

Datakafé



A technology services company
helping organizations build
modern data and AI products

Who We Are

- Founded in 2015 by ex-Accenture technologists
- Specialized in Databricks, Open Source Frameworks & all major clouds
- Implemented multiple Petabyte-scale Data & AI Solutions for Fortune 500 companies
- End-to-End Systems Integration and Use Case Development
- Strong focus on Innovation & AI Product Development: Rapid Prototyping + Shipping Fast
- Design and execute Data & AI Strategies (including Data Governance) for Organizations

Management Team



Robert Leach | Head of Delivery

Over 20 years of experience delivering data analytics solutions for Fortune 500 companies and the US Government, Robert has led global teams in implementing data modernization programs.

He is a **Databricks Platform Architect** and serves as adjunct faculty and guest lecturer on Data and AI at Columbia University and the University of Maryland.



Jon Saginaw | CTO

Over 20 years in strategy, management, and technology consulting, including 16 years at Accenture, Jon has helped major companies solve the most complex problems. He holds three patents, actively contributes to open source projects, and has launched several consumer Data and AI products in recent years.

Delivery

Our team is comprised of seasoned delivery professionals and certified Databricks practitioners:

- Databricks Platform Administrators
- Databricks Platform Architects
- Azure Databricks Architects
- Databricks Data Engineers Associates
- Databricks Data Engineer Professionals
- ML Specialists
- Generative AI Engineers

 **databricks**
Consulting Partner

 **databricks**
Authorized Training Partner

 **Microsoft**
Partner

Past Performance

We have delivered Data & AI Solutions and shipped modern data products for numerous Fortune 500 companies and startups



Services

1

Data & AI Strategy

We work with customers to create Data and AI roadmaps, governance frameworks, data architectures, and tech stacks to drive business outcomes with data.

2

Data Professional Services

From Lakehouse build-outs to Unity Catalog migrations, our team will leverage our solution accelerators to deliver end-to-end data solutions.

3

AI Product Development

This is our lean-startup methodology, where we deliver MVP products in days to weeks, not months.

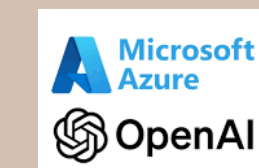
4

Data & AI Training

We've empowered 100's of Data Engineers with the skills and knowledge to build data analytics and AI solutions on the Databricks, AWS, and Microsoft stack.

Technical Focus

- Databricks Infrastructure and DevOps
- Data Architecture (medallion)
- Data Lakehouse implementations
- Data Governance, compute governance
- Power BI to Databricks Integration
- Unity Catalog Configuration and Migrations
- Lakehouse Monitoring
- Databricks GenAI model serving
- Genie Spaces
- Databricks Asset Bundles
- LangChain
- CrewAI
- Railway
- FastAPI
- Streamlit
- REACT front-end development
- OAuth



Why Datakafé

Our top 5 differentiators

- 1 Team Experience:** Our team has a proven track record and a combined 40+ years of delivery experience. We're lean and hands on with our customers.
- 2 Specialized Knowledge:** We're not everything to everyone. We focus on Databricks and open-source frameworks, as well as major cloud platforms.
- 3 Innovation:** We stay hands-on with the latest technologies, developing and delivering first-hand insights directly to our customers, beyond just PPT presentations.
- 4 Comprehensive Solutions:** We have battle tested reference architectures and solutions accelerators to delivery performant and maintainable solutions.
- 5 Strategy Led:** We know the importance of starting with the end in mind, and developing smart Data & AI strategies that drive outcomes.

Data, Analytics, Traditional ML Case Studies

Select case studies 2020 to 2024



Public Demos:

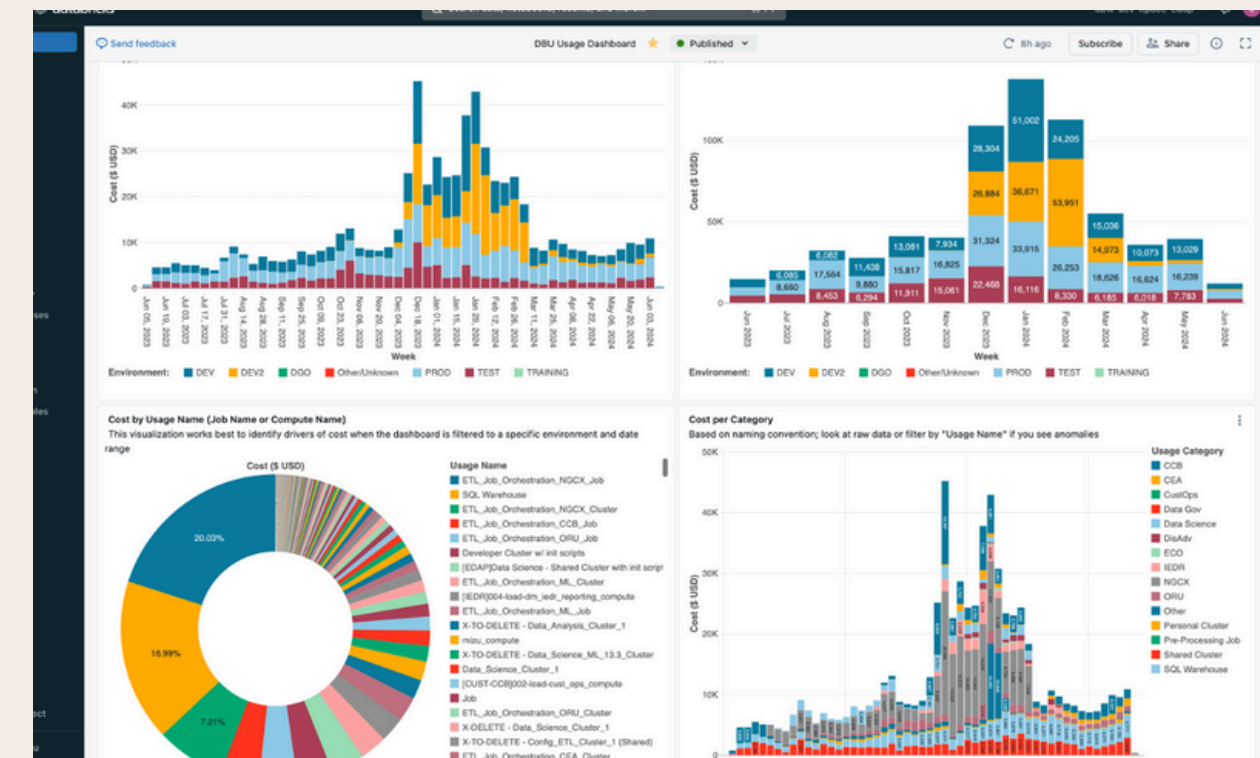
- [UtilitiesGPT](#)
- [GreenMarbles](#)
- [MyDailyPod](#)
- [Full Portfolio](#)

Customer Experience Lakehouse Optimization

Situation

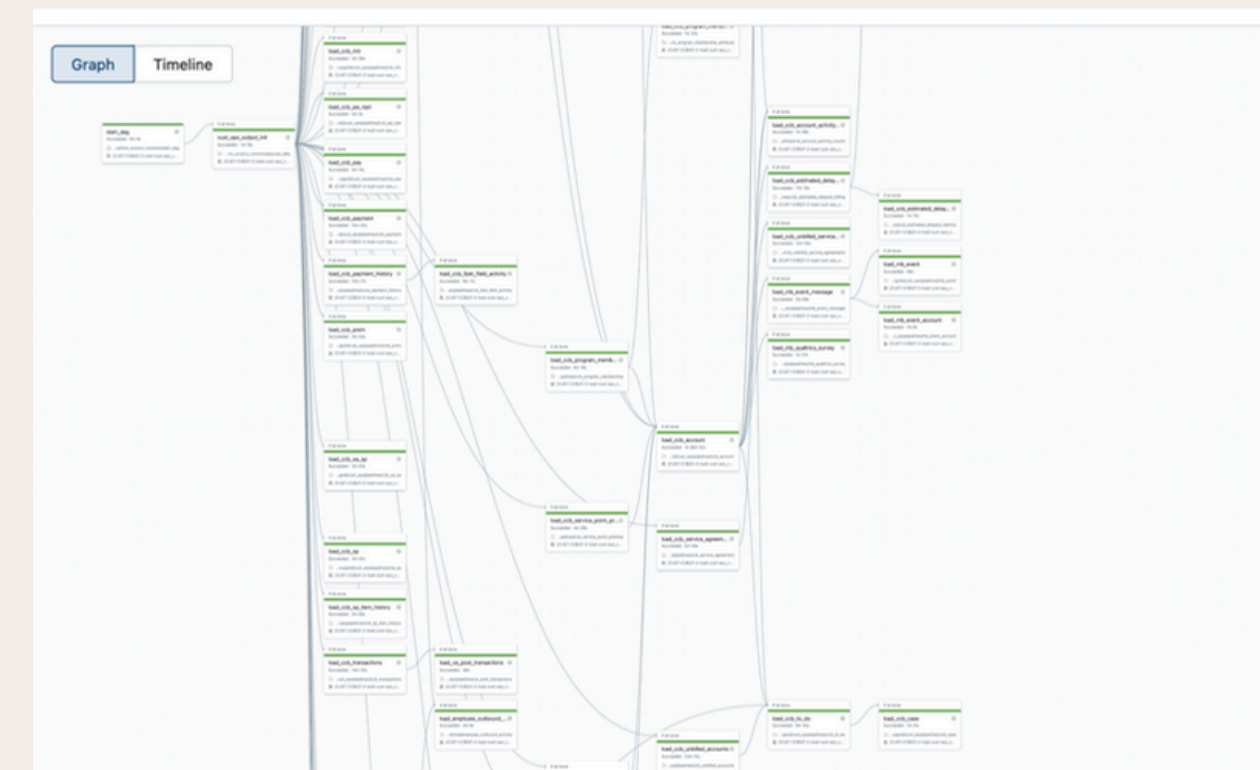
A large utility company faced rising compute costs, unreliable data pipelines, and mounting technical debt for their Customer Analytics Lakehouse. Datakafé was brought in to rearchitect the solution, reduce costs, and improve performance

- Petabyte-scale Lakehouse serving 200+ users
- 14 analytics use cases; 90 data sources, 500+ files, APIs
- Power BI front-end
- Custom SQL Database Orchestration



Solution

Simplified the orchestration architecture, migrating from a metadata driven SQL Database and Python design, to using Databricks Workflows with JSON config and PySpark. Migrated 250+ workloads to new orchestration



Outcomes

Increased performance by 30% and lowered costs by 50% due to right sizing clusters. Additionally, increased stability and maintainability of solution by simplifying the end-to-end architecture.

Unity Catalog Implementation / Migration

Situation

A large company needed to improve their data governance capabilities to ready themselves for scaling out more projects, and for expected GenAI use cases. With an existing footprint of over 500 data feeds, and 2000 tables, the migration to unity needed to be carefully planned and implemented to not impact production workloads.

