



Improving AI accuracy with AI-ready unstructured and structured data

Hilton Rosenfeld
IBM, Senior Solution Engineer

13 October 2025
Perth



AI gives us real-time insights from all our data.


Amazing!
What data?



TOM
FISH
BURNE


Accuracy is the new Currency

Enterprise Data



Enterprise Data

90%
unstructured¹




Enterprise Data

90%

unstructured¹

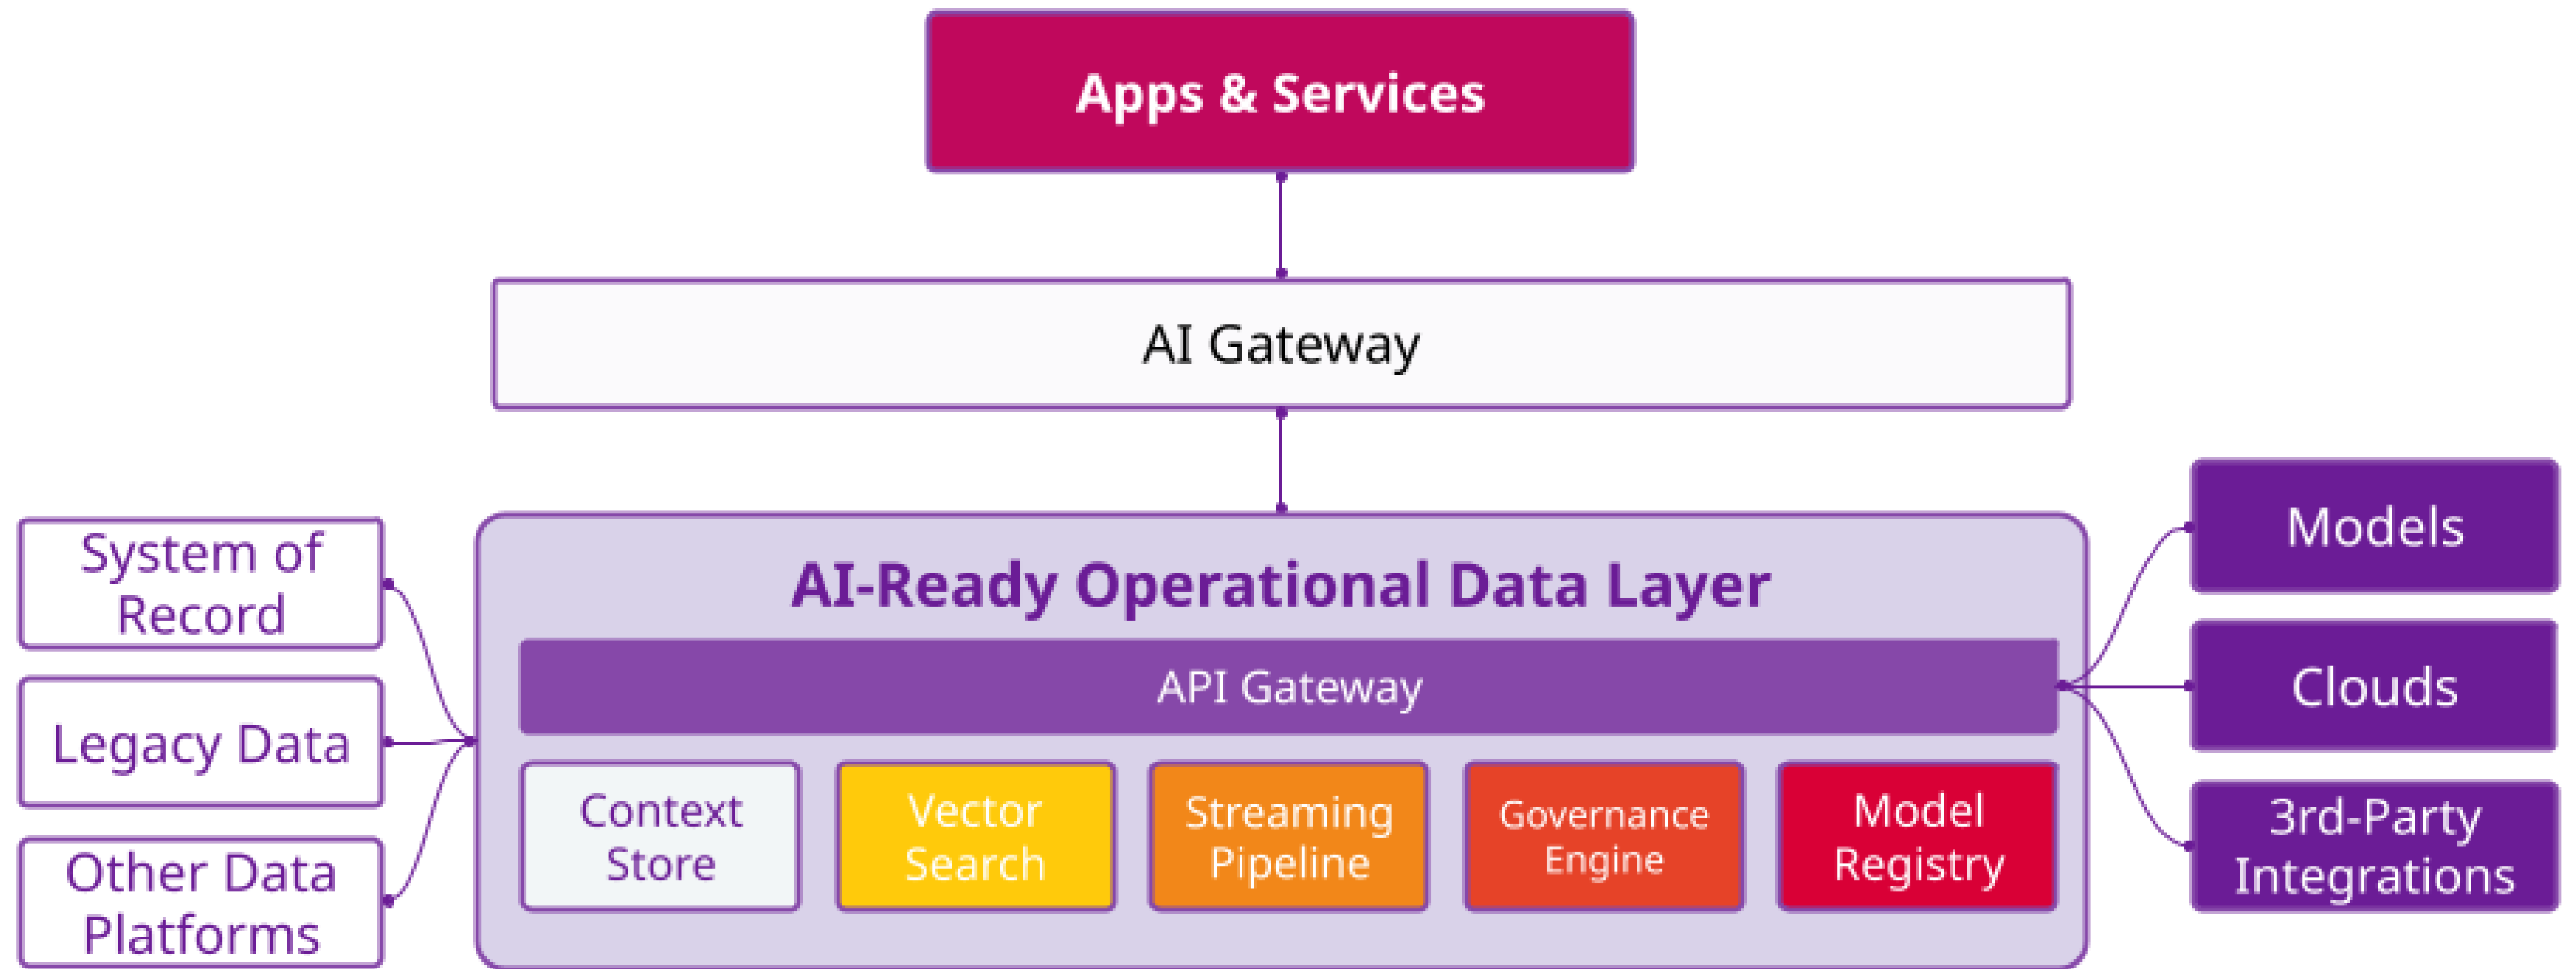
1%

currently being
leveraged by gen AI²

