

# Datakafé



*Brewing up data goodness*

[datakafe.com](http://datakafe.com)



**Brian** 

**Product | Utilities SME**

15 Yrs Industry Experience in Finance,  
Corp Strategy, MBA



**Robert Leach** 

**Product**

3x Founder, Columbia Adjunct Faculty,  
Former Accenture Innovation Center



**Jon Saginaw** 

**Technology**

Former Accenture Partner,  
Tech Innovator (3 patents)

# A little bit about us

- **Regional Data/AI Company, since 2015**

Strong focus on Utility Industry, High Tech, Telco, Retail, Federal

- **Help companies build the modern data stack, focused on Databricks**

Data Engineering, Data Governance, ML & AI Application Development

- **We're Builders & Open Source Contributors (LangChain)**

We ship products and build with the modern tech stacks (Git, HuggingFace, Auth0, Azure, AWS, GCP, Railway, Docker, Mongo, Colab, Pinecone)

# What we do...


Data Engineering

Data/AI Strategy & Data Governance

AI Product Development



Now focused on  
advancing the Utility Industry  
forward with...



**utilityGPT**

AI Applications for Utilities

Clean Energy Transition

| Resiliency

| Reliability

# The future is domain specific AI

## The future of generative AI is niche, not generalized

ChatGPT has sparked speculation about artificial general intelligence. But the next real phase of AI will be in specific domains and contexts.

By Ken Mugrage

MIT  
Technology  
Review

Starting with

Rate Case AI...

# Situation - Rate Case Process

- Utilities don't operate in open markets to determine rates
- 10-12 mth process requiring written testimony, exhibits, Data Requests (DRs)
- This happens in the public domain (public data, protests, petitions)
- The Public Service Commission determines fair increases



# Problem

- The effort is organizationally exhausting and lengthy
- Decentralized story creates quality, consistency, and strategic issues
- Clean Energy Transition driving more change & attention

**Rate requests by energy utilities hit another all-time record in 2022**

# Benefits of Utility AI

an AI copilot for utilities built for

- Better rate case outcomes
- Lower cost research and analysis
- Increased consistency and control
- Increase productivity (less time searching)

# Features

- **Ask** questions of your data with citations
- **Tools** for benchmarking, composition & more
- **Threads** for organizing research & collaboration
- **Infobase** preloaded with public documents

# DEMO



Utility AI

 Utility AI



Copy link

## AI Powered Productivity for Utilities

Questions about the C  nergy Transition, Rate Cases, Benchmarking. **Our chat has the answers**

Login

Sign up

Watch on  YouTube

<https://youtu.be/EBbNaTsTca0>

# Why

- Regulatory industries have natural use cases for AI copilots

*Vast amounts of data, complex processes, underserved market > massive change over next 30 years with clean energy transition*

- ChatGPT, Bing is general AI

## The future of generative AI is niche, not generalized

ChatGPT has sparked speculation about artificial general intelligence. But the next real phase of AI will be in specific domains and contexts.

By Ken Mugrage

MIT  
Technology  
Review

- Provide a low-risk way for customers to experience AI use cases

*Build supportable product/accelerator focused on UI/UX*


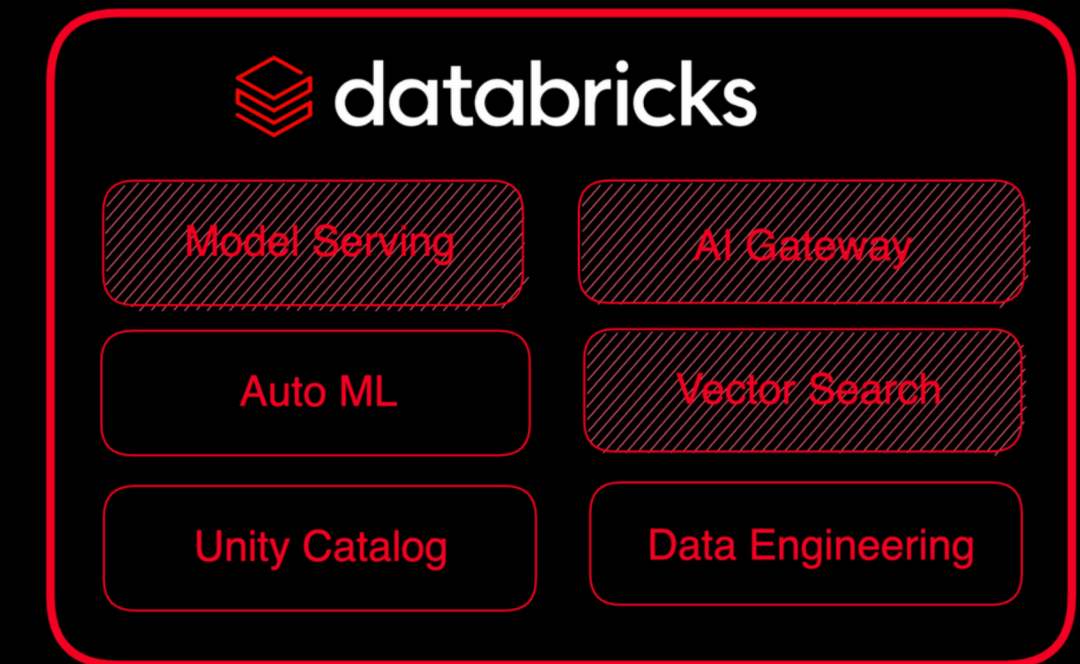
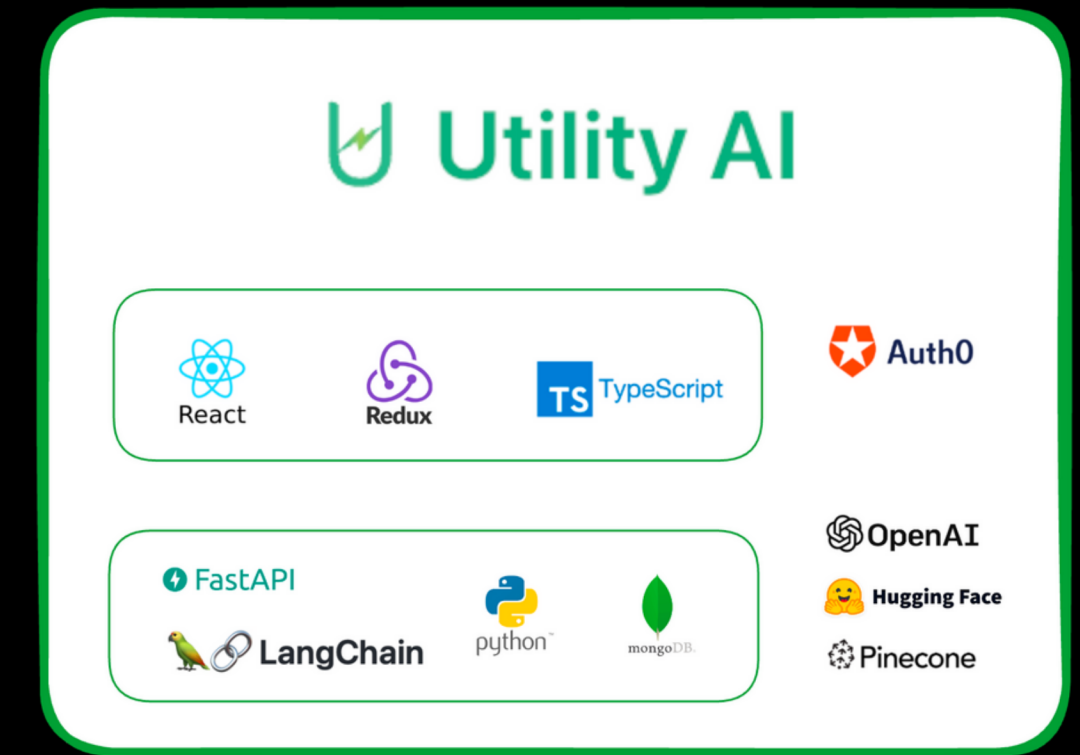
# Features