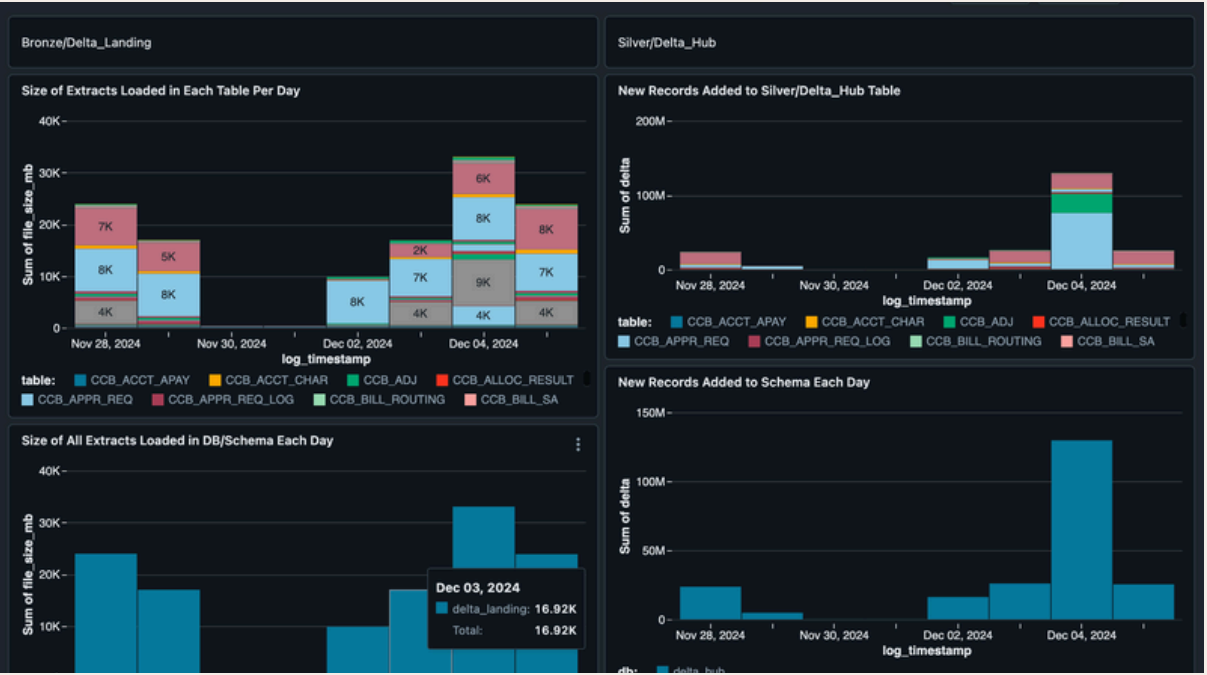
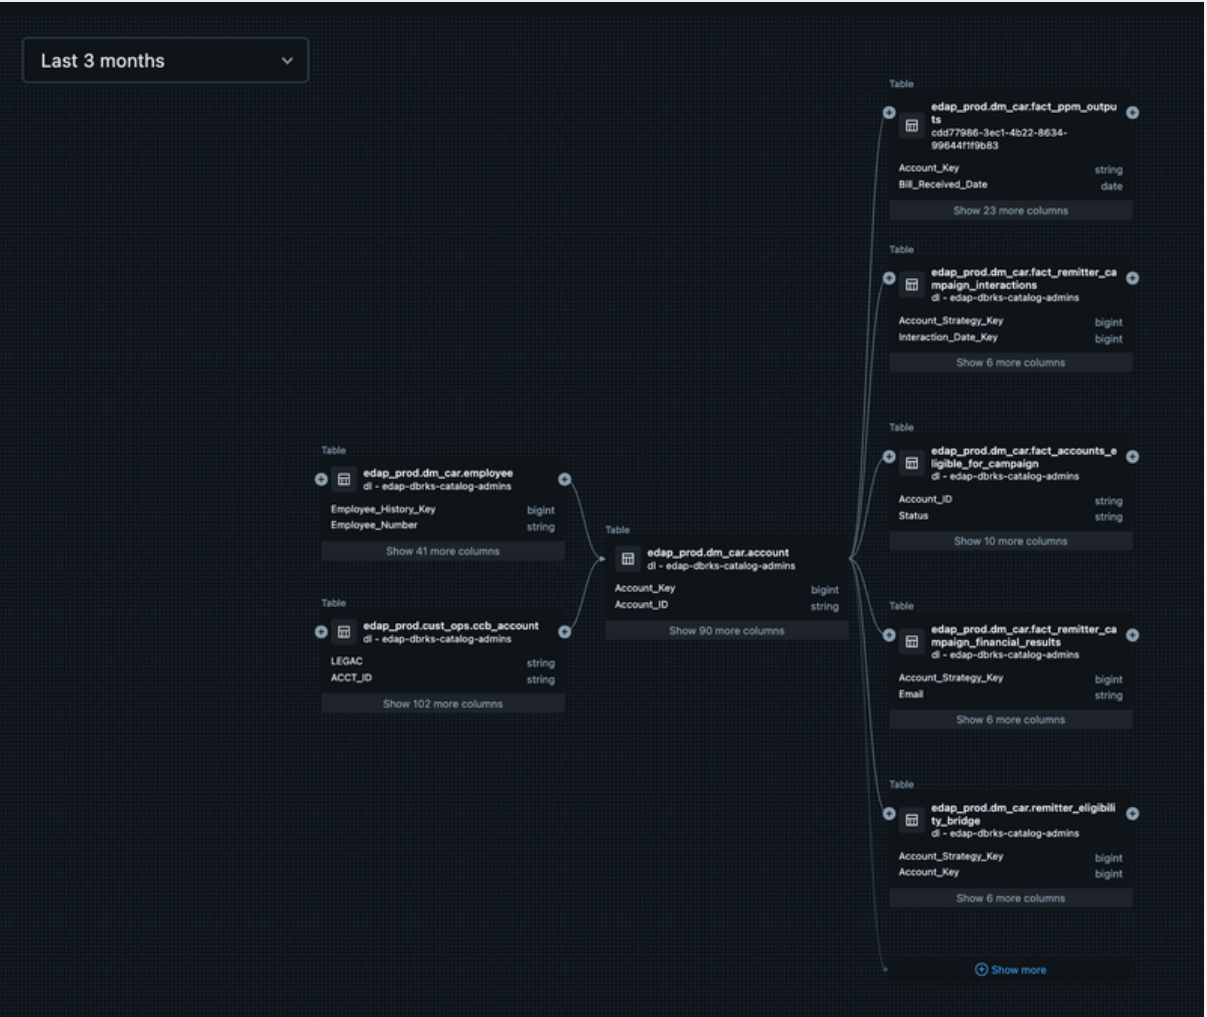
Solution

We designed a new taxonomy and naming standard alongside a POC to showcase improved organization, data discovery, and features like Lakehouse Monitoring and Databricks AI Assistant, which could be delivered with little impact. This was in addition to showing the administration benefits of user management.

Next we used the UCX utility to catalog tables, mount points, and ACLs, and took a phased approach to implementation: first syncing the Hive tables and changing over pipelines and job configs for Unity, then migrating the data to Unity completely.

Outcomes

Successfully rolled out Unity Catalog, enabled monitoring and alerting to improve SLAs, lineage for traceability, and the Databricks AI assistant for increased productivity.



Customer Wellness Check-in

Situation

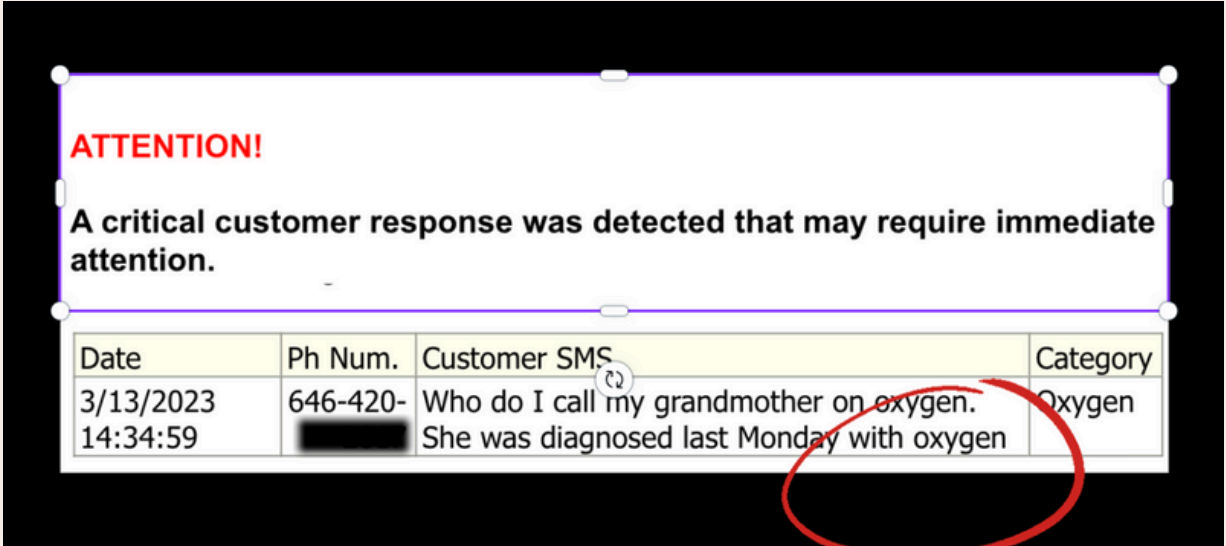
A customer experience organization was looking to identify basic KPIs on mobile communications opt-in and opt-out for campaigns. In researching these KPIs and performing data analysis, it was discovered that users were replying to texts stating they were in need of emergency services, as they were stuck in elevators, or in need of oxygen.

Solution

Developed a data streaming application that ingested text message data, analyzed it for emergency keywords, and send automated alerts to the call center to perform wellness checks.

Outcomes

This award winning application is credited with saving human lives, increasing CSAT scores, and helping update the customer master record, as all customers requiring life sustaining equipment needs to be captured for proper handling during emergencies.



The screenshot shows a monitoring interface with a black background. At the top, a white box contains the text "ATTENTION!" in red, followed by "A critical customer response was detected that may require immediate attention." Below this is a table with four columns: Date, Ph Num., Customer SMS, and Category. The first row of data shows a date of 3/13/2023 14:34:59, a phone number 646-420- followed by a redacted area, and a message "Who do I call my grandmother on oxygen." The second row shows the same date and phone number, with the message "She was diagnosed last Monday with oxygen". The "Oxygen" category is listed in the fourth column. A red circle highlights the "Oxygen" category in the second row.