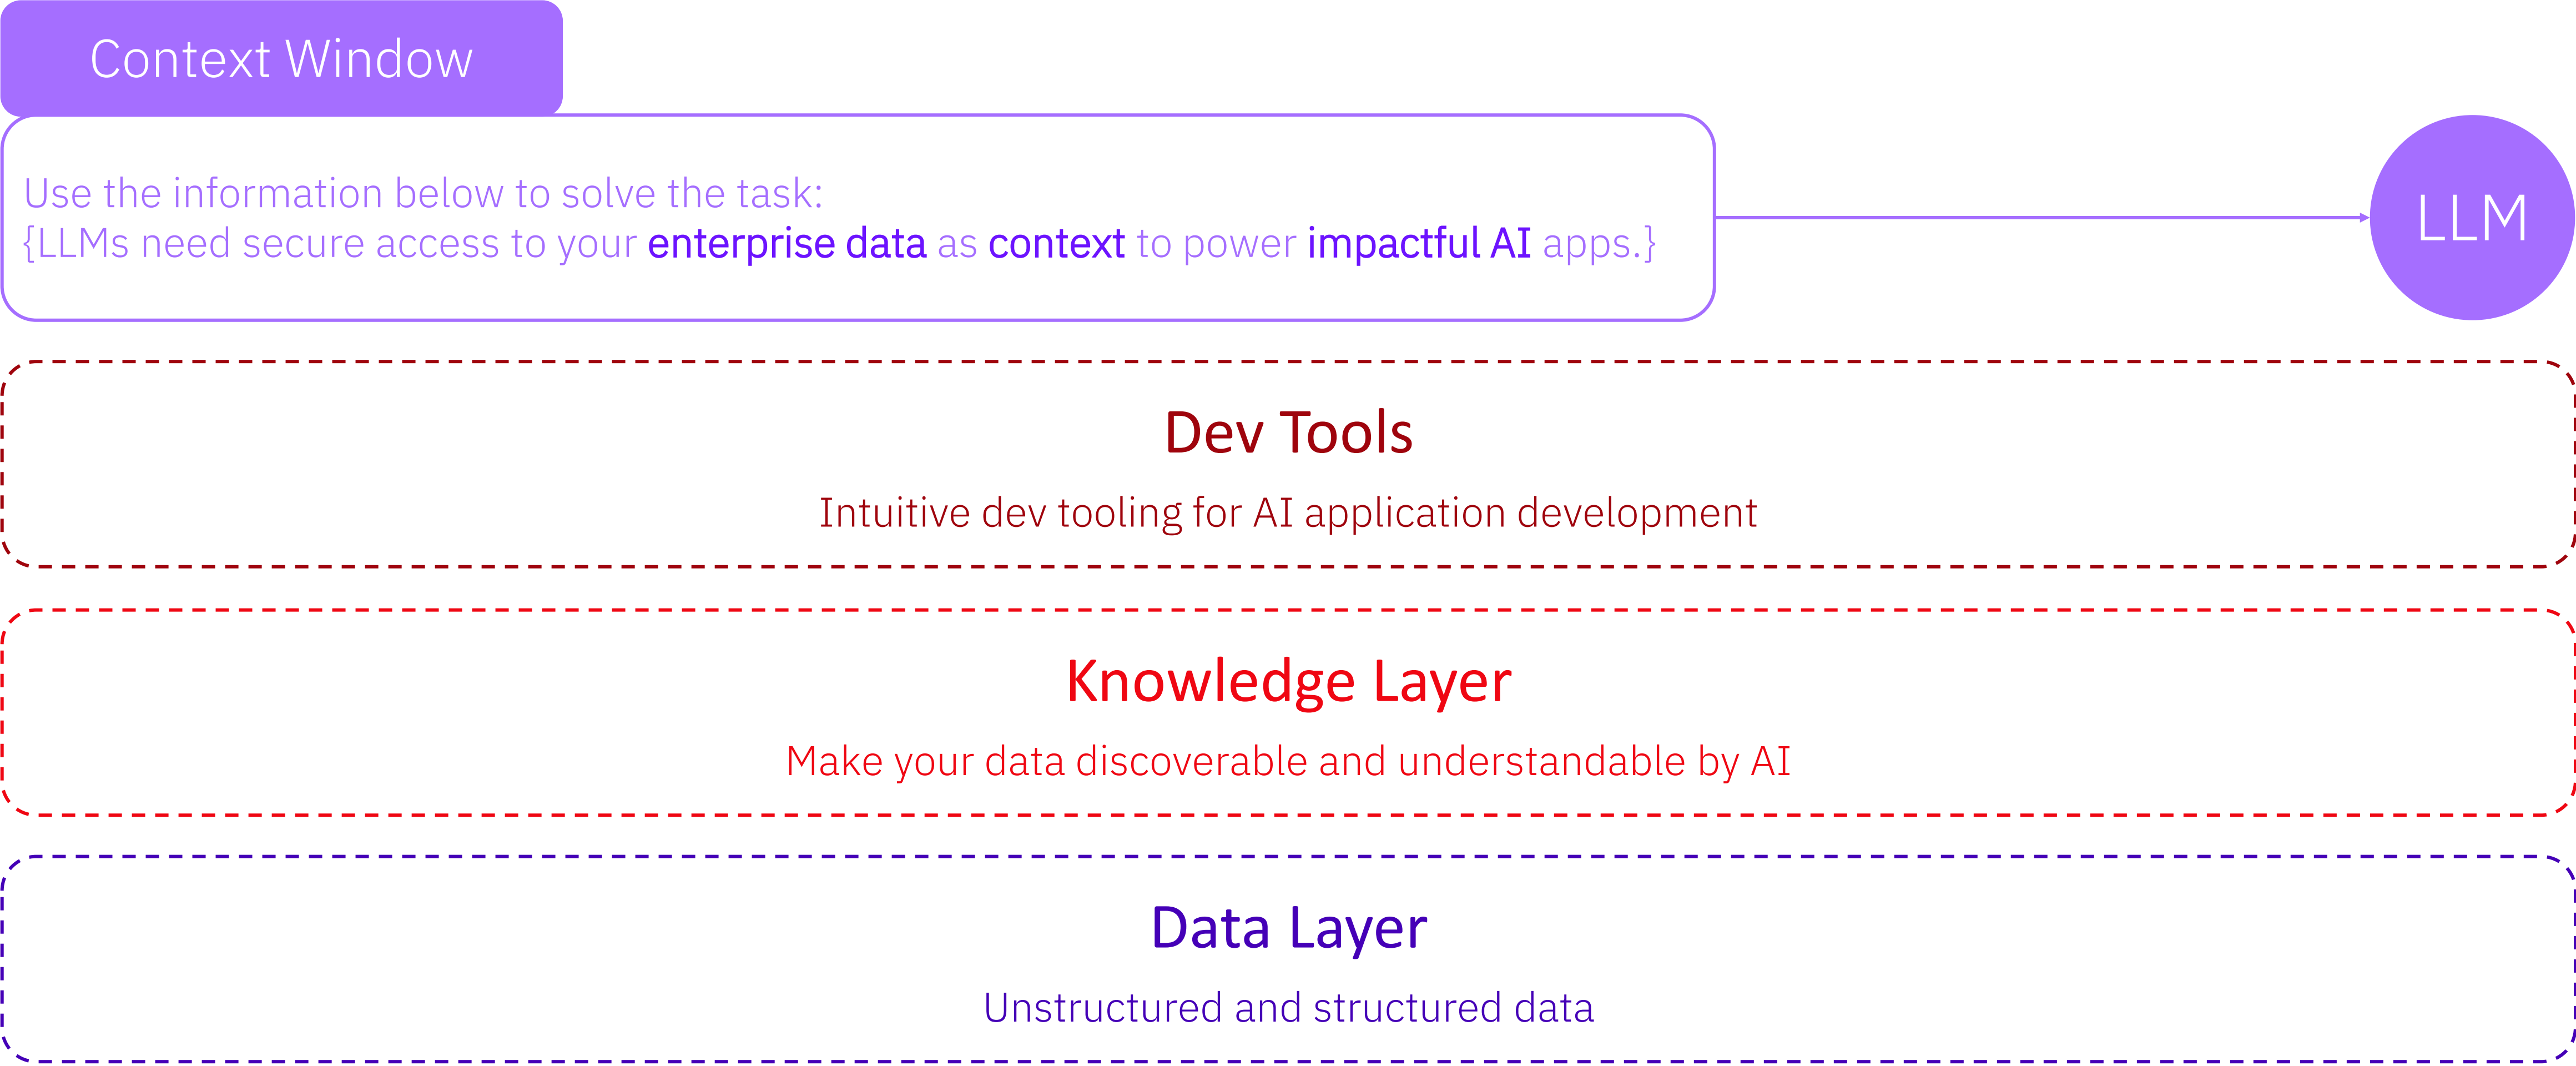


How do you unlock
AI-ready unstructured
and structured data
for your agents?

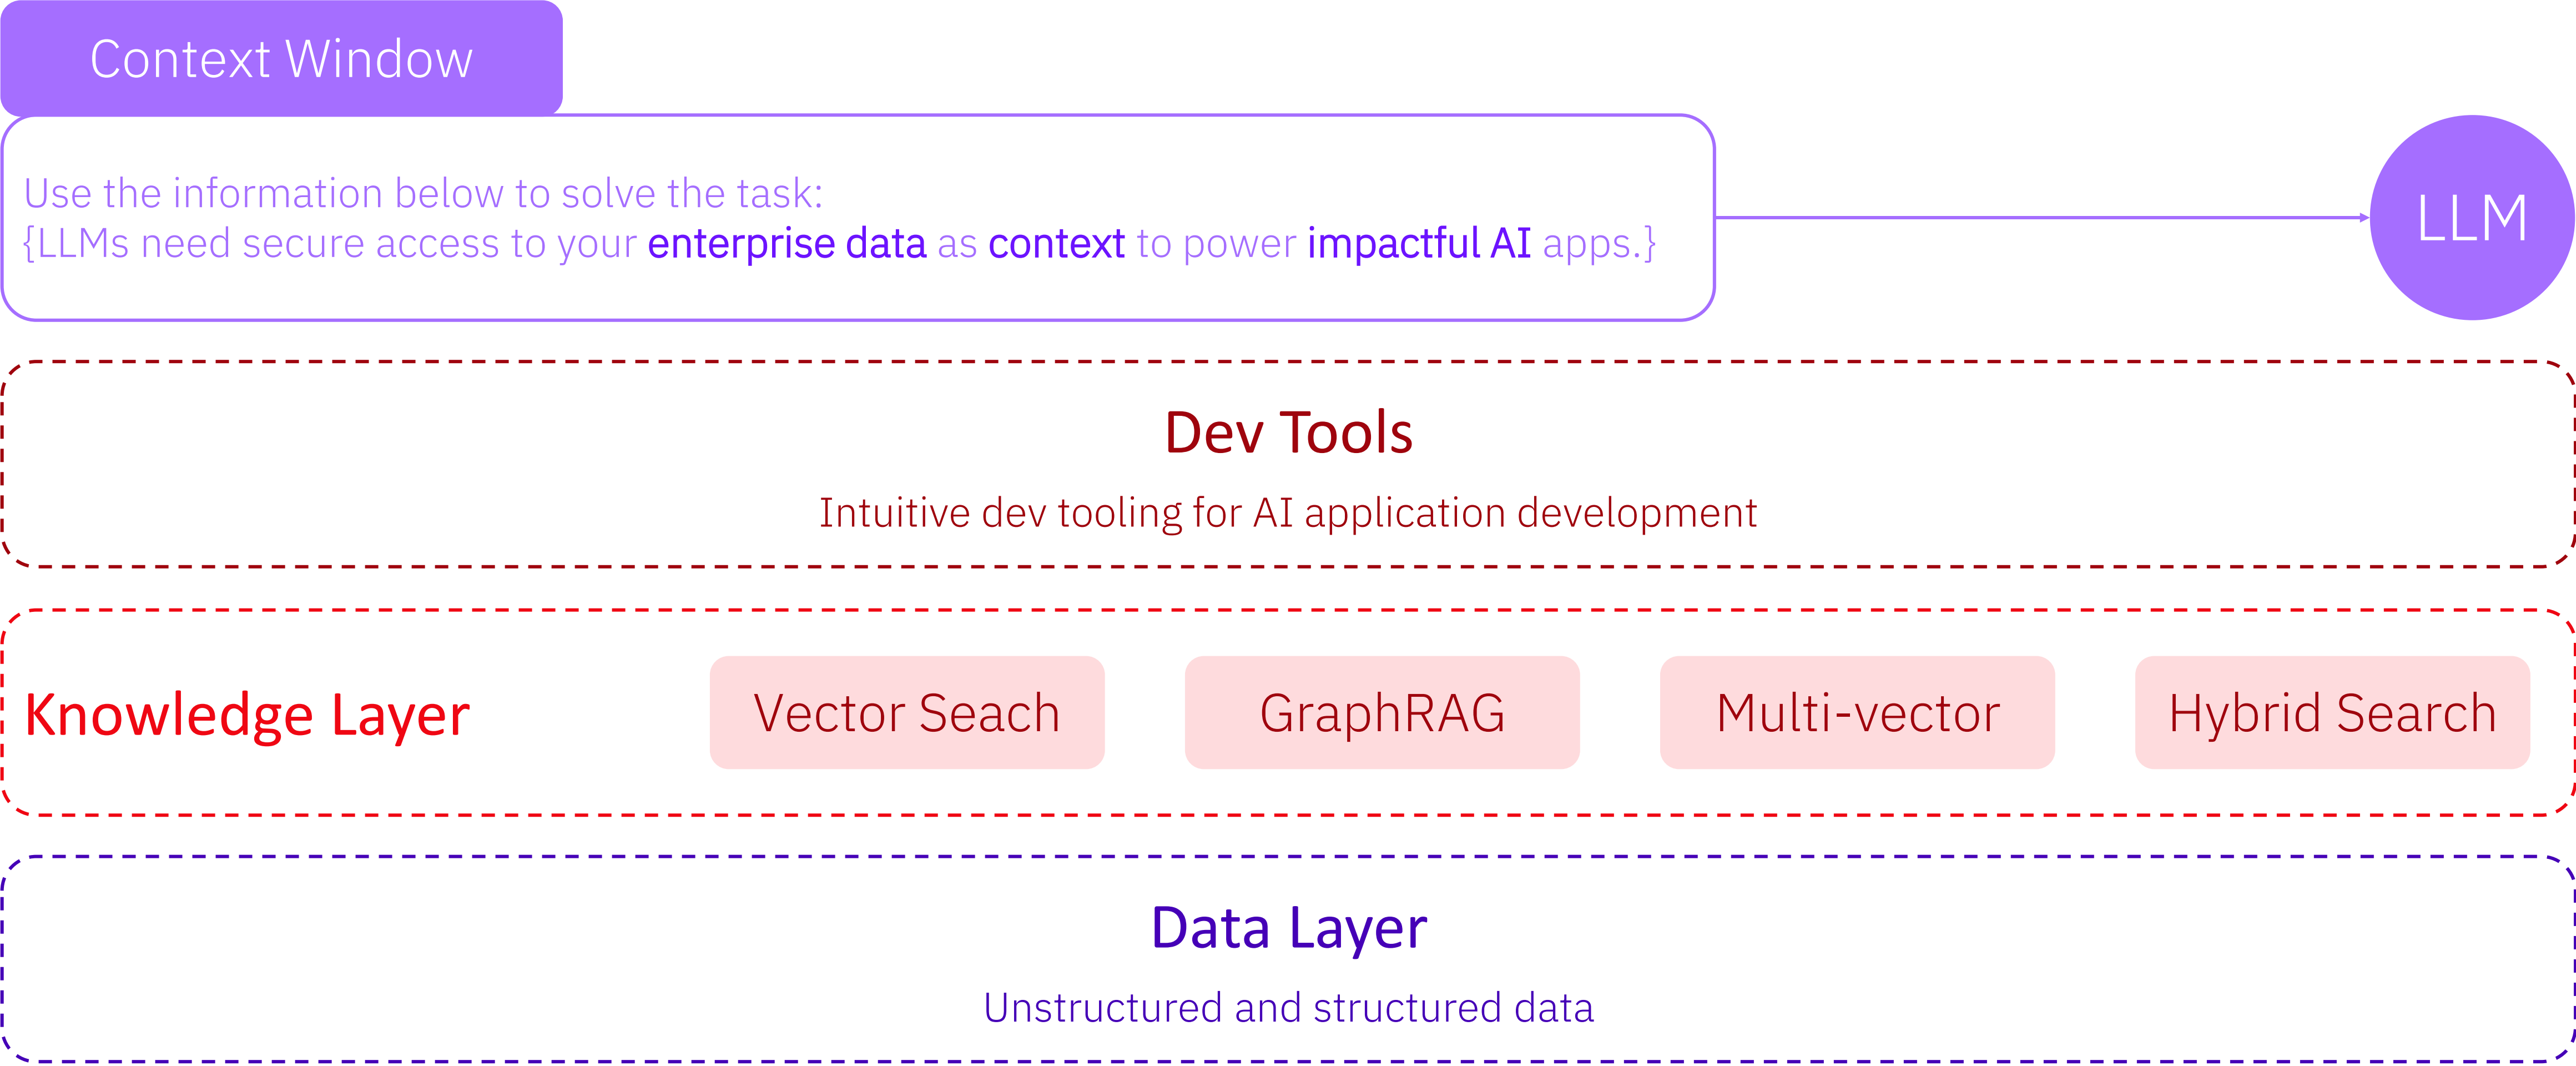
Start with an AI-ready operational data layer