- **Ask** questions of your data with citations
- **Tools** for benchmarking, composition & more
- **Threads** for organizing research & collaboration
- **Infobase** preloaded with public documents

# Architecture

# Tech Stack

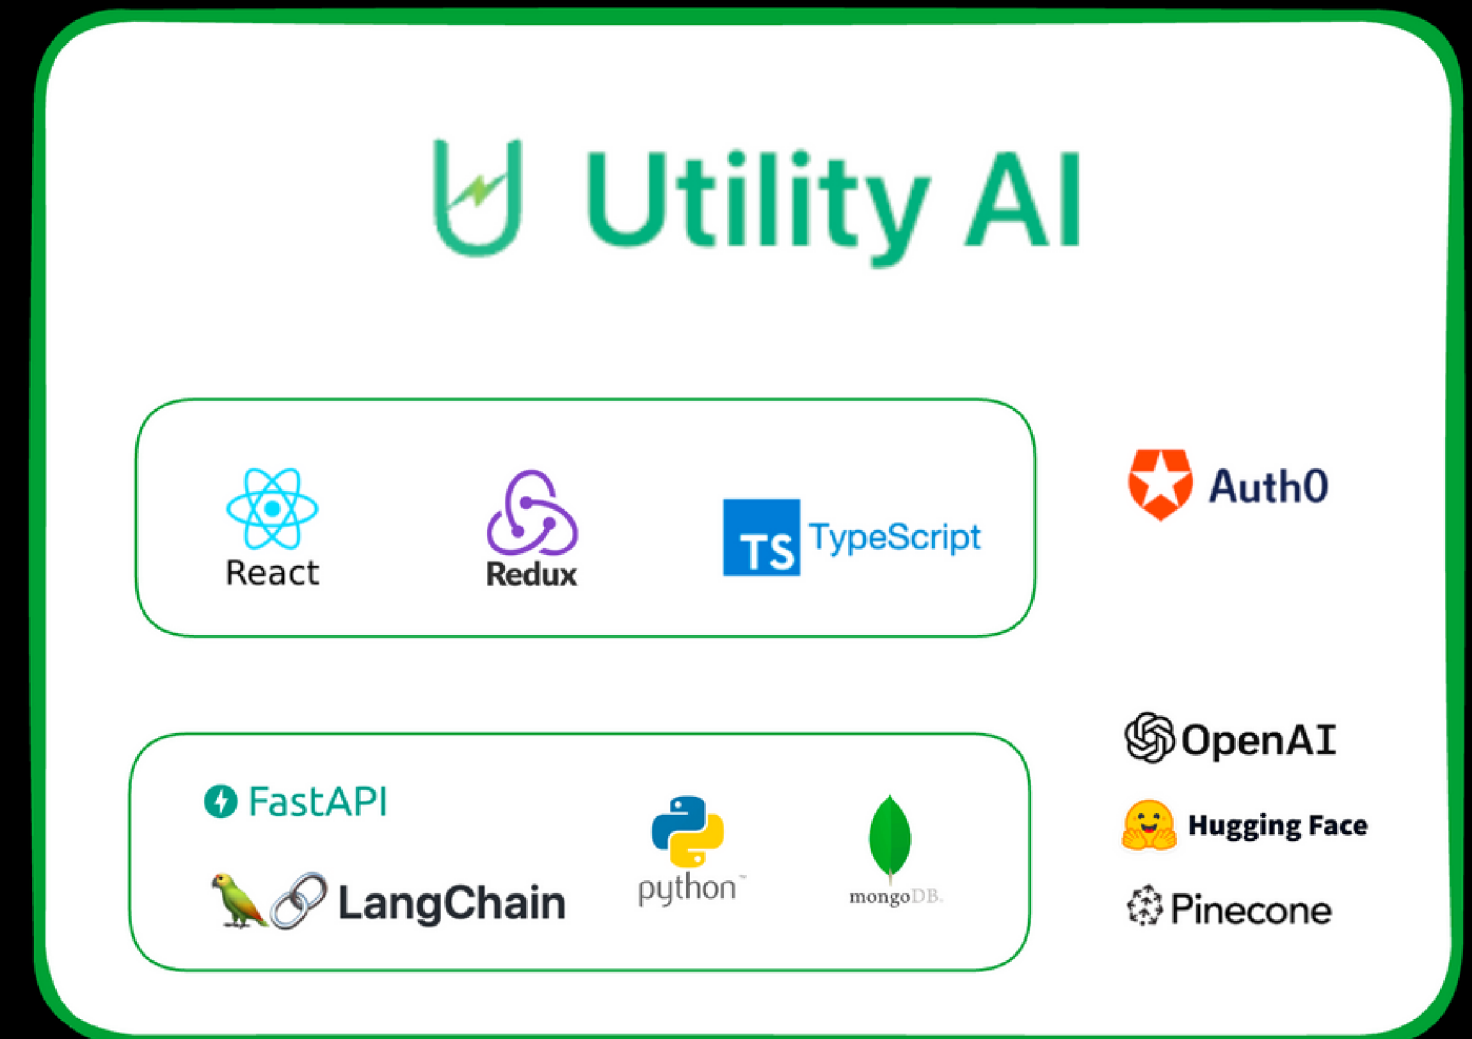
- LangChain: for LLM building blocks
- LangSmith: quality control, monitoring
- LLMs: OpenAI, Azure Open AI, OSS Models (Mosaic 7B)
- AQE: Automatic Query Expansion > dictionary for better AI results
- Pinecone: for vector store
- Databricks: Data Processing, Vector Creation, Data Governance
- Microsoft: Hosting and Security