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



AI case studies

Select case studies from 2023-24

UtilitiesGPT - Rate Case AI

Situation

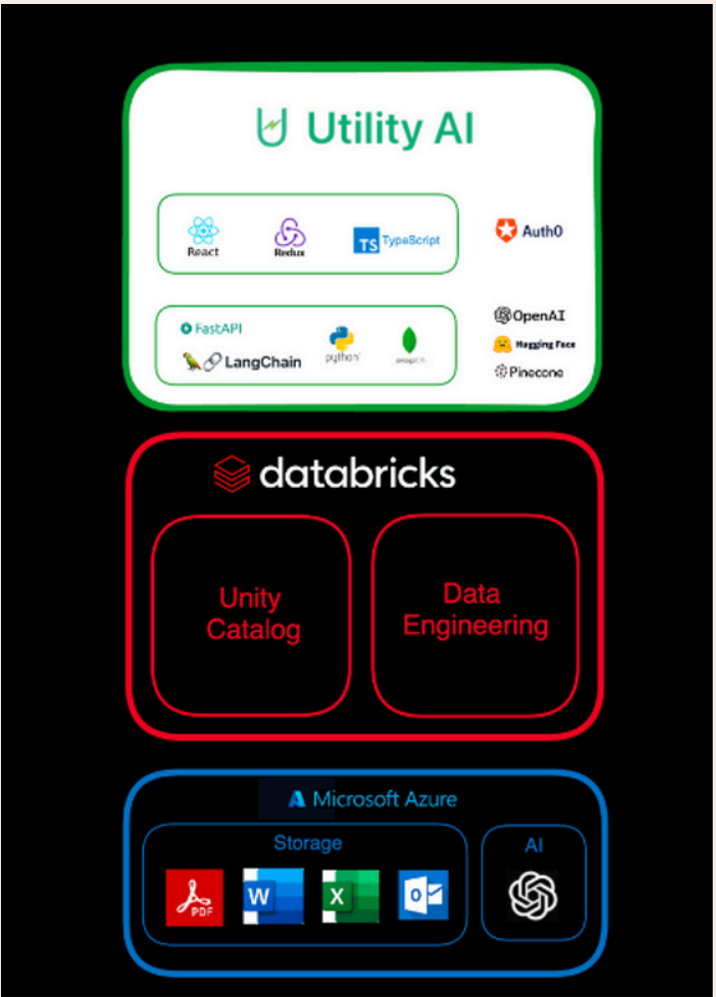
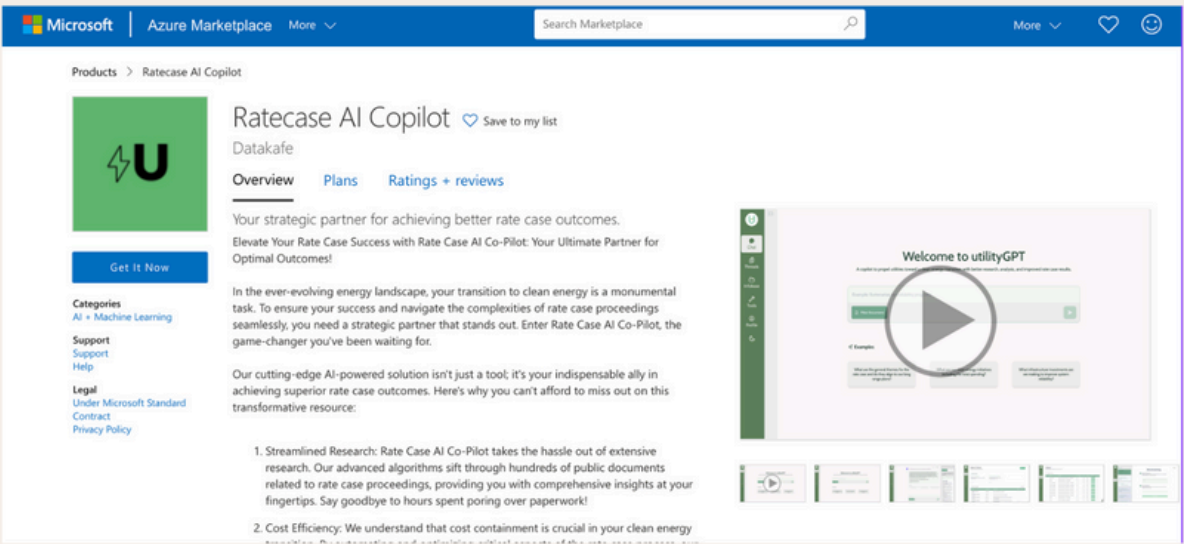
Several large utilities had a goal of improving their Rate Case process, a procedure that is crucial for justifying the rates charged to customers while ensuring regulatory compliance and operational efficiency. This process is organizationally exhausting requiring business users to research hundreds of pages of regulatory documents, and author hundreds of pages of legal testimony.

Solution

Developed Utility AI, an AI application with Retrieval Augmented Generation capabilities allowing users to chat over historical PDF regulatory documents, as well as agent based tools for benchmarking and writing high quality testimony.

Outcomes

Developed and deployed Utility AI to Azure Marketplace, integrating with Databricks, and launched it to three utility companies. This uncovered new benchmarking insights and increased efficiency in researching historical regulatory filings.



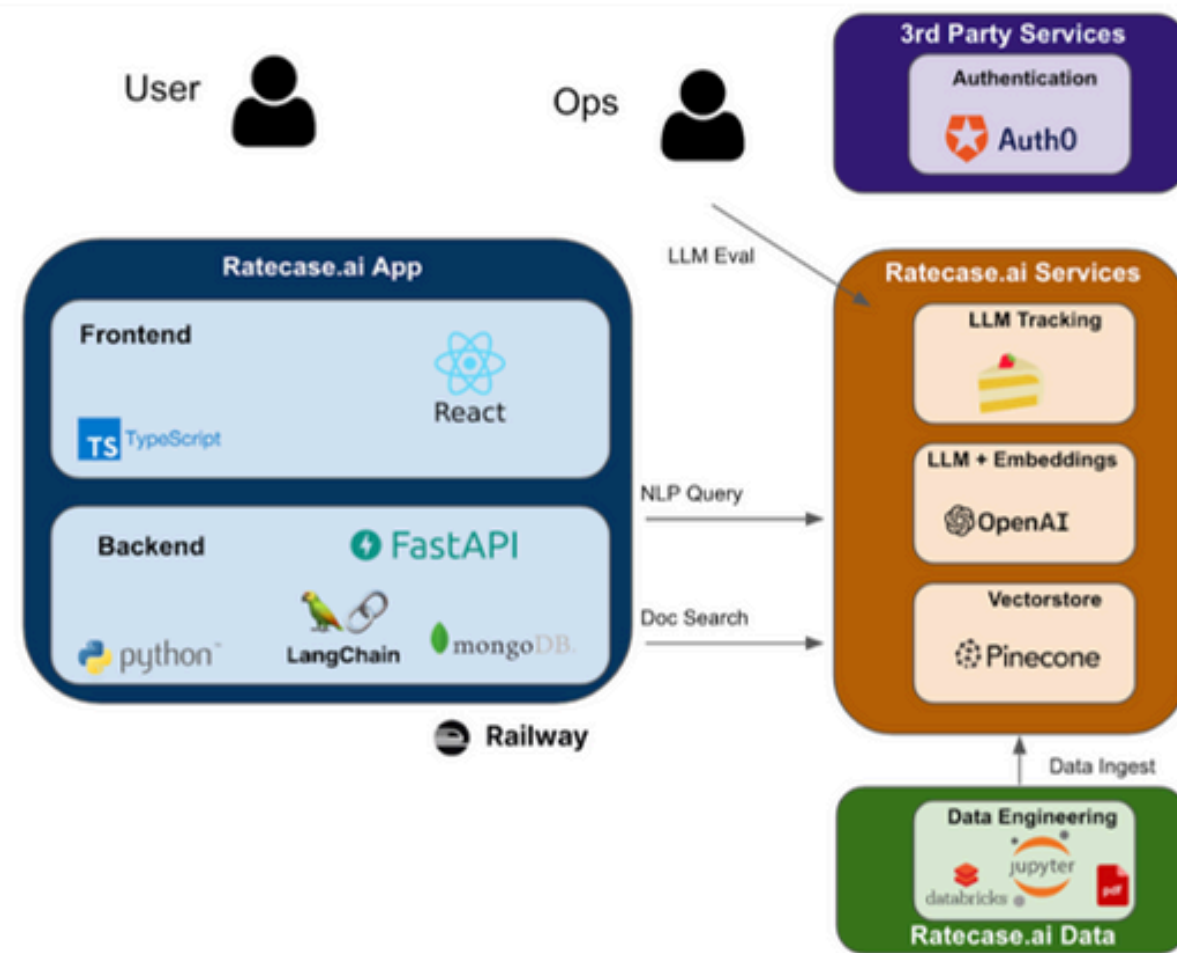
UtilitiesGPT - Rate Case AI (technical)

Demo: [See video demo on utilitiesGPT.com](https://utilitiesGPT.com)

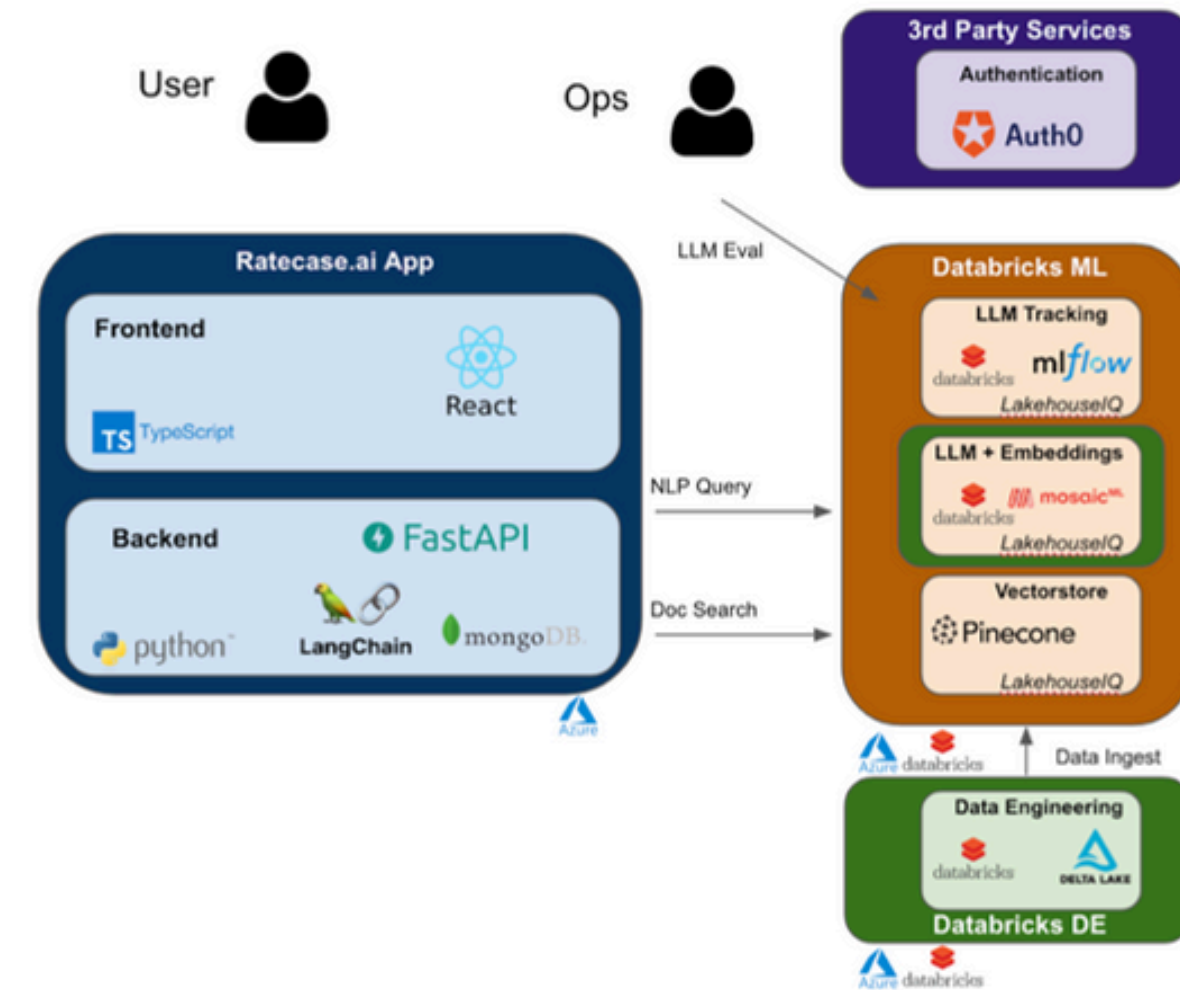
[Slide demo](#) (slide 52)

