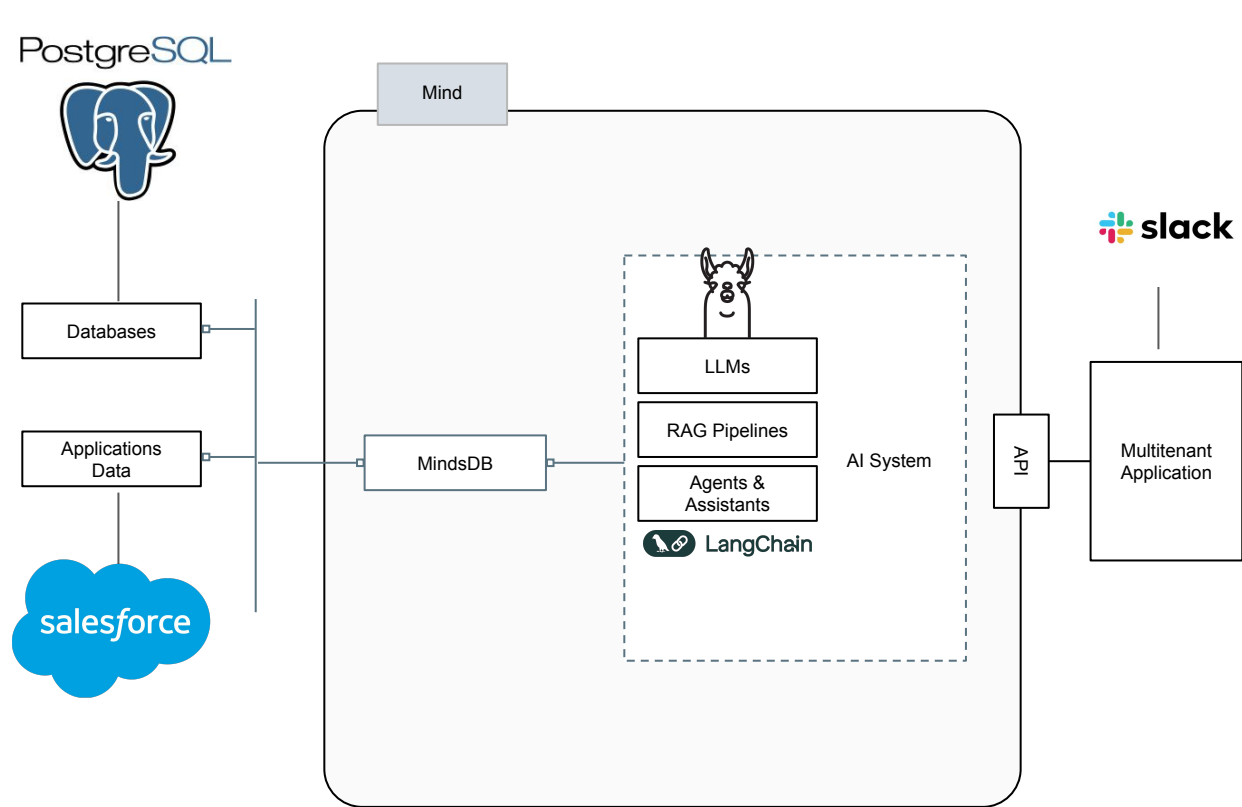
MindsDB data capabilities can help langchain developers build best in class AI systems to answer questions from all kinds of data.

It is very convenient for AI systems to only have to be really good at one dialect to access all data.



Demo

Answering questions over CRM and database data and expose that system in slack.



Demo

```
300
301 CREATE MODEL mindsdb_inference_endpoints_model
302 PREDICT answer
303 USING
304   engine = 'mindsdb_inference_engine',
305   model_name = 'mistral-7b',
306   prompt_template = 'this prompt will be overridden';
307
308 DESCRIBE mindsdb_inference_endpoints_model;
309
310 CREATE MODEL my_conv_model_mindsdb
311 PREDICT answer
312 USING
313   engine = 'langchain_engine',
314   provider = 'mindsdb',
315   model_name = 'mindsdb_inference_endpoints_model',
316   mode = 'conversational',
317   user_column = 'question',
318   assistant_column = 'answer';
319   max_tokens = 100,
320   verbose = True,
321   prompt_template = 'Answer the user input in a helpful way';
322
323 DESCRIBE my_conv_model_mindsdb;
```