Accurate AI Requires a Knowledge Layer



Accurate AI Requires a Knowledge Layer



Real-World Use Cases

Ahead of the Curve in Multiple Clouds

Macquarie Counts on DataStax

Industry: Banking
Customer Since: 2017

“In 2017, we started to use DataStax Enterprise.

We started using a smart data store that is middleware that is running on the digital platform. That middleware is receiving all the customer info from all of our legacy systems and all interactions. Having that approach gives you a sense of how important the product is.

For the first time in the history of the bank, we can provide a real 360-degree view of customers.”

Luis Uguina, Chief Data Officer



1M+

Customers served on the digital banking experience

<2 yrs

to transform from no retail presence to digital banking juggernaut

100%

Uptime

Skypoint Delivers GenAI-Driven Efficiency Revolution in Healthcare with Astra DB and Vector Search

Industry: Software

Customer Since: 2023


“Skypoint Enterprise AI is bringing about a significant transformation in the senior living healthcare industry, valued at \$100B annually and currently burdened with nearly 70% operational and labor costs.

This transformation is achieved by leveraging DataStax's Astra DB vector database features on Microsoft Azure. Employing generative AI and columnar data lakehouse technology, Skypoint AI ensures seamless access to resident health data and administrative insights. Envision it as a ChatGPT equivalent for senior living enterprise data, maintaining full HIPAA compliance, and significantly improving healthcare for the elderly.”