Architecture:

Option 1: Hosted Platform
Current State



Option 1: Hosted Platform
Databricks Roadmap

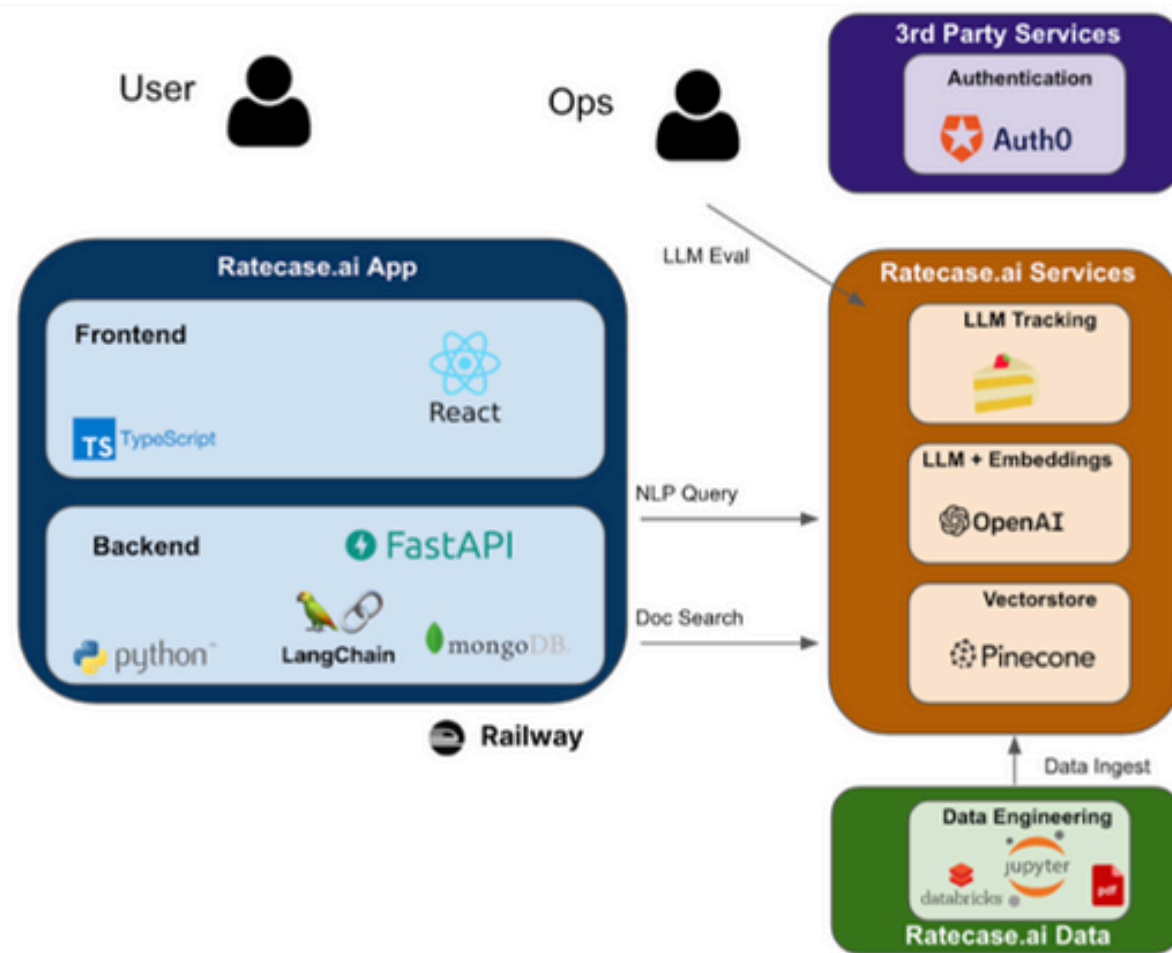


UtilitiesGPT - Rate Case AI (technical)

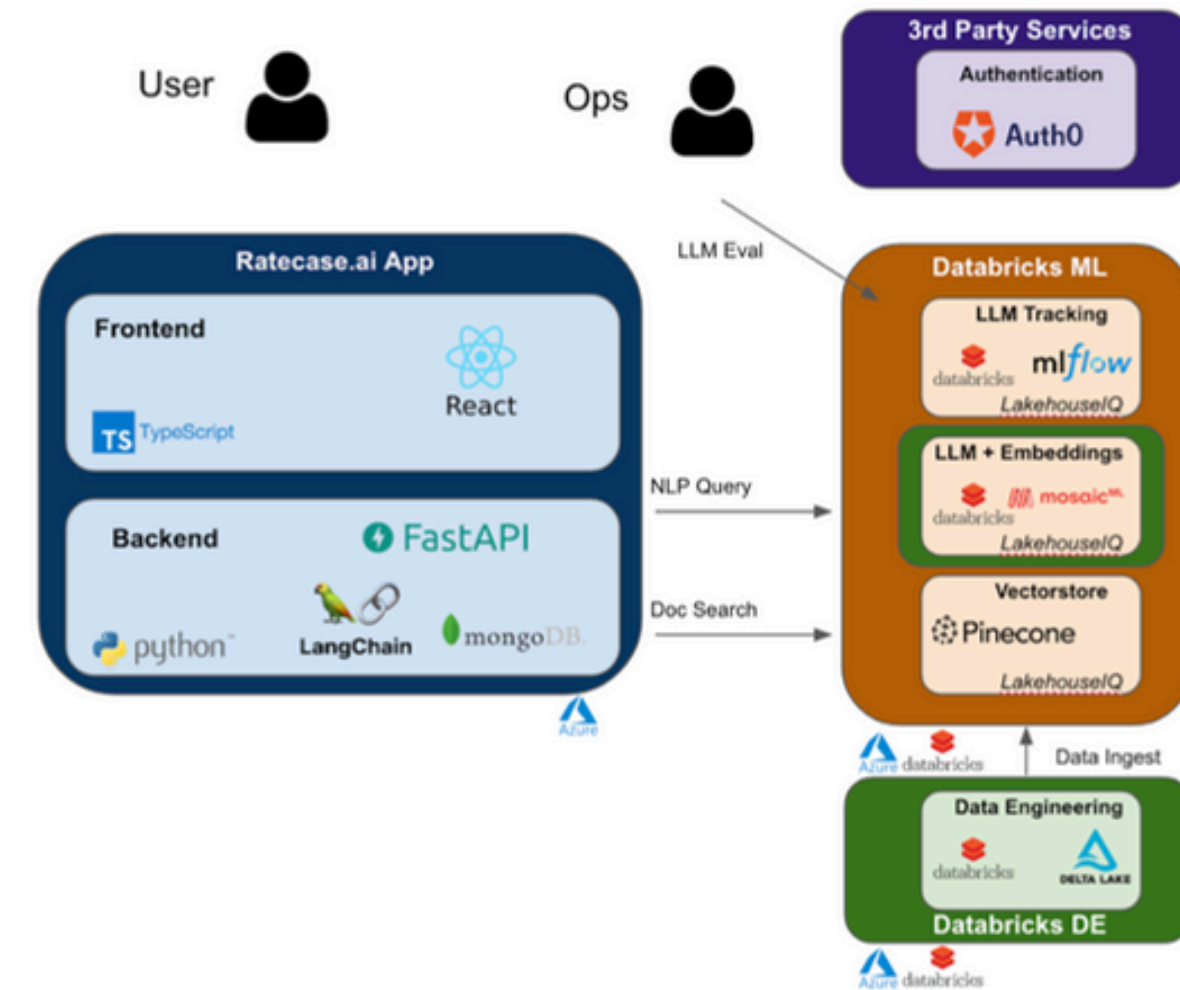
Demo: [See video demo on utilitiesGPT.com](https://utilitiesGPT.com)

Architecture:

Option 1: Hosted Platform
Current State



Option 1: Hosted Platform
Databricks Roadmap

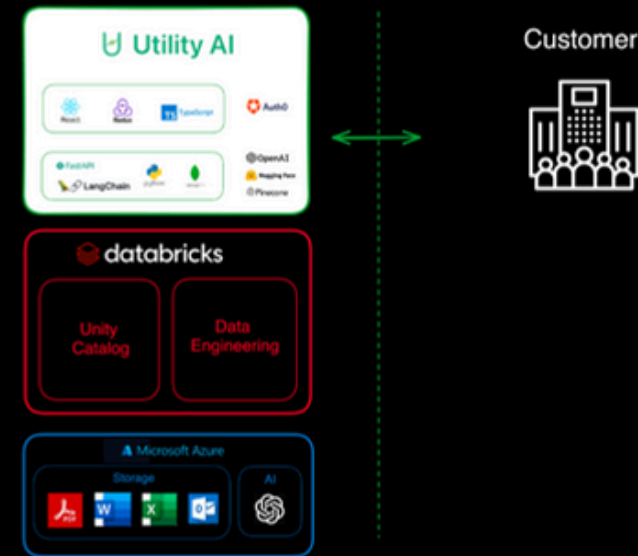


UtilitiesGPT - Rate Case AI (implementation)

Implementation Options:

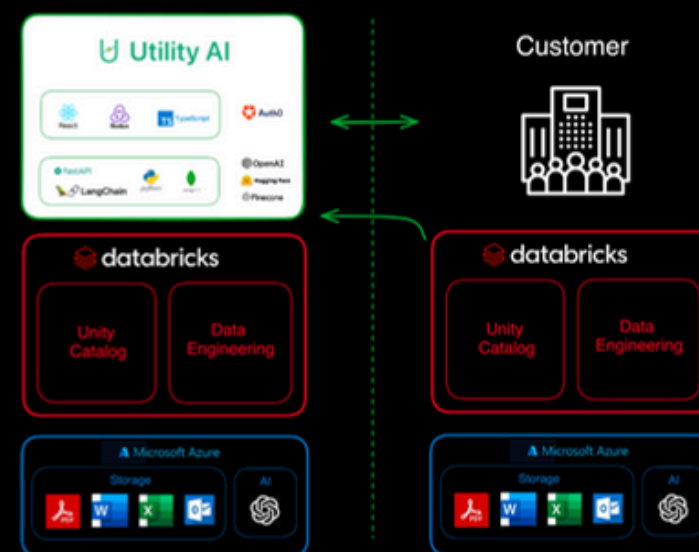
Option 1: Hosted SaaS

- Utility AI 100% hosted (i.e. Salesforce model)
- Integrate with Customer Single Sign-On (SSO)
- Customer uploads data
- All compute on Utility AI Databricks



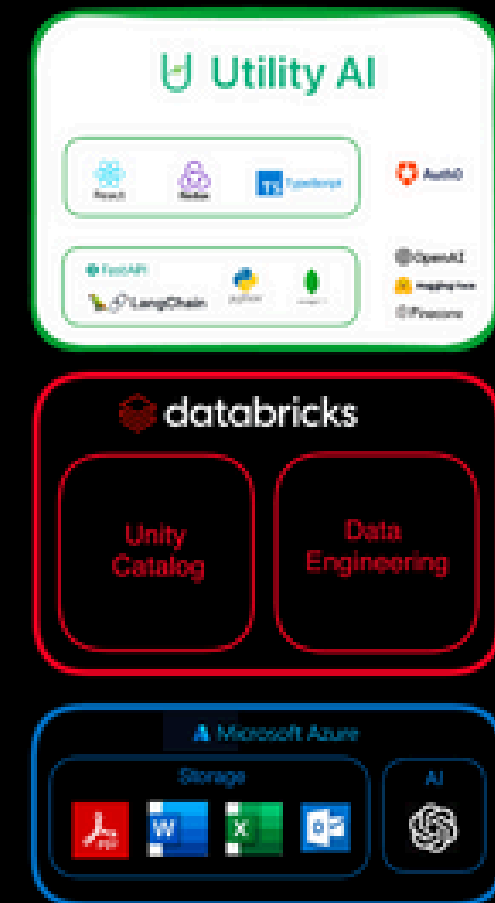
Option 2: Multi-Tenant SaaS

- Leverage Customer Cloud Provider
- Integrate with Customer Single Sign-On (SSO)
- Integrate with Customer Databricks
- Customer governs & processes their data



Option 3: In-tenant

Customer



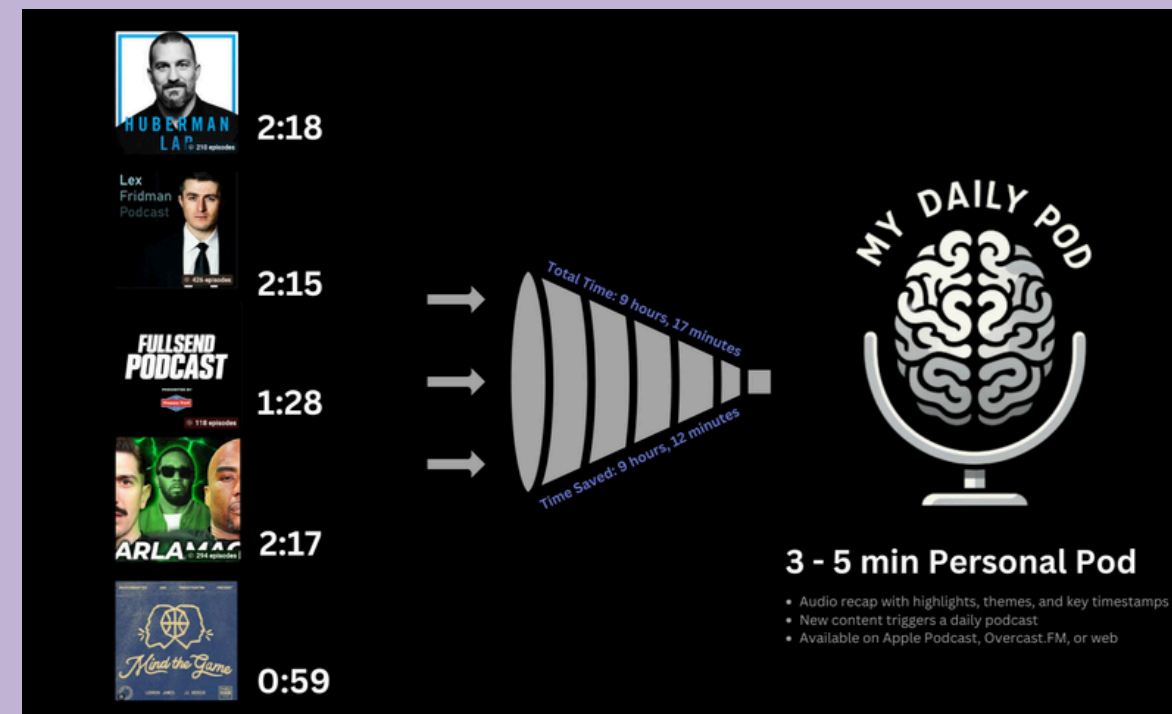
My Daily Pod - AI personalized podcasts

Situation

Not enough time to watch your favorite podcasts and YouTube content?

Solution

Our AI creates ~5min podcasts that summarize what you missed. Listen to your personalized podcast in your phone's podcast app.



Outcomes

2 to 5 minute personal podcasts delivered to your email, phone, or podcast app



My Daily Pod

Not enough time to watch your favorite YouTube content?

Our AI creates ~5min podcasts that summarize what you missed.