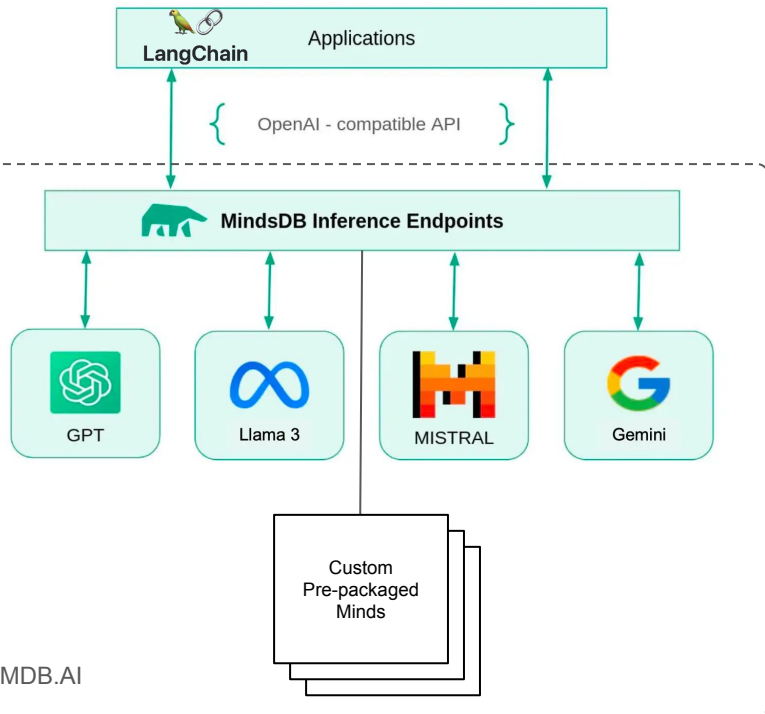
TABLES	NAME	ENGINE	PROJECT	ACTIVE	VERSION	STATUS	ACCURACY	PREDICT
1						complete	[NULL]	sentiment

and assistant columns means this column will be provided.

<https://www.loom.com/share/25dec6f224b7420f973e6db965a7883e>



MDB.AI



The screenshot shows the MindsDB web interface. The top navigation bar includes the MindsDB logo and the URL "mdb.ai/llm-serve". The left sidebar contains navigation links for "API Key", "Profile", and "Documentation". The main content area is divided into two sections:

- Manage your MindsDB API key:** This section shows an API key generated on "Apr 10, 2024" with the value "a4***f611". It includes buttons for "Regenerate" and "Revoke".
- Compare LLMs:** This section displays a list of LLMs on the left, including "claude-3-haiku", "codellama-70b", "dbrx", "firefunction-v1", "firellava-13b", "gemini-1.5-pro", "gemma-7b", "gpt-3.5-turbo", and "hermes-2-pro". The "claude-3-haiku" model is selected, showing details such as "Owner: Anthropic", "License: Proprietary", "Context window: 200k tokens", "Throughput: 88.9 tokens/second", "Quality score: 78%", "Usage: structured information extraction", and "Latency: 0.41 seconds". A "Model details" link points to "artificialanalysis.ai".

On the right side of the LLM comparison section, there is a code block showing a curl command for API interaction:

```
curl --location 'https://llm.mdb.ai/chat/completions' \
--header 'Content-Type: application/json' \
--header 'Authorization: Bearer a4***f611' \
--data '{
  "model": "claude-3-haiku",
  "messages": [
    {
      "role": "user",
      "content": "What LLM are you?"
    }
  ]
}'
```

Next Steps



We can help you implement AI strategy quickly and simply

- Identify use cases
- Help with Reference Architecture
- Delegate engineers, training and resources
- Build your custom solution

Book a Demo

<https://mindsdb.com/enterprise#contact>

Questions? Get in touch!

jorge@mindsdb.com



Questions?



AI Challenges Solved with MindsDB



SAMSUNG

- Samsung uses [MindsDB](#) to predict demand and optimize partner cellular network traffic.
- Project led by a team of Engineers, not Data Scientists, using [MindsDB](#).
- The demand predictions and optimized traffic have helped make better use of existing infrastructure with seven-figure cost savings and happy partners.



Goldman
Sachs

- Goldman Sachs uses [MindsDB](#) to spot anomalies in the central data platform.
- Goldman Sachs IT and dev team have integrated [MindsDB](#) directly their data stores. This is securely deployed on-premise.
- Their Full-Stack Engineer definition has expanded to include ML via [MindsDB](#)
- This has lead to an increase in data quality, reduction of errors (that can cost millions) and a simplified system with less admin and maintenance.



olla.

- Olla.ai is a fast moving startup that takes an 'AI-first' approach
- [MindsDB](#) allowed Olla.ai developers to implement advanced AI features into their core platform quickly
- Olla.ai uses the power of OpenAI's GPT-4 to build an expert assistant for elderly care support
- [MindsDB](#) allows Olla.ai engineers to quickly connect OpenAI, their product platform, multiple data sources, multiple 3rd party systems and manage and maintain their AI features in production with little effort