Tisson Mathew, Founder & CEO, SkyPoint



No Ops, Reduce TCO



High Reliability with Scalability



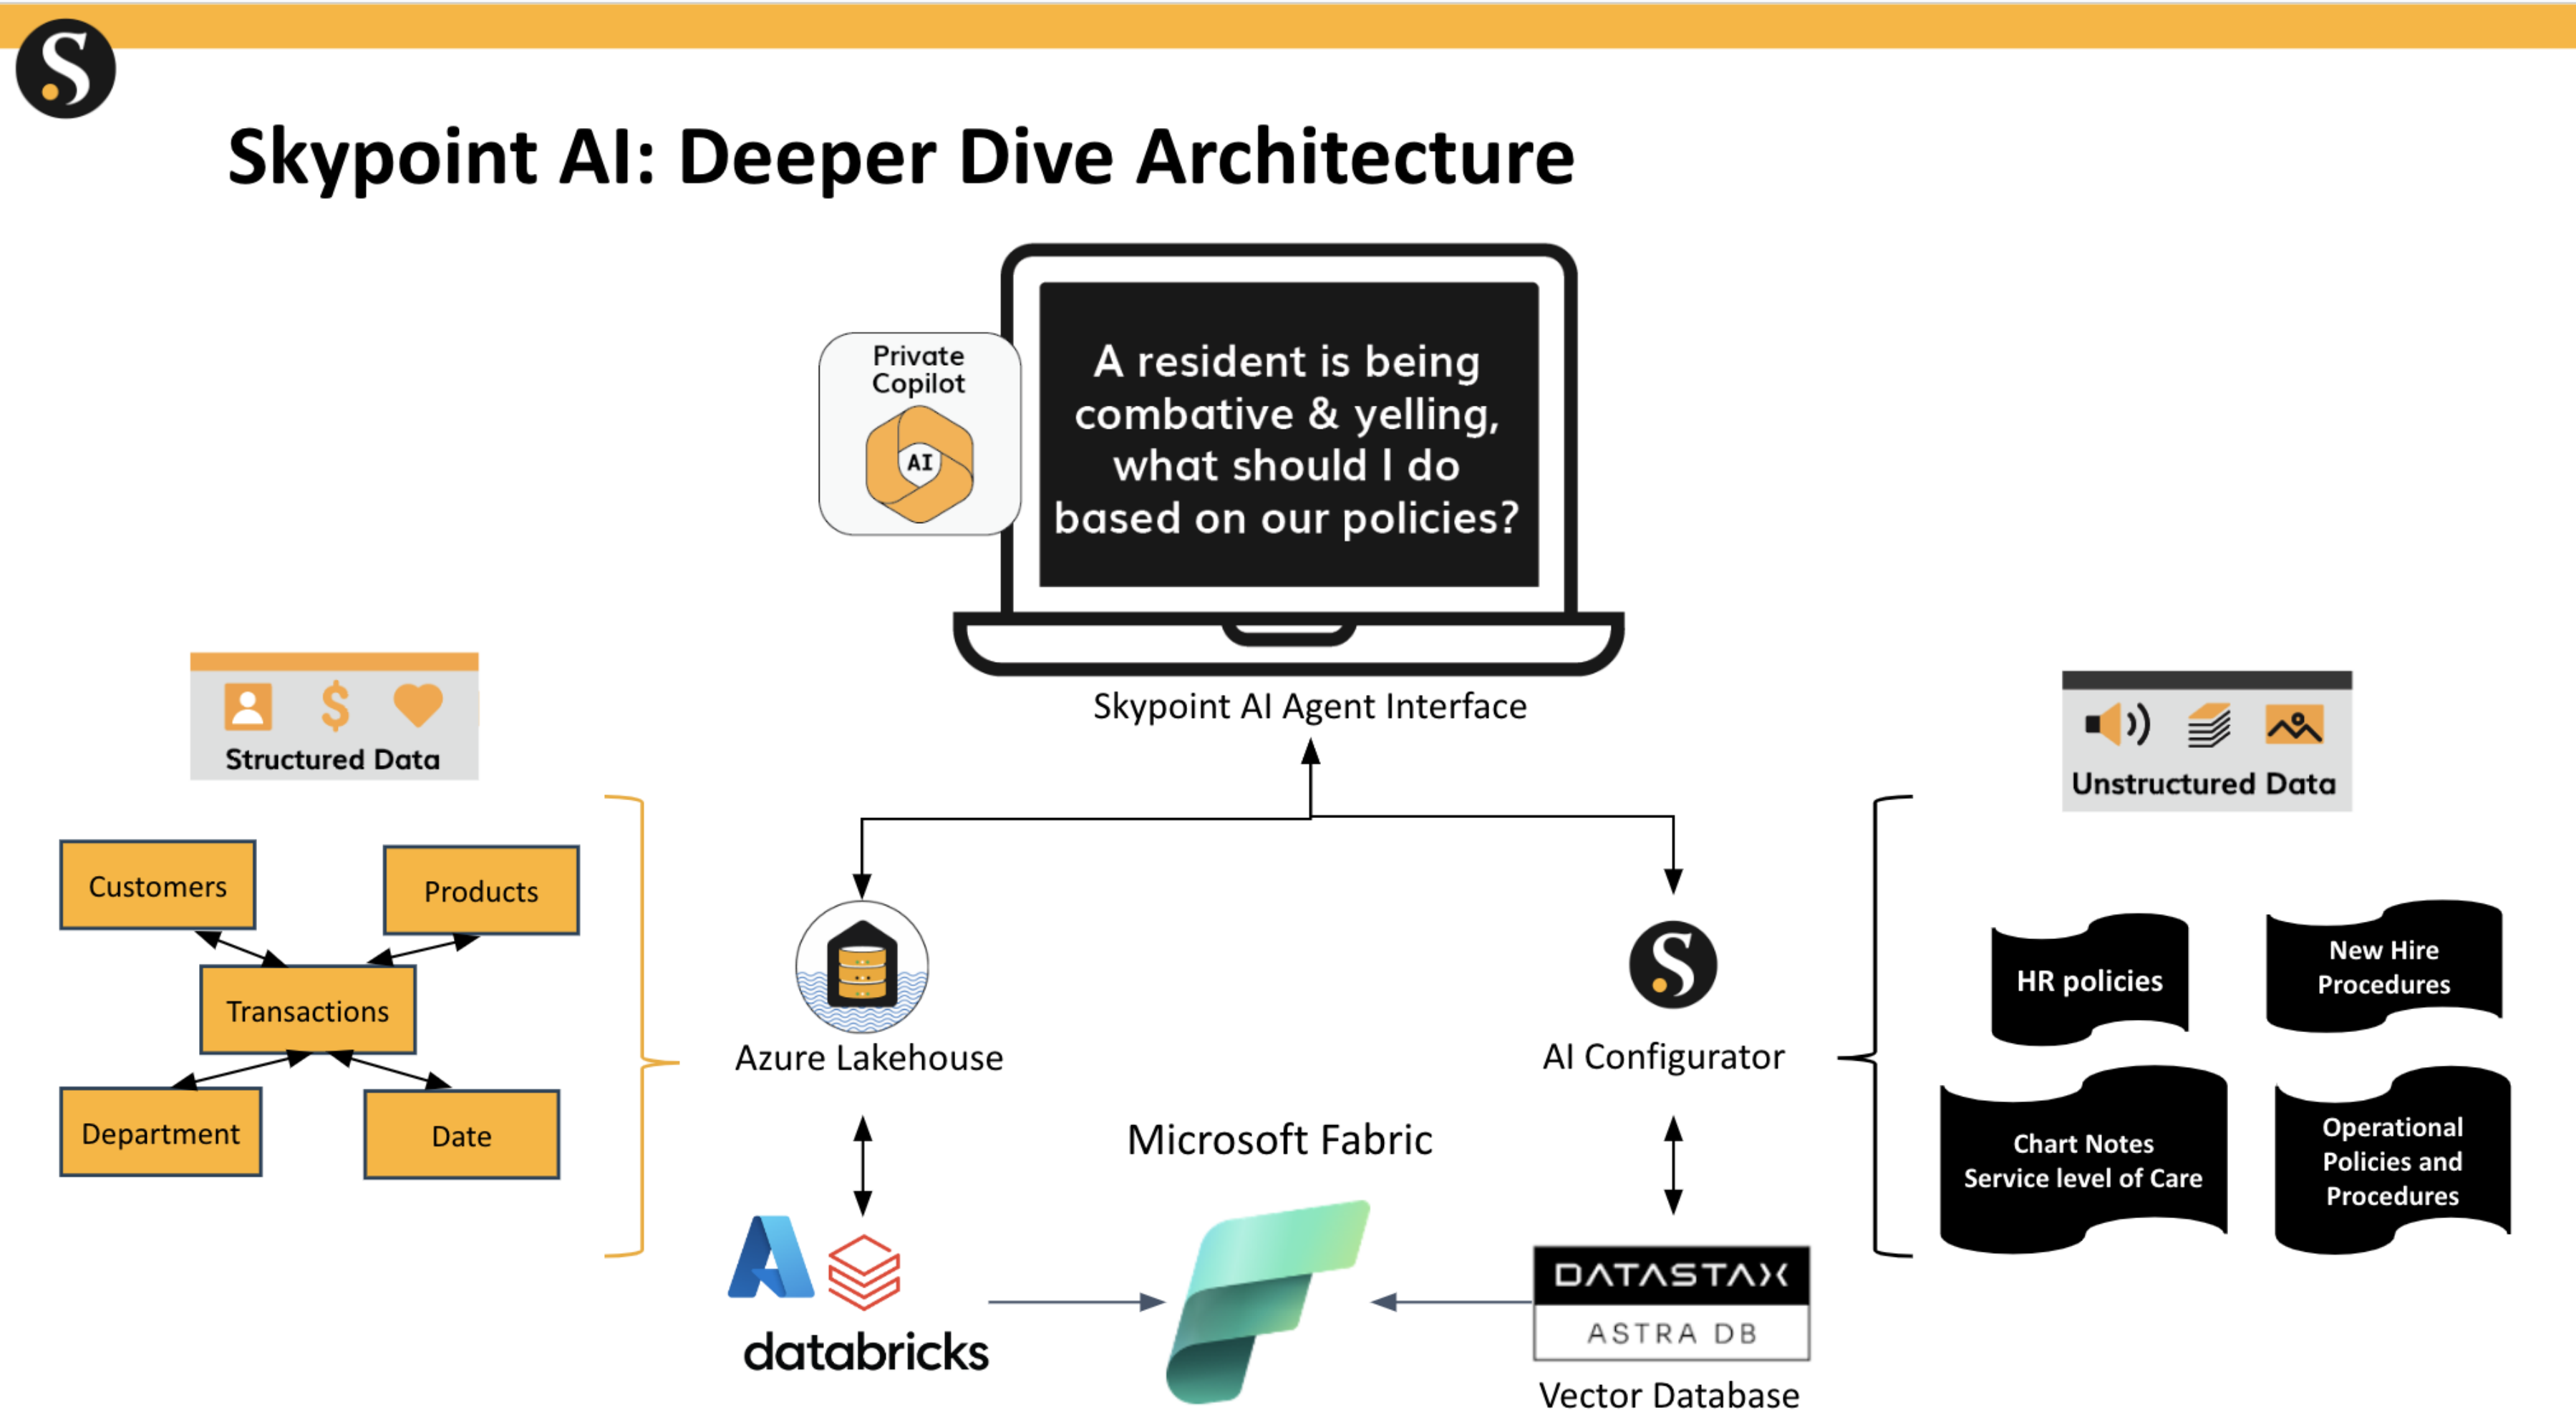
Speed

Combination of structured and unstructured data.

Structured data (elderly care information and assisted living ratings, for example) is stored in a lakehouse.

Unstructured data (such as policy PDFs and legal documents) are stored in Astra DB with vector search.

*“Reduced answer retrieval time from **hours to seconds** using DataStax and AI.”*



How Physics Wallah Uses Vector Search to Revolutionise Education

Industry: Software

Customer Since: 2023

“At PhysicsWallah, we're dedicated to delivering high-quality and affordable education.