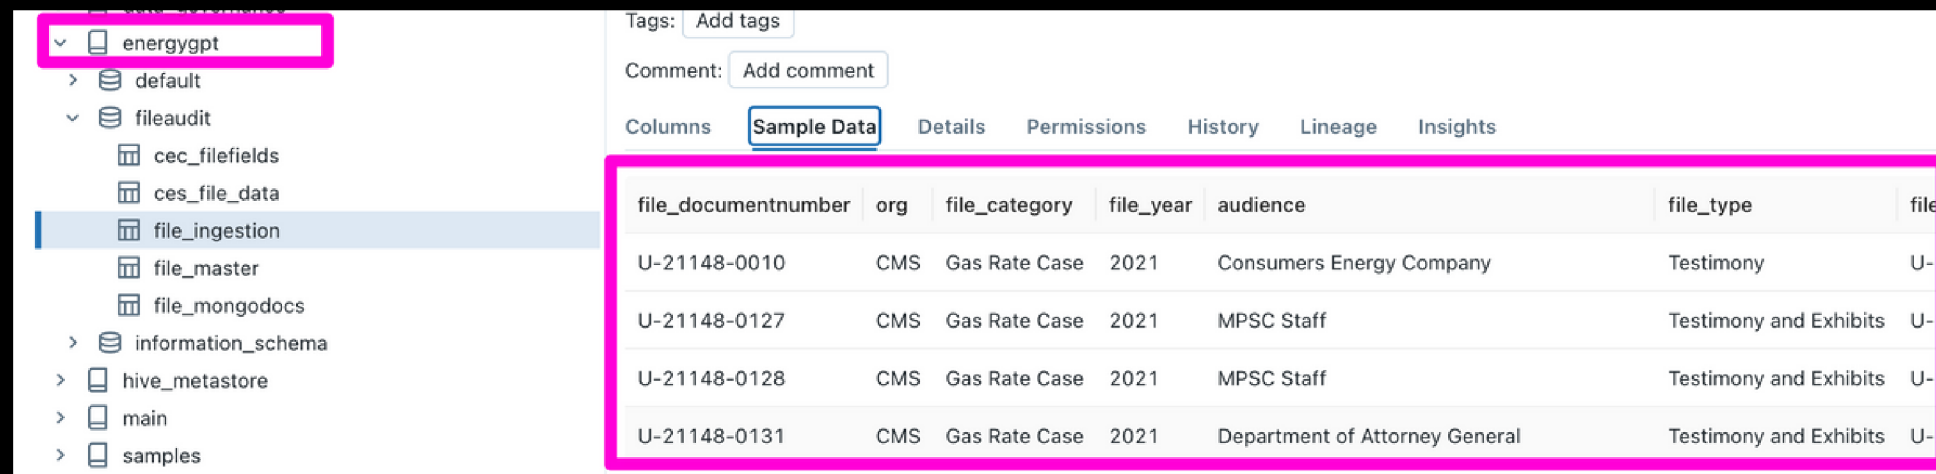
# Frontend & Backend

- React frontend for optimized cust experience
- Python backend with Fast API
- Auth0 for enterprise single sign-on
- Building off industry leading packages