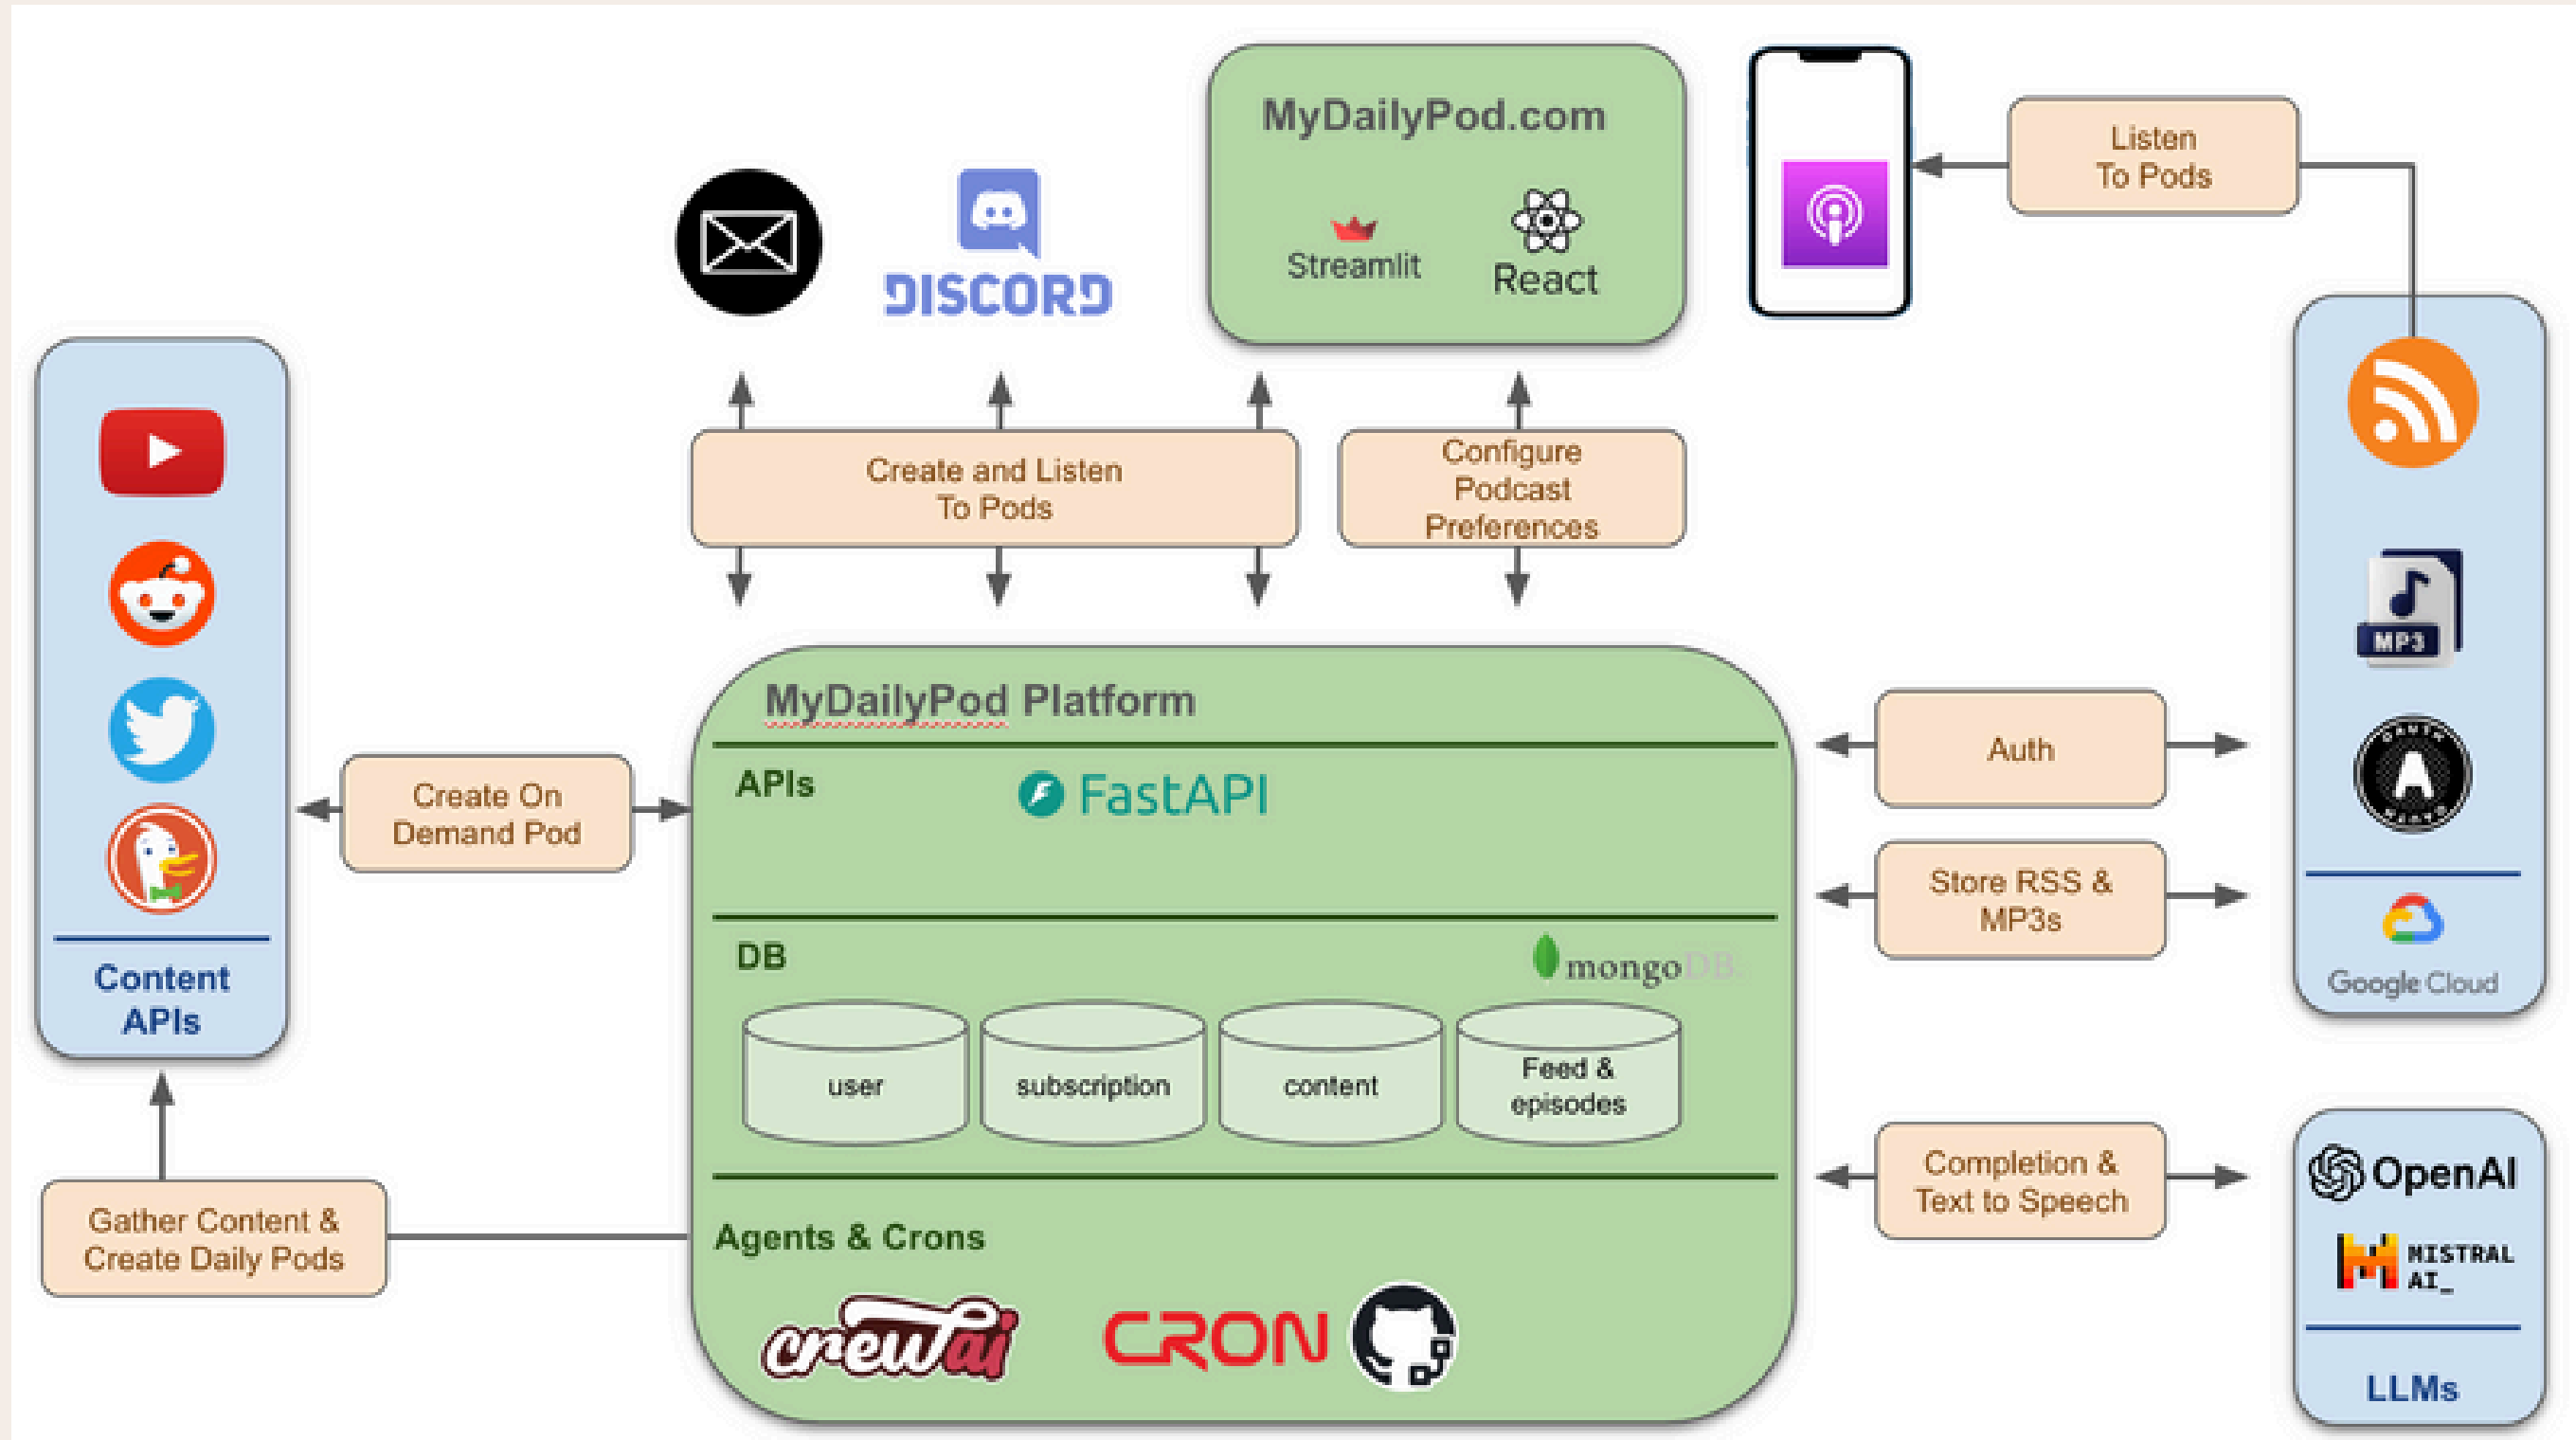
Listen to your personalized podcast in your phone's podcast app.

Try the App! ➤

My Daily Pod - AI personalized podcasts

Demo: [Video Demo](#)

Architecture:



DiscoverAde - LLM App for internal audit

Situation

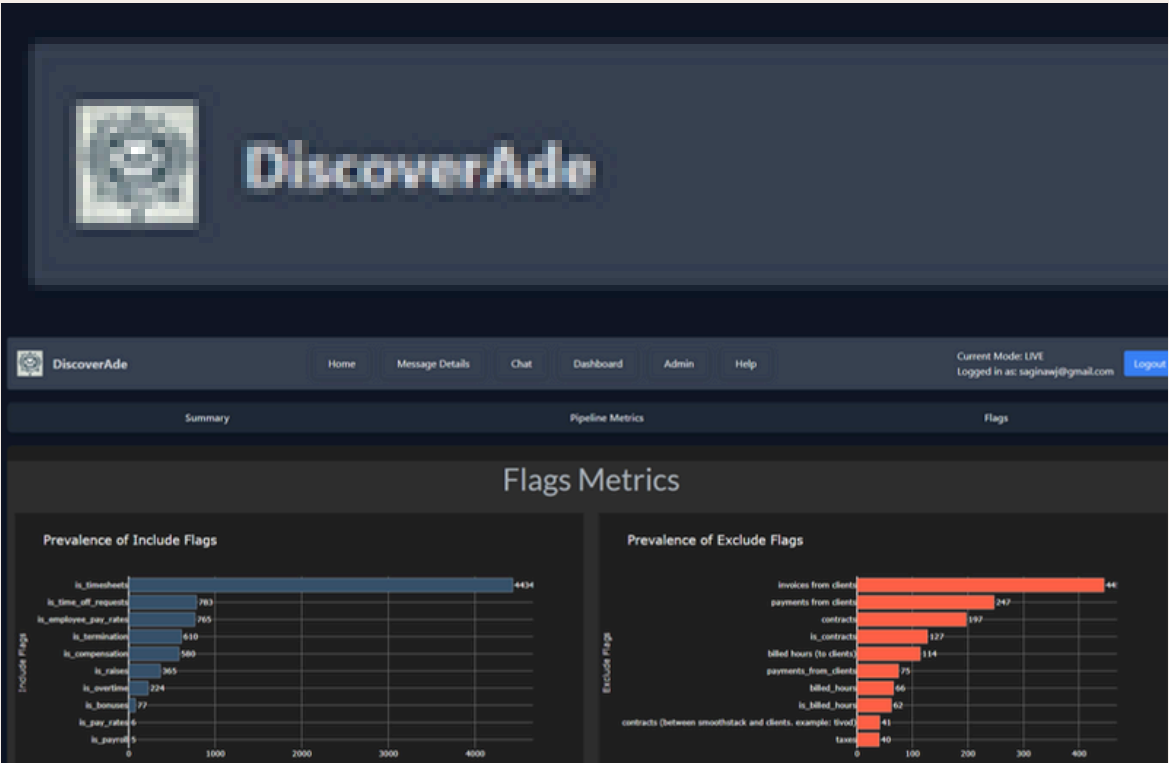
A large tech company conducting an internal audit needed to analyze millions of employee and HR emails related to paid time off, hourly rates, and bonuses. The existing manual and traditional ML pipelines lacked scalability, accuracy, and the ability to interpret nuanced email content, causing delays and missed insights into potential fraud or phishing activities.

Solution

Developed DiscoverAde, an agentic LLM app leveraging MongoDB, and Azure to streamline email analysis with a four-stage pipeline, achieving high fraud detection precision. Deployed with Docker and Railway, it provided scalable fraud detection, contextual clustering, and secure database integration.