We built a generative AI-driven chatbot powered by the Astra DB vector database and LangChain to be a one-stop solution for every student's learning needs.

We employ Astra DB’s semantic search for advanced support queries, enriching our students’ learning experience, and RAGStack facilitates seamless deployment of RAG-based applications.”


Sandeep Penmetsa, head of data science and engineering, PhysicsWallah.



No Ops, Reduce TCO



High Reliability with Scalability



Speed

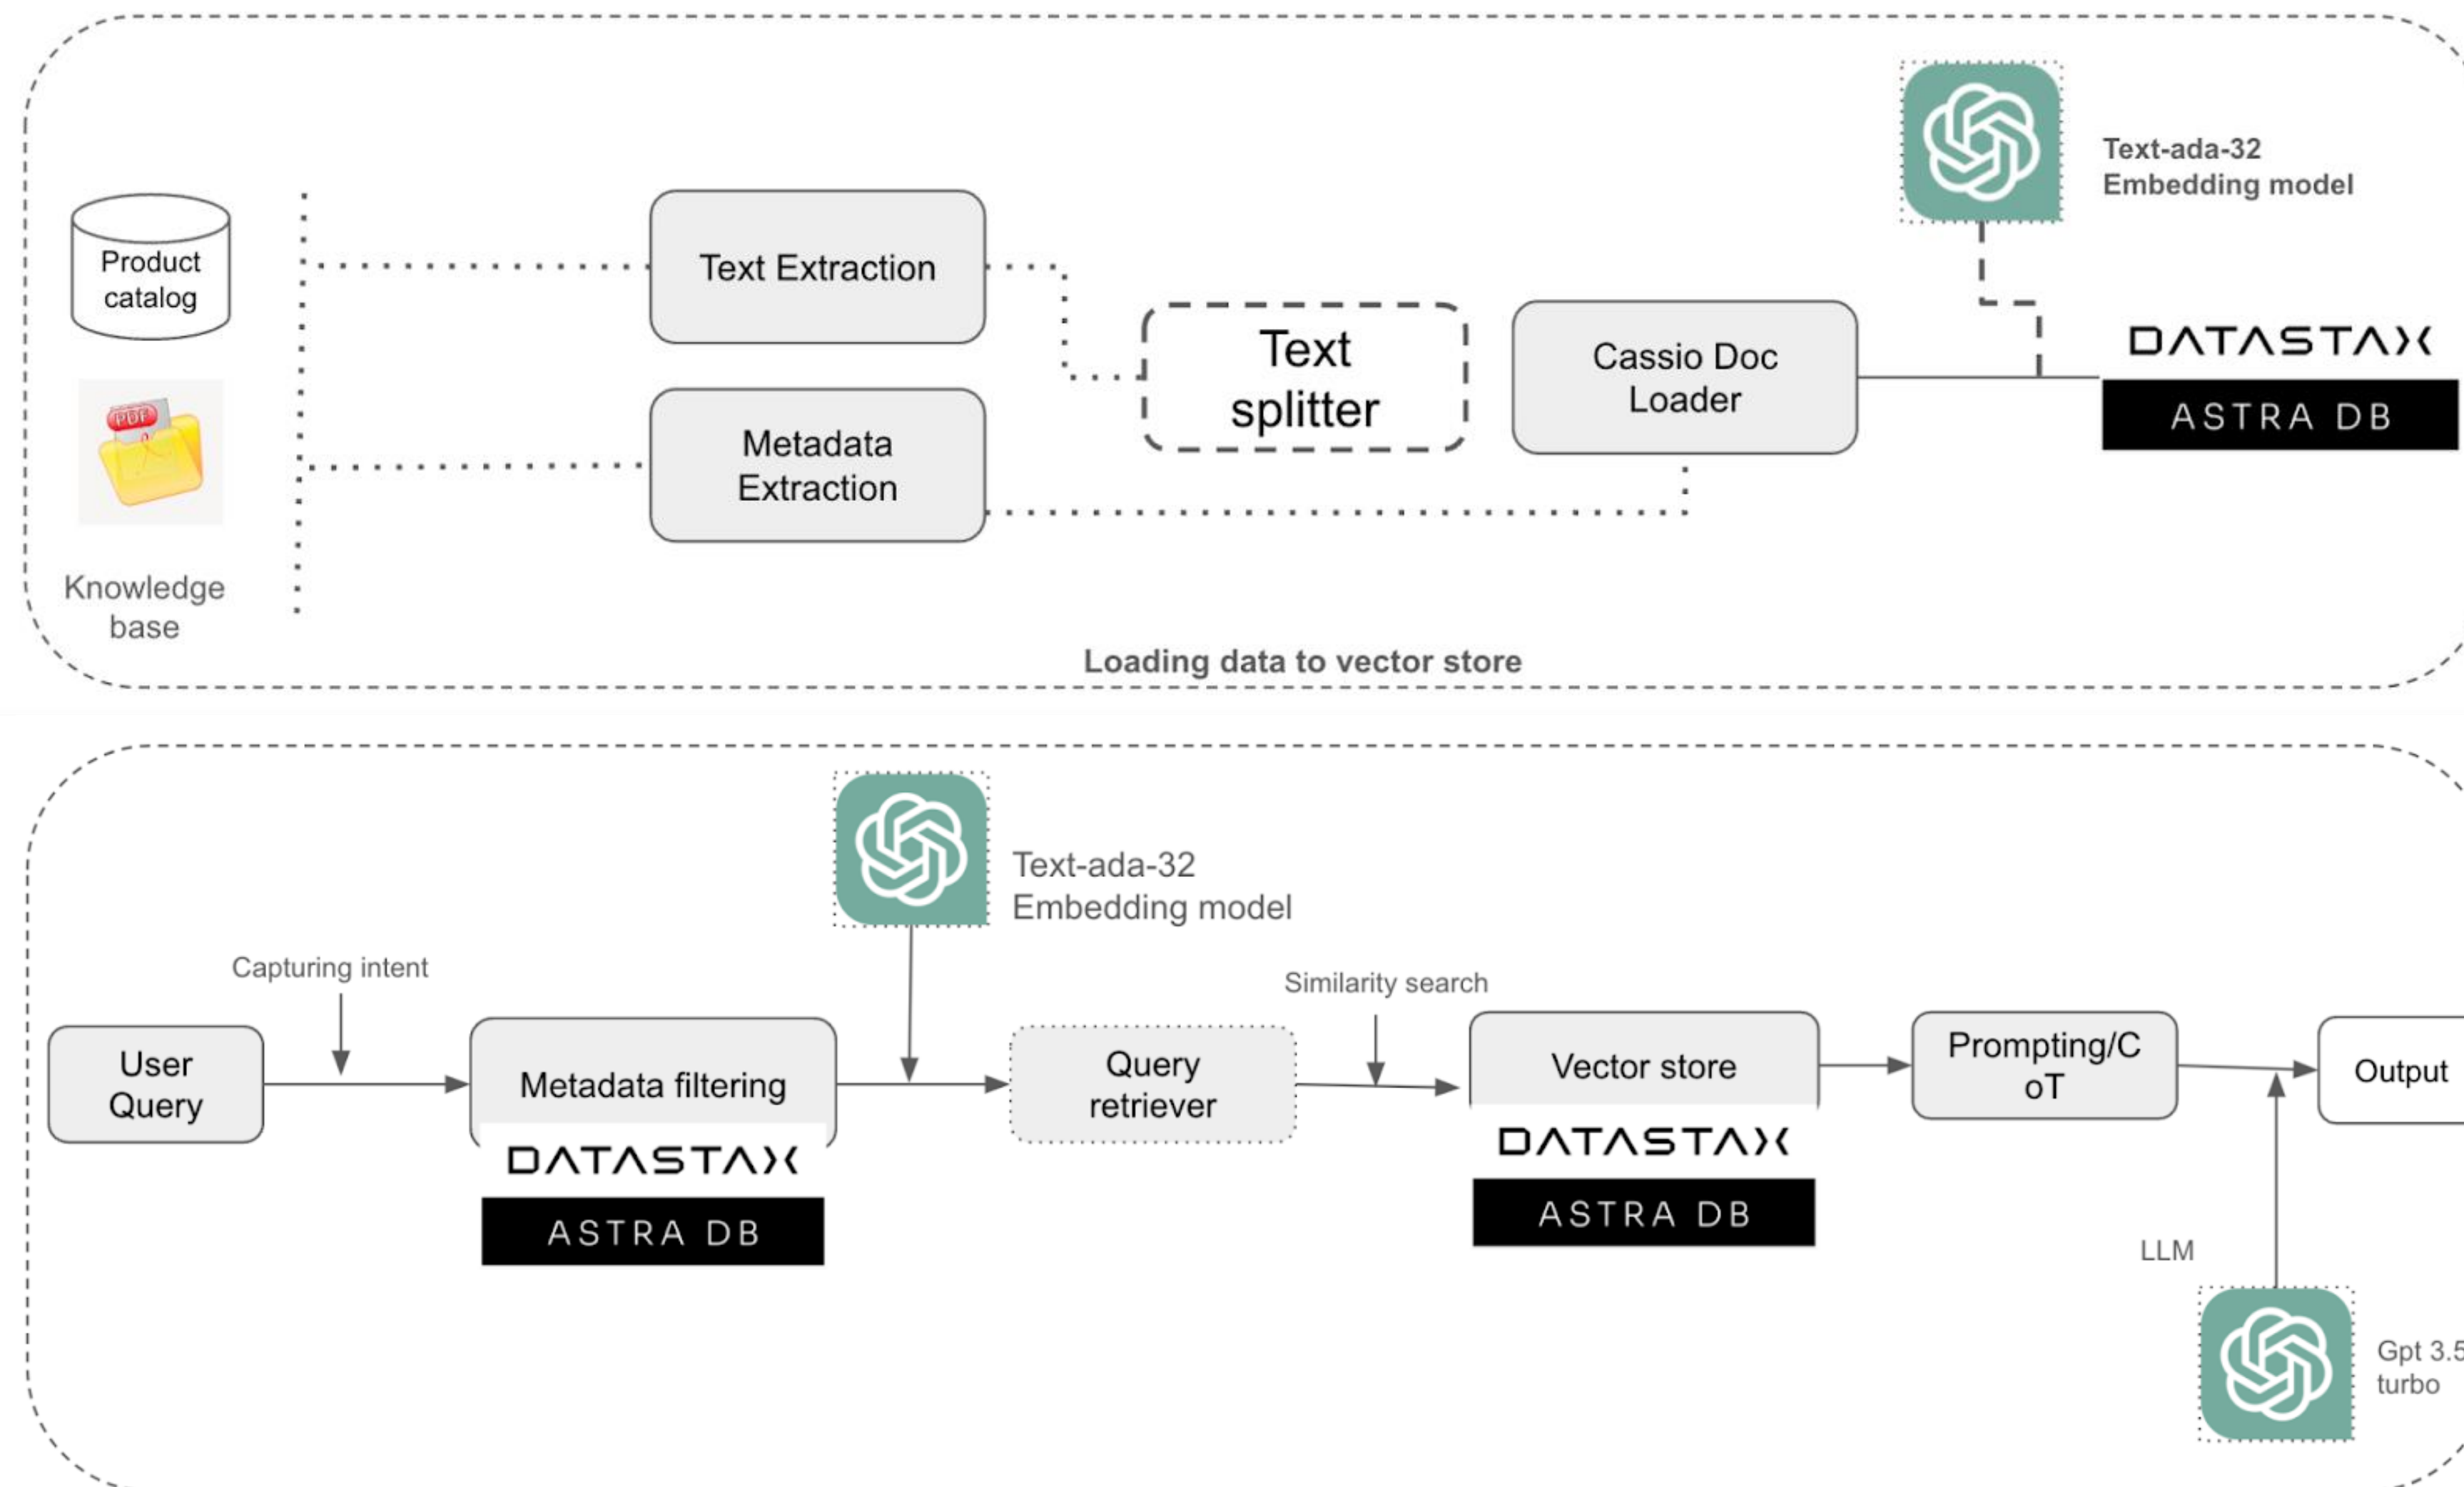
Product Architecture



Personalized AI tutor application with 6M DAU to assist students with their academic and support queries.

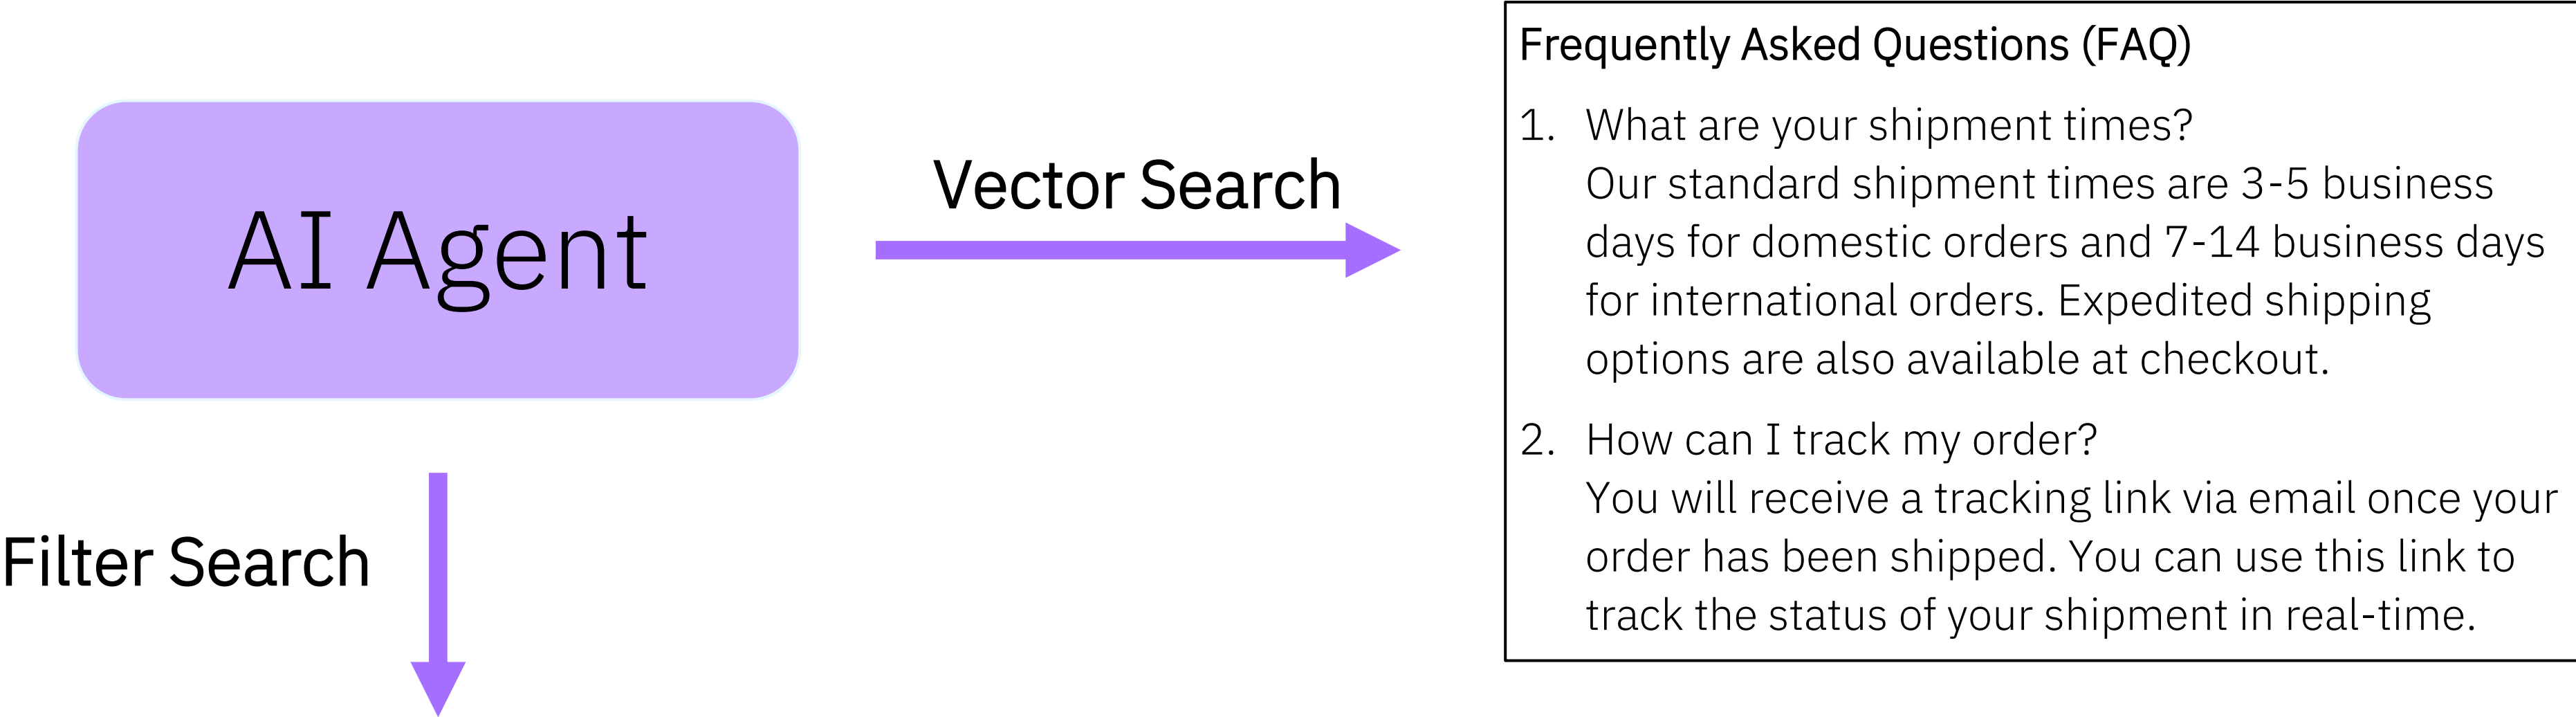
The application also collects feedback for every question and stores it as vectors in Astra DB.