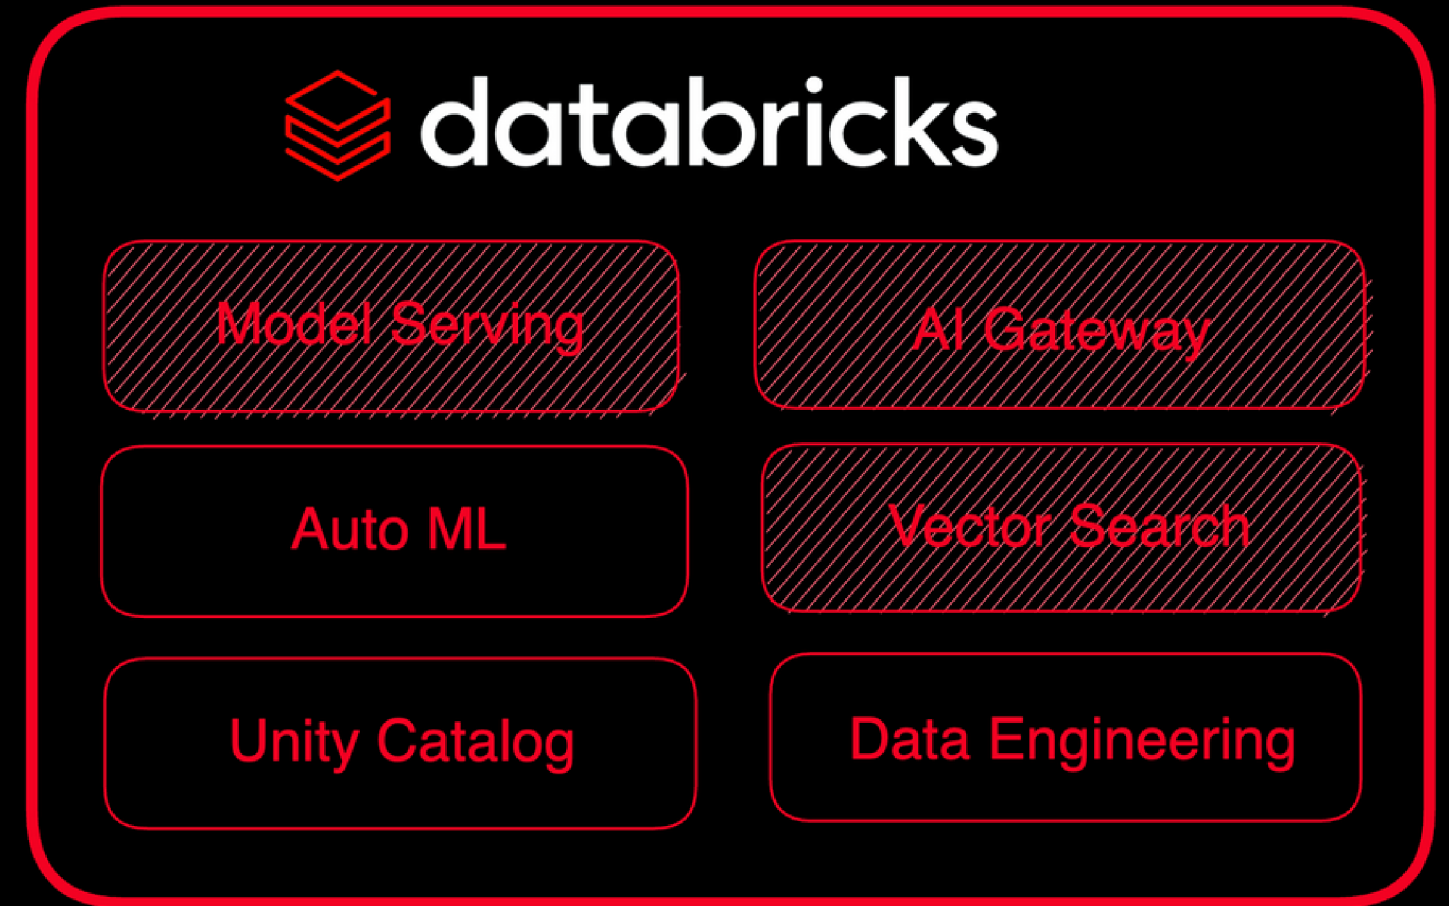
# ML & Data Engineering

- Configure & evaluate different LLMs
- Evaluate, Fine-tune OSS LLMs
- Leverage emerging Databricks features
- Govern Data with Unity



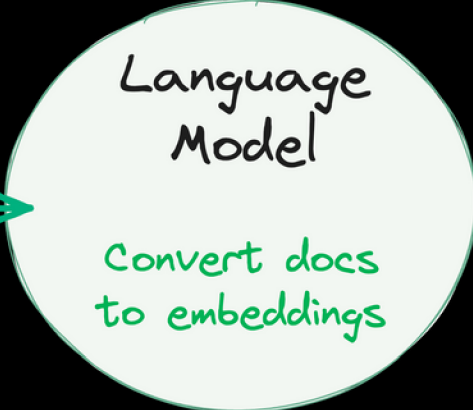
The screenshot shows the Databricks Unity Catalog interface. On the left, a navigation pane shows a folder structure under 'energygpt', with 'file\_ingestion' selected. The main area displays a table with columns: file\_documentnumber, org, file\_category, file\_year, audience, file\_type, and file\_id. The table contains four rows of sample data.

file_documentnumber	org	file_category	file_year	audience	file_type	file_id
U-21148-0010	CMS	Gas Rate Case	2021	Consumers Energy Company	Testimony	U-
U-21148-0127	CMS	Gas Rate Case	2021	MPSC Staff	Testimony and Exhibits	U-
U-21148-0128	CMS	Gas Rate Case	2021	MPSC Staff	Testimony and Exhibits	U-
U-21148-0131	CMS	Gas Rate Case	2021	Department of Attorney General	Testimony and Exhibits	U-