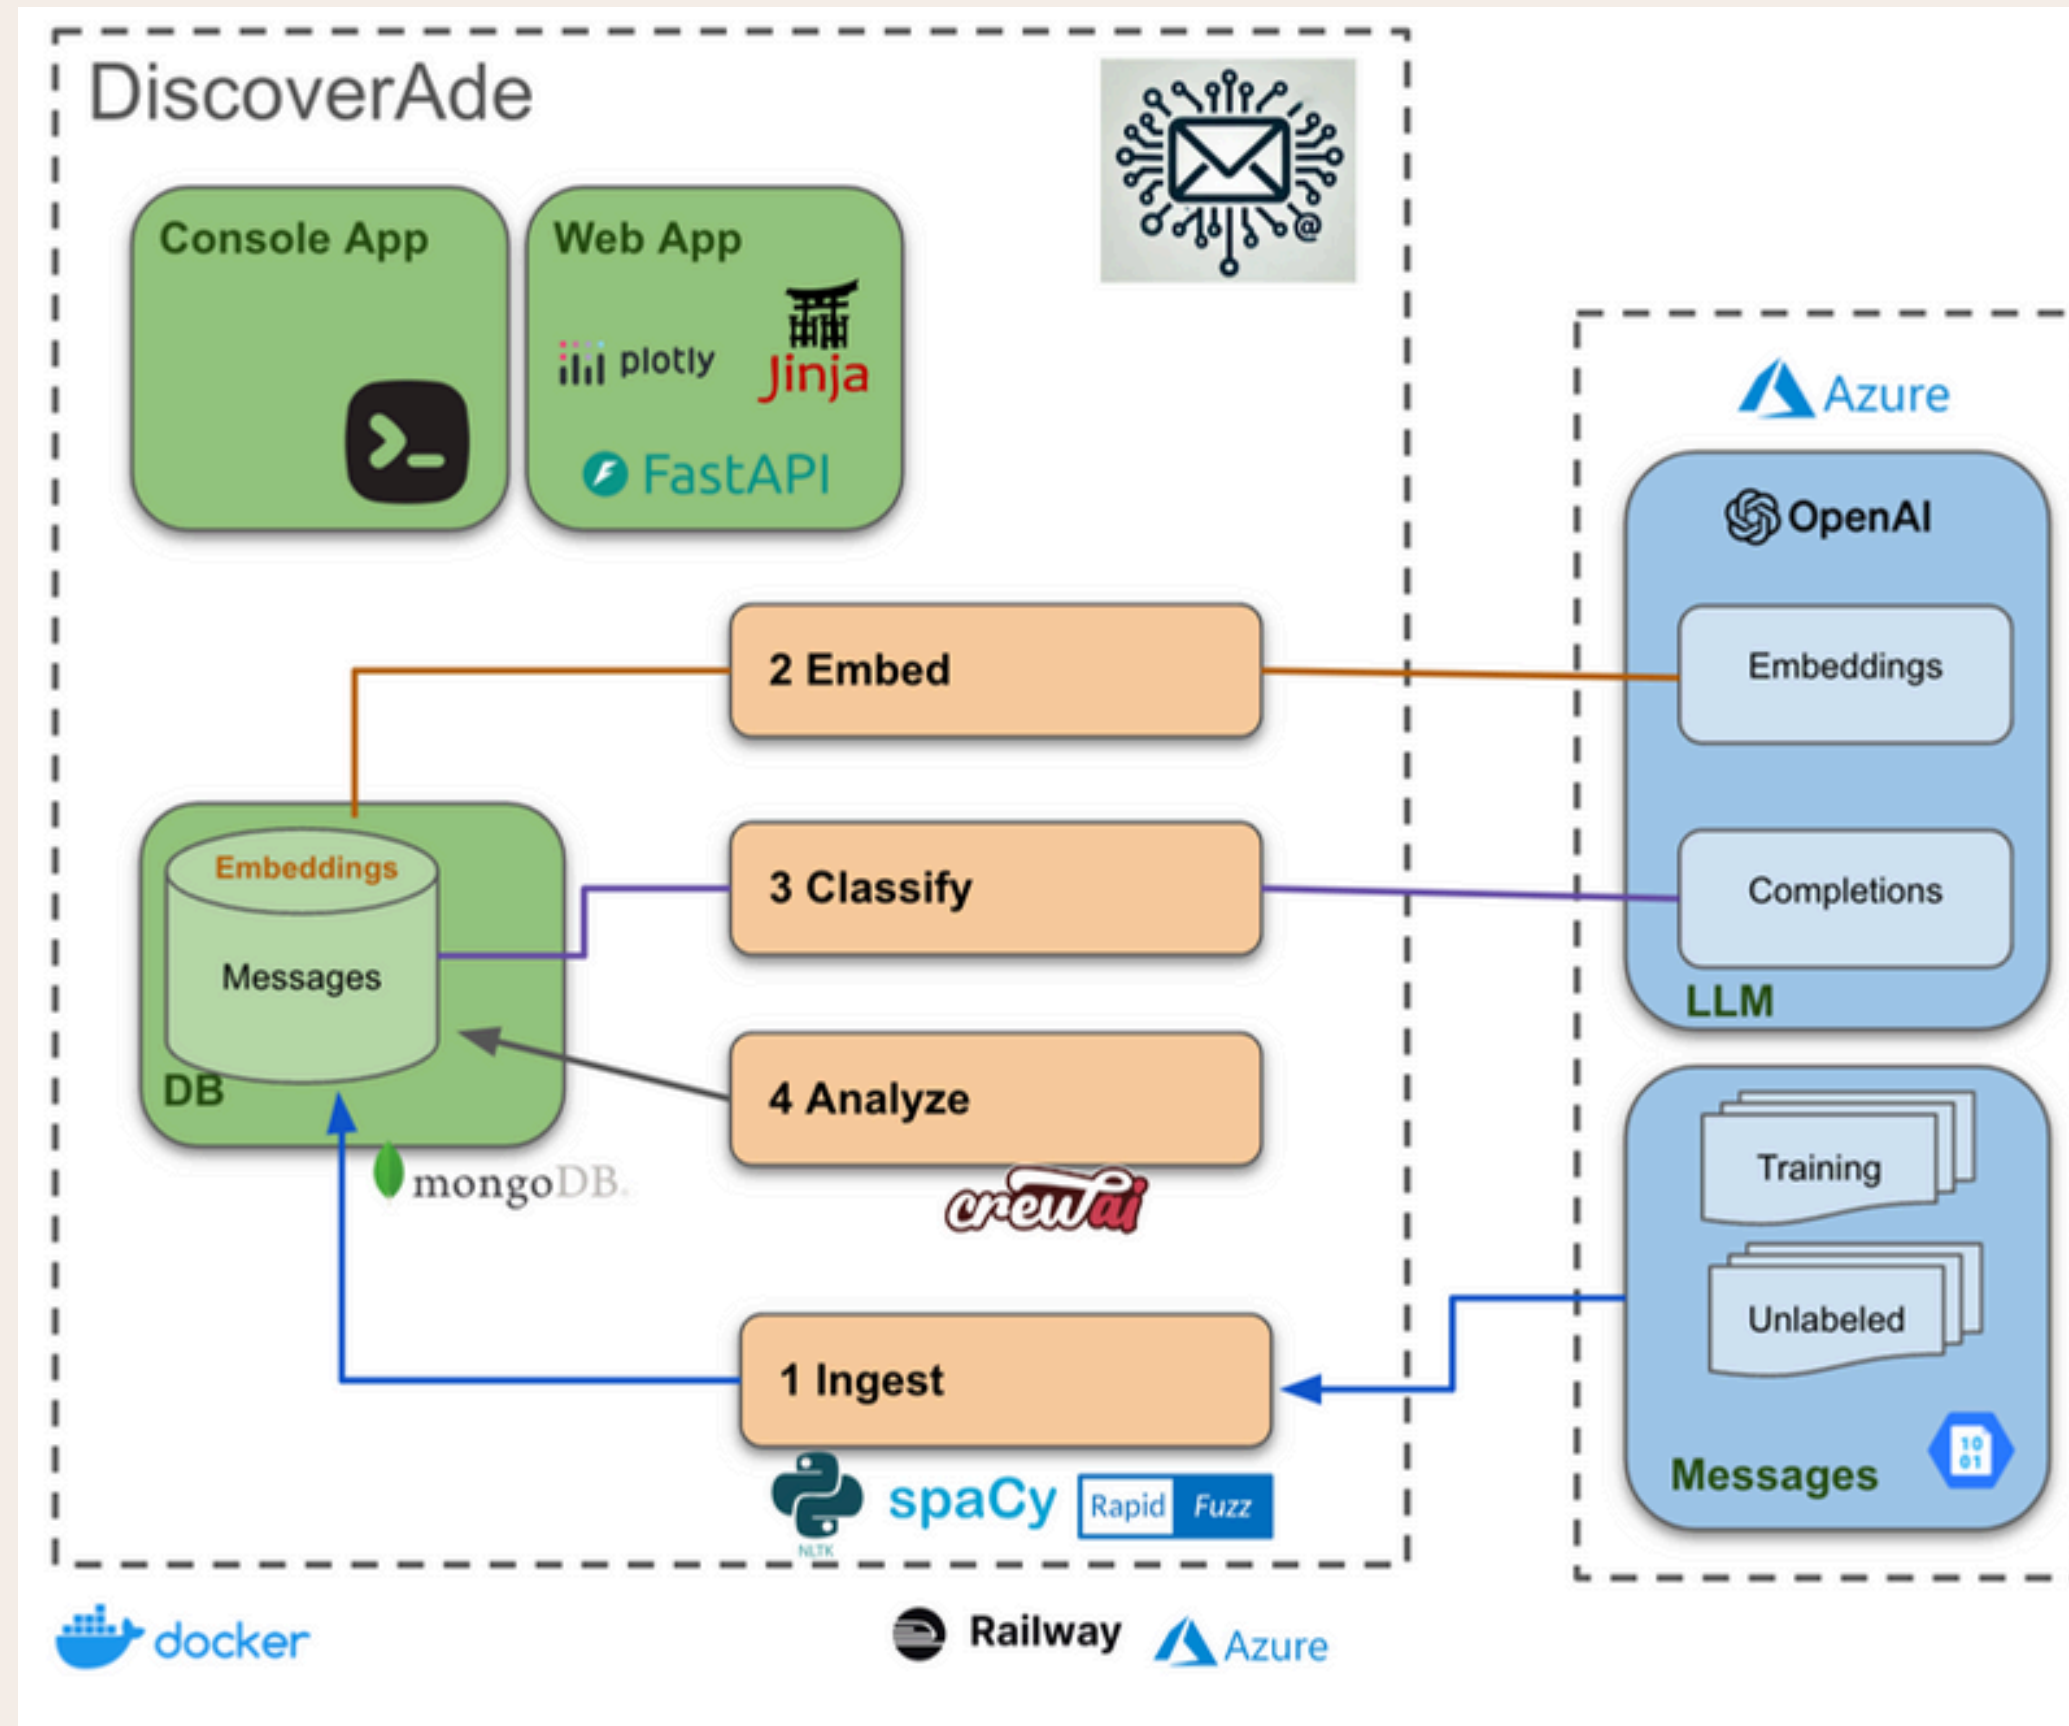
Outcomes

Developed and deployed Utility AI to Azure Marketplace, integrating with Databricks, and launched it to three utility companies. This uncovered new benchmarking insights and increased efficiency in researching historical regulatory filings.