“Handled 50x increase in traffic with zero downtime across 2M students using DataStax.”



DEMO: How to handle both Structured and Unstructured Data

E-Commerce AI Agent: retrieving from both Structured and Unstructured Data



orderNumber	customerFirstName	customerLastName	date	price	status	productIdsOrdered
1001	Sophia	Doe	2024-11-20	91.92	Processing	[349]
1002	John	Lee	2024-11-21	279.15	Cancelled	[516, 189]

productId	name	description
101	Wireless Mouse	Ergonomic wireless mouse with adjustable DPI.
103	Mechanical Keyboard	RGB mechanical keyboard with customizable keys.

Structured Data

Astra DB Tool

Tool to run hybrid vector and metadata search on DataStax Astra DB Collection

Tool Name*

ProductLookup

Tool Description*

A tool used to look up a product b...

Collection Name*

products

Astra DB Application Token*

ASTRA_API_KEY x

Tools Parameters

Open table

Tool

▼

≡

Tool

Astra DB Tool

Tool to run hybrid vector and metadata search on DataStax Astra DB Collection

Tool Name*

OrderLookup

Tool Description*

A tool used to look up an order ba...

Collection Name*

orders

Astra DB Application Token*

ASTRA_API_KEY x

Tools Parameters

Open table

Tool

▼

≡

Tool

UnStructured Data

Astra DB

Ingest and search documents in Astra DB

Astra DB Application Token*

ASTRA_API_KEY x

Collection*

company_faq

Reranker

nvidia/llama-3.2-nv-rerank...

☒

Actions

SEARCH_DOCUMENTS AS_DATAFRAME AS_VECTOR_STORE

Toolset

≡

Agent

Define the agent's instructions, then enter a task to complete using tools.

Model Provider

OpenAI

Model Name

gpt-4.1-mini

OpenAI API Key

OPENAI_KEY x

Agent Instructions

You are a skilled customer service a...

Tools

Toolset

Input

Receiving input

Response

≡

URL tool

Calculator tool

Chat Input

Chat Output

Agent

Define the agent's instructions, then enter a task to complete using tools.

Model Provider ⓘ

OpenAI



Model Name ⓘ

gpt-4.1-mini



OpenAI API Key ⓘ

OPENAI_KEY ×

Agent Instructions ⓘ

You are a skilled customer service agent



Tools ⓘ

Input ⓘ

Receiving input



Response



UnStructured Data

DS Astra DB



Ingest and search documents in Astra DB

Astra DB Application Token* ⓘ

ASTRA_API_KEY ×

Collection* ⓘ



company_faq



Reranker ⓘ



nvidia/llama-3.2-nv-rerank...



Actions ⓘ



SEARCH_DOCUMENTS

AS_DATAFRAME

AS_VECTOR_STORE

Toolset



Structured Data

DS Astra DB Tool ▶

Tool to run hybrid vector and metadata search on DataStax Astra DB Collection

Tool Name* ⓘ

ProductLookup

🌐

Tool Description* ⓘ

A tool used to look up a product b...

🌐

Collection Name* ⓘ

products

🌐

Astra DB Application Token* ⓘ

ASTRA_API_KEY ×

Tools Parameters ⓘ

📄 Open table

DS Astra DB Tool ▶

Tool to run hybrid vector and metadata search on DataStax Astra DB Collection

Tool Name* ⓘ

OrderLookup

🌐

Tool Description* ⓘ

A tool used to look up an order ba...

🌐

Collection Name* ⓘ

orders

🌐

Astra DB Application Token* ⓘ

ASTRA_API_KEY ×

Tools Parameters ⓘ

📄 Open table




New chat

Test your flow with a chat prompt

Tell me about order 1001|




Send



User

Tell me about order 1001



Agent




Accessing **OrderLookup**

) 1.4s



Send a message...




Stop





User

Tell me about order 1001



Agent




Accessing **ProductLookup**




3.9s



Send a message...



Stop 



AI gpt-4.1-mini

✓ **Finished**

3.8s ▼

Order 1001 was placed by John Doe on November 20, 2024. The order status is "Shipped" and the total price is \$120.50. The order includes the following products:

- 1. Wireless Mouse - Ergonomic wireless mouse with adjustable DPI.
- 2. Mechanical Keyboard - RGB mechanical keyboard with customizable keys.

If you need more details or assistance with this order, please let me know!

|



Send




New chat

Test your flow with a chat prompt

What is the policy on returns?




Send



User

What is the policy on returns?



Agent




Accessing **search_documents**


🔄 2.8s



Send a message...




Stop 🔄



User

What is the policy on returns?




AI gpt-4.1-mini

✓ Finished

3.6s ▼

The return policy is that returns are accepted within 30 days of purchase for items in their original condition. For more details on how to initiate a return, please visit the website. If you need further assistance, feel free to ask!



Send




New chat

Test your flow with a chat prompt

How can I cancel order 1002, and what is the shipping policy?




Send



User

How can I cancel order 1002, and what is the shipping policy?



Agent




Accessing **OrderLookup**

) 1.4s




Send a message...




Stop





User

How can I cancel order 1002, and what is the shipping policy?



Agent




Accessing **search_documents**

🔄 2.6s



Send a message...



Stop 🔄

✓ **Finished**

5.5s



Regarding your order cancellation for order number 1002: The order is currently in "Pending" status. Typically, changes or cancellations can be made if the order has not yet been processed. It is best to contact customer support as soon as possible to request cancellation.

As for the shipping policy:

- Standard shipment times are 3-5 business days for domestic orders and 7-14 business days for international orders.
- Expedited shipping options are available at checkout.
- You will receive a tracking link via email once your order has been shipped to track the status in real-time.

|



Send

Thank You

Hilton Rosenfeld
IBM, Senior Solution Engineer
hilton.rosenfeld@ibm.com