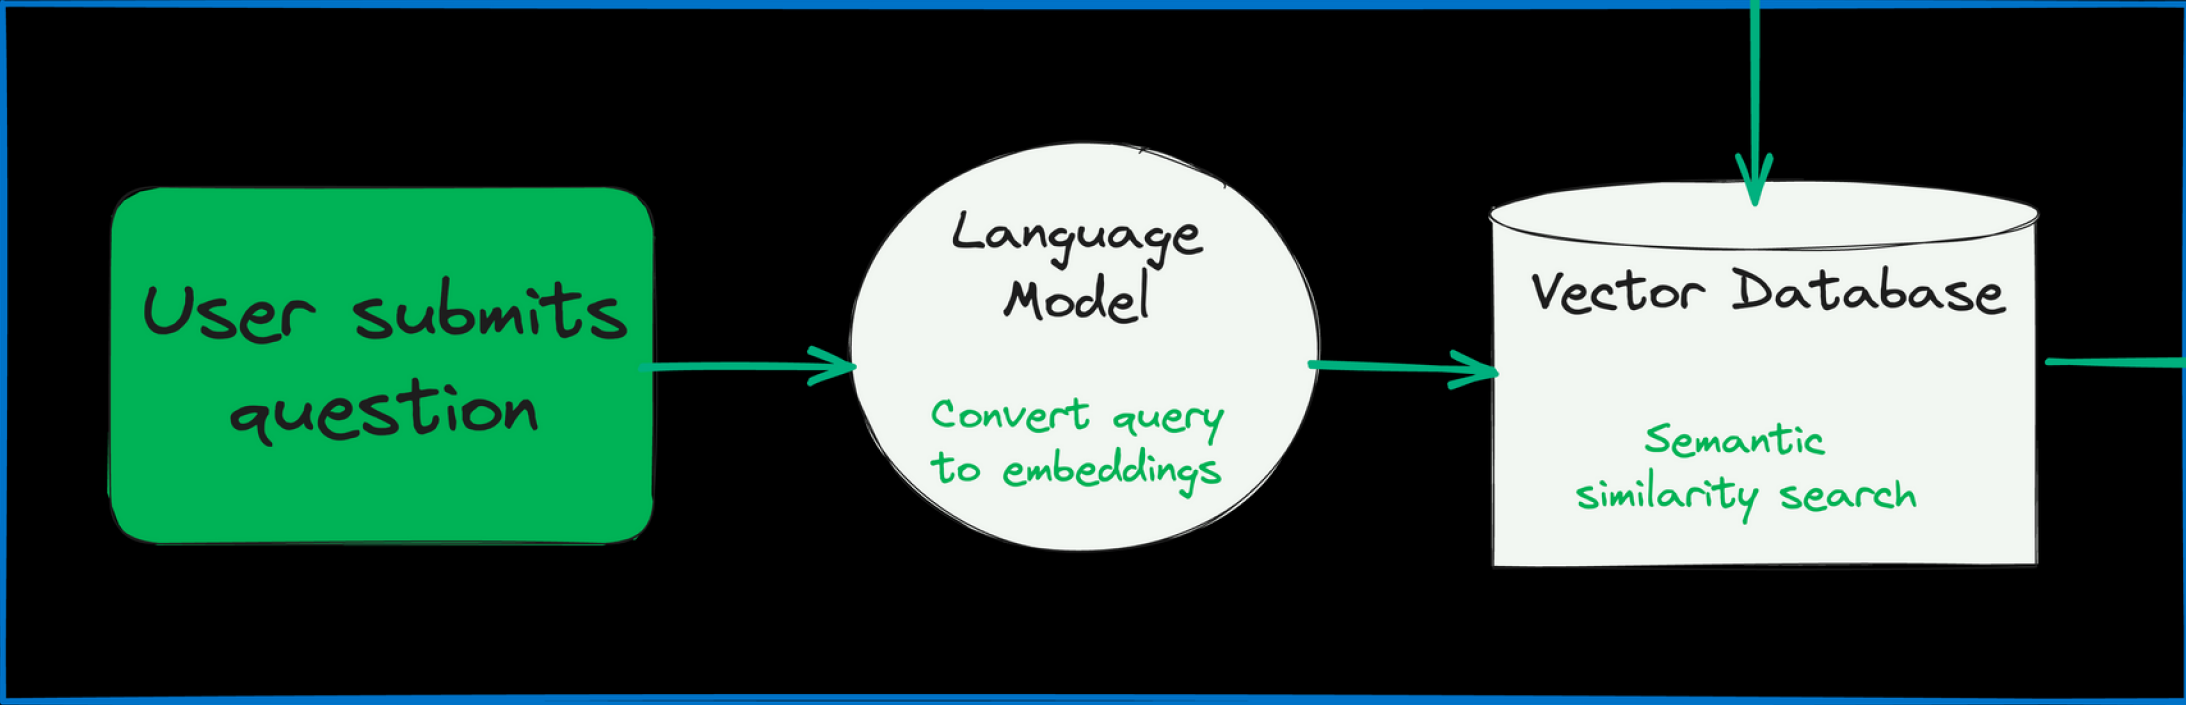
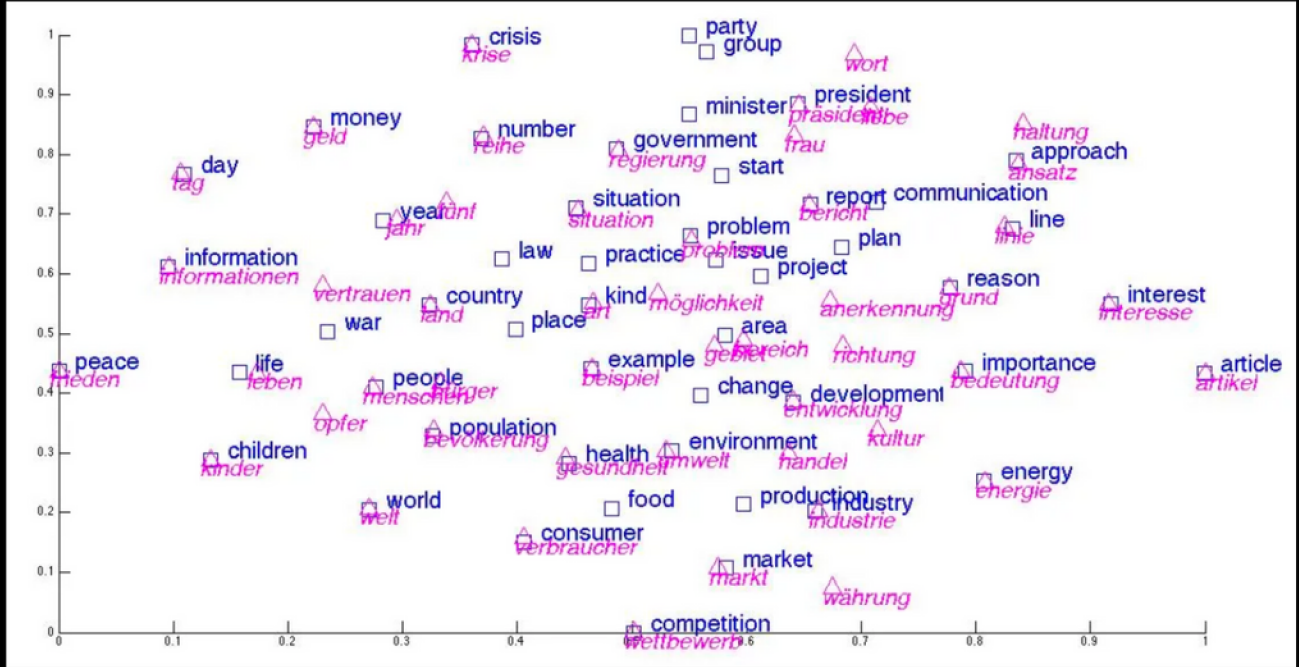