DiscoverAde - LLM App for internal audit

Architecture:





Appendix

Additional case studies

Energy Efficiency - Sales Prospecting

Situation

GreenMarbles, a company that sells energy efficiency devices (smart thermostats, smart refrigerators, smart light bulbs), wants to research territories they can penetrate.

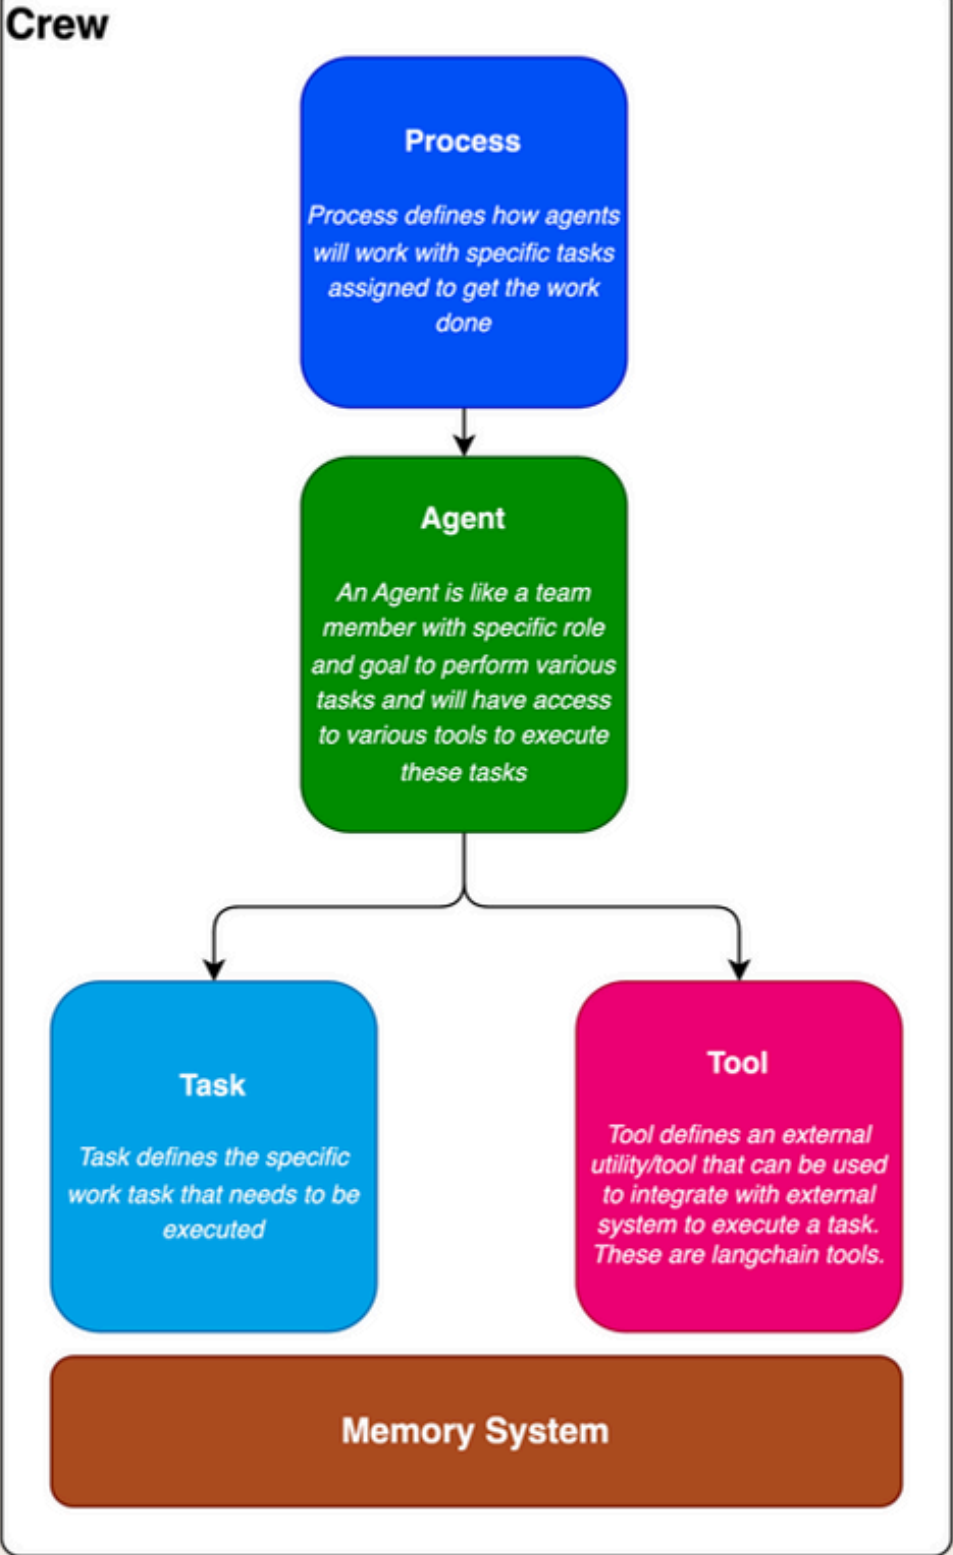
Identify which utilities offer rebates for energy efficiency devices, meet with them to sign them up for the Green Marbles marketplace and do co-marketing with them.

Solution

Developed an Agent based solution to predict the best sales opportunities

Outcomes

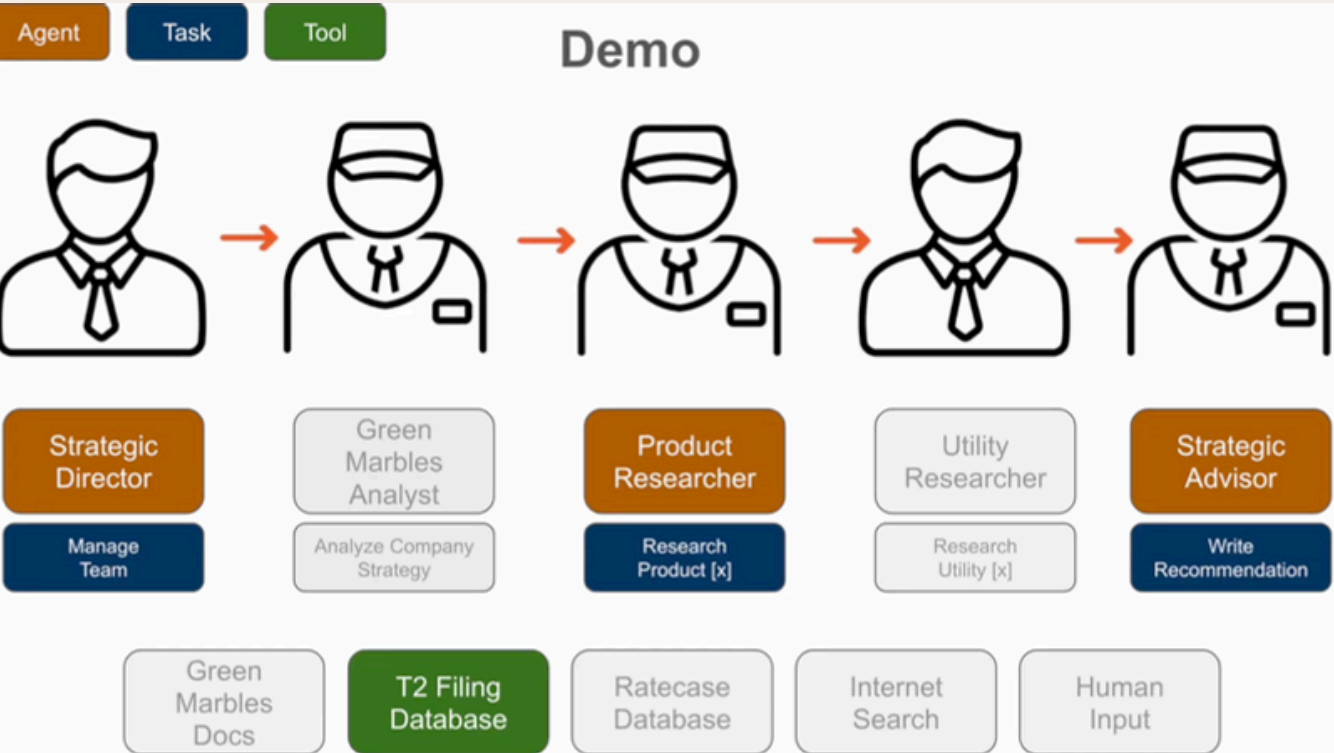
Developed and deployed use case for GreenMarbles, which helped them direct their very lean sales team toward opportunities in Massachusetts, ultimately enabling their go-to-market approach in the Northeast.



Energy Efficiency - Sales Prospecting (technical)

Demo: Sales Prospecting Demo

Architecture:



Example: Smart Thermostat to JCP&L

I'm looking to pitch Smart Thermostat initiatives to JCP&L. Based on how their Programs are structured in the T2 filing, is this a good idea? If so, please provide details on the specific programs where Smart Thermostat initiatives are mentioned, along with any expectations or specifics from JCP&L

Based on the details from JCP&L's Energy Efficiency and Conservation Plan (EE&C Plan II), pitching Smart Thermostat initiatives to JCP&L could be a good idea, as smart thermostats are mentioned in their Load Optimization & Peak Demand Reduction (PDR) program. This program aims to target energy and peak demand savings through Behavioral and Load Control components, with Load Control specifically targeting customer-connected devices, including smart thermostats. This component is designed to achieve energy and peak demand savings by managing these devices .

JCP&L's plan includes a comprehensive portfolio of energy efficiency (EE), PDR, and building decarbonization (BD) programs for various customer sectors. Smart thermostats, as part of the Load Optimization & PDR program, are aligned with JCP&L's objectives to increase energy efficiency and reduce peak demand across its customer base. The program is designed to be inclusive, offering opportunities for all customers, including residential, multifamily, and commercial and industrial sectors .

Therefore, proposing Smart Thermostat initiatives to JCP&L aligns well with their strategic direction in EE&C Plan II, particularly under the Load Optimization & PDR program. The initiative should focus on energy savings, customer engagement, and peak demand reduction, which are core aspects of JCP&L's plan.

Weather Impact ML Model

Situation

A set of meteorologists were tasked with predicting the number of power outages based on 20 years of historical outage data and 20 years of historical weather data. This prediction is critical, as it informs how Utilities prepare crews for responding to power outages.

Solution

Developed an ML regression model to predict weather based on 50+ weather parameters. Developed front-end application to perform real-time scenario planning to allow crew managers to forecast different potential outage outcomes.

Outcomes

Model had R2 value of .86 and was the first model productionized as an Enterprise application, transforming the way operations planned for outages. This had huge benefits taking ML from a back office activity, to embedding it into the business operations.