# Retrieval Augmented Generation (RAG)


Documents



0.7	0.3	0.8	-0.7	0.1	-0.5	-0.4
-----	-----	-----	------	-----	------	------

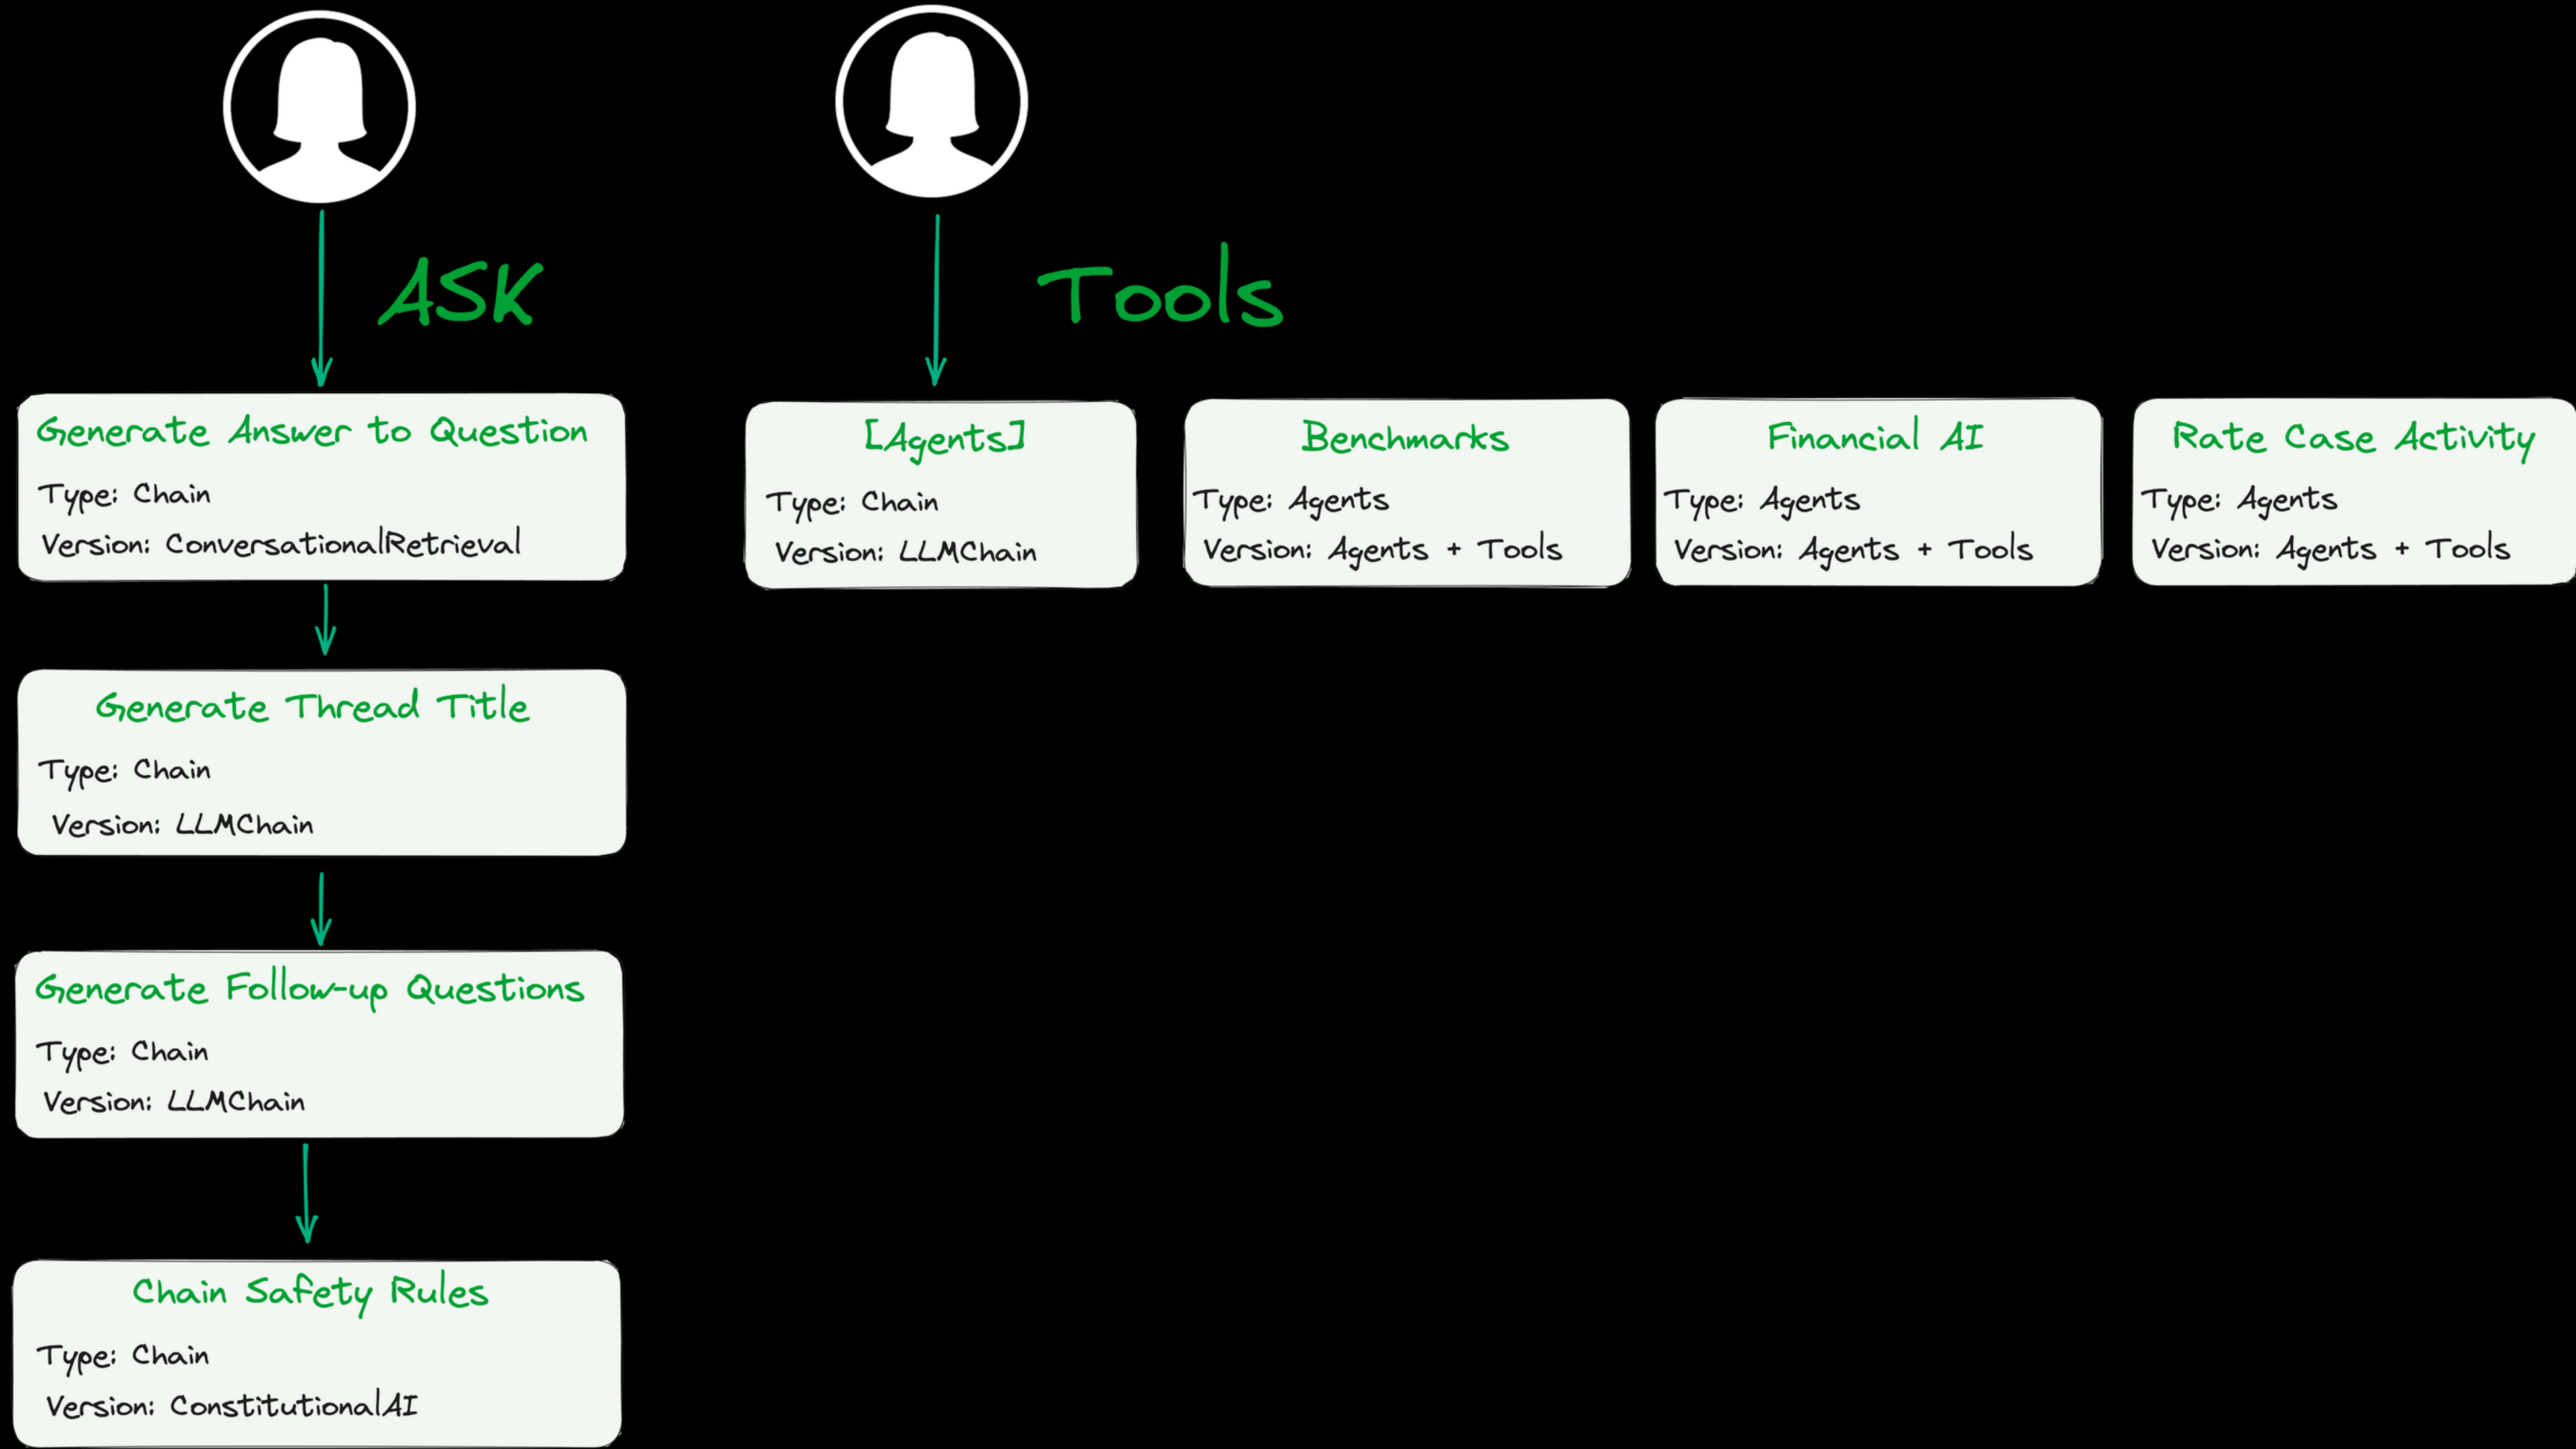


Similar docs are passed as context



User receives response

# Tools Architecture



# Security, Privacy, Safety

- All data encrypted in transit and at rest
- All data systems SOC2 Type II certified
- Implementations using Azure Open AI
  - Strong Enterprise Standards
  - Content Filters
  - Abuse Monitoring\*

## 📌 Important

Your prompts (inputs) and completions (outputs), your embeddings, and your training data:

- are NOT available to other customers.
- are NOT available to OpenAI.
- are NOT used to improve OpenAI models.
- are NOT used to improve any Microsoft or 3rd party products or services.
- are NOT used for automatically improving Azure OpenAI models for your use in your resource (The models are stateless, unless you explicitly fine-tune models with your training data).
- Your fine-tuned Azure OpenAI models are available exclusively for your use.

The Azure OpenAI Service is fully controlled by Microsoft; Microsoft hosts the OpenAI models in Microsoft's Azure environment and the Service does NOT interact with any services operated by OpenAI (e.g. ChatGPT, or the OpenAI API).

*Azure Open AI Data Privacy Terms*

*\*May include human reviewers (MS Employees), but option to opt-out*

# Sample Work

# Weather Impact Model

Goal: Predicting outages based on the weather events for crew staffing

2.02 Weather Forecast & Impact Home Catalog Favorites Dashboards New Open Signed In As Leach, Robert (Admin Accou

Weather QuickLook ConEd Forecast ConEd vs. External Forecasts Observations **Outage Scenario Planner** Historical Outages ConEd Forecast App Quicklook Forecast App

### Weather Parameters

- \* Calculate Predictions: YES
- Forecast Date: 07/23/2020
- \* Region: BW
- \* Temperature: 30
- \* Precipitation Forecast: 0.5
- \* Wind Gust: 50
- \* Peak Wind Gust Direction: NE
- \* Foliage %: 100
- \* Rainfall Amount: 0.5
- \* Lightning Forecast: NO
- \* Severe Thunderstorm Forecast: NO
- \* Flood Forecast: NO
- \* Tropical System Forecast: NO
- \* Yesterday's Rainfall Amount: 0.5
- \* Last 5 Days Rainfall Amount: 0
- \* Snow Fall Amount: 0
- \* Snow Ratio: 9
- \* Coastal Flooding: 10

Apply Reset

### Weather Impact

Forecasted Job Range			
Rules	Low	Expected	High
The Low and High are calculated using the values 0.5 and 1.5 )	112	223	335

Forecasted Customer Range			
Rules	Low	Expected	High
a. Expected : # of Predictions Expected * 68			
b. Low : Expected - ( # of Predictions Expected * 68*0.28)	10,935	15,188	19,440
c. High: Expected + ( # of Predictions Expected * 68*0.28)			

Soil Saturation	
Rules	Predicted
Sum of the Values from CUS_BW_MAXDISCHARGE_COEFF table divided by the respective median and standard deviation for that day.	95.35%

Foliage	
Rules	Predicted
Value from the CUS_BW_MAXDISCHARGE_COEFF table.	-0.24%

### Forecast Probably Curve

Value	Prediction %
0	0%
112	~25%
223	50%
335	~25%
402	0%

# Customer Experience Analytics

Goal: Insight into Customer Experience, during a storm





# Critical Customer Alerts

Goal: Identify customers with critical needs during outage event, and proactively reach out

## ATTENTION!

A critical customer response was detected that may require immediate attention.

Date	Ph Num.	Customer SMS	Category
3/13/2023 14:34:59	646-420- [REDACTED]	Who do I call my grandmother on oxygen. She was diagnosed last Monday with oxygen	Oxygen