



20 May, 2025

Gen AI-Powered Industry Use-Cases with IDMC

- Masaki Nakasima, Sr. Solutions Architect, CSA
- Anita Ayyagari, Senior Solutions Architect, CSA
- Sreehari Desai, Senior Solutions Architect, CSA
- Ravi Tubati, Solutions Architect, CSA

**Where data
& AI come to** 

Housekeeping Tips



- Today's Webinar is scheduled for **1 hour**
- The session will include a webcast and then your questions will be answered live at the end of the presentation
- All dial-in participants will be muted to enable the speakers to present without interruption
- Questions can be submitted to "All Panelists" via the **Q&A option** and we will respond at the end of the presentation
- The webinar is **being recorded** and will be available on our [Success Portal](#) - where you can download the **slide deck** for the presentation. The link to the recording will be emailed as well.
- Please take time to complete the **post-webinar survey** and provide your feedback and suggestions for upcoming topics.

Feature Rich Success Portal



Bootstrap trial and
POC Customers



Enriched Customer
Onboarding
experience



Product Learning
Paths and Weekly
Expert Sessions



Informatica
Concierge



Tailored training and
content
recommendations

More Information



Success Portal

<https://success.informatica.com>



Communities & Support

<https://network.informatica.com>



Documentation

<https://docs.informatica.com>



University

<https://www.informatica.com/in/services-and-training/informatica-university.html>

Safe Harbor

The information being provided today is for informational purposes only. The development, release, and timing of any Informatica product or functionality described today remain at the sole discretion of Informatica and should not be relied upon in making a purchasing decision.

Statements made today are based on currently available information, which is subject to change. Such statements should not be relied upon as a representation, warranty or commitment to deliver specific products or functionality in the future.



May 20th 2025

Gen AI-Powered Industry Use-Cases with IDMC

Masaki Nakasima, Ravi Tubati, Anita Ayyagari, Sreehari Desai

Customer Success Architects

Where data & AI come to 

Agenda

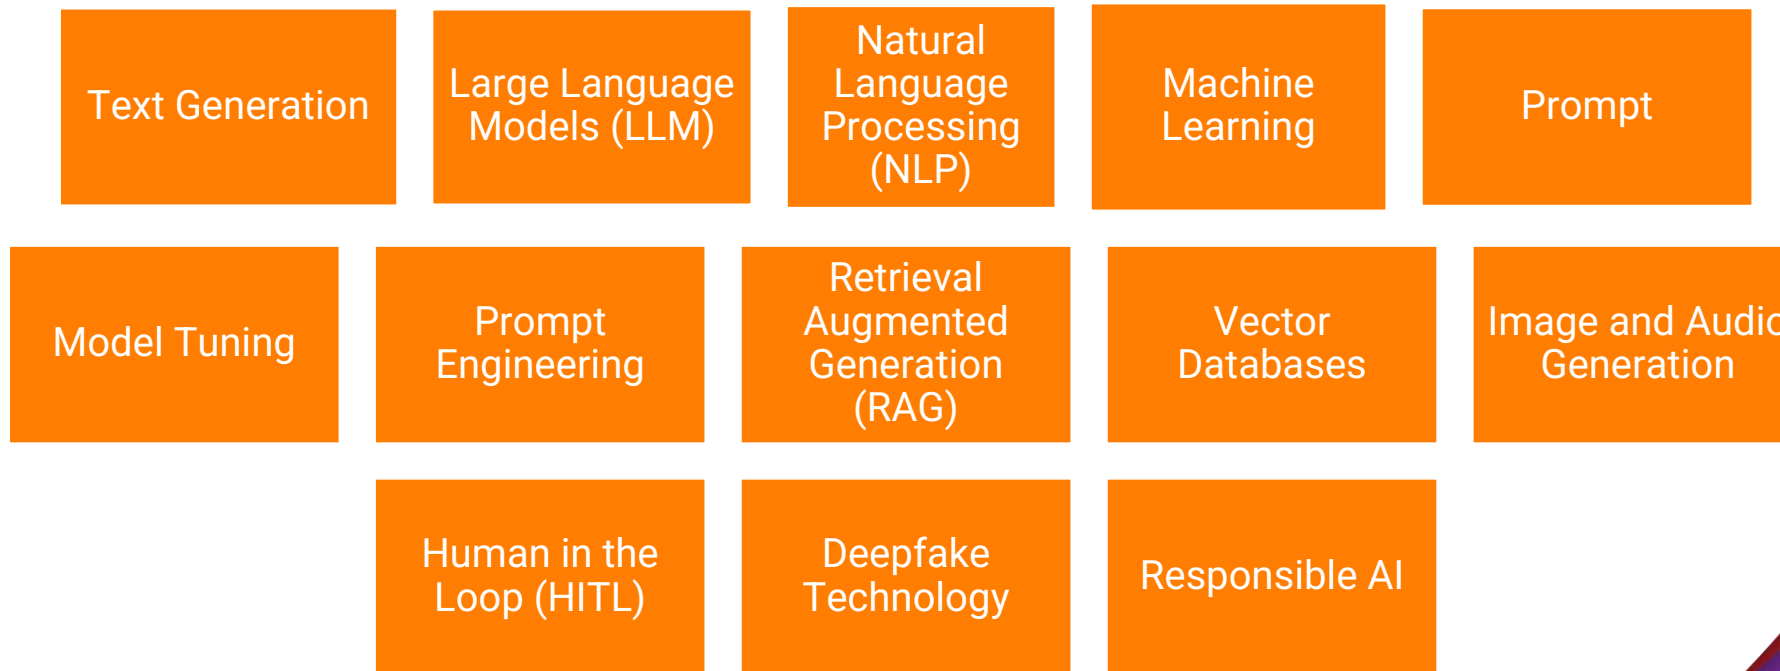
- ❑ **Recap of Episode 1**
- ❑ **Understanding Gen AI Project Flow – Breaking down the stages of a Gen AI implementation**
- ❑ **Deep Dive into Informatica's blueprints of Gen AI Architectures**
- ❑ **Gen AI Use Case with Industry examples:**
 - ❑ **Healthcare-Automating Patient Admission**
 - ❑ **Insurance- Claims Processing with Gen AI**
- ❑ **Demos:**
 - ❑ **Ingesting documents to Vector Database using CDI**
 - ❑ **Creating a RAG Ingestion Pipeline with Low Code No-Code**
- ❑ **Q&A Session – Address queries and discuss key takeaways**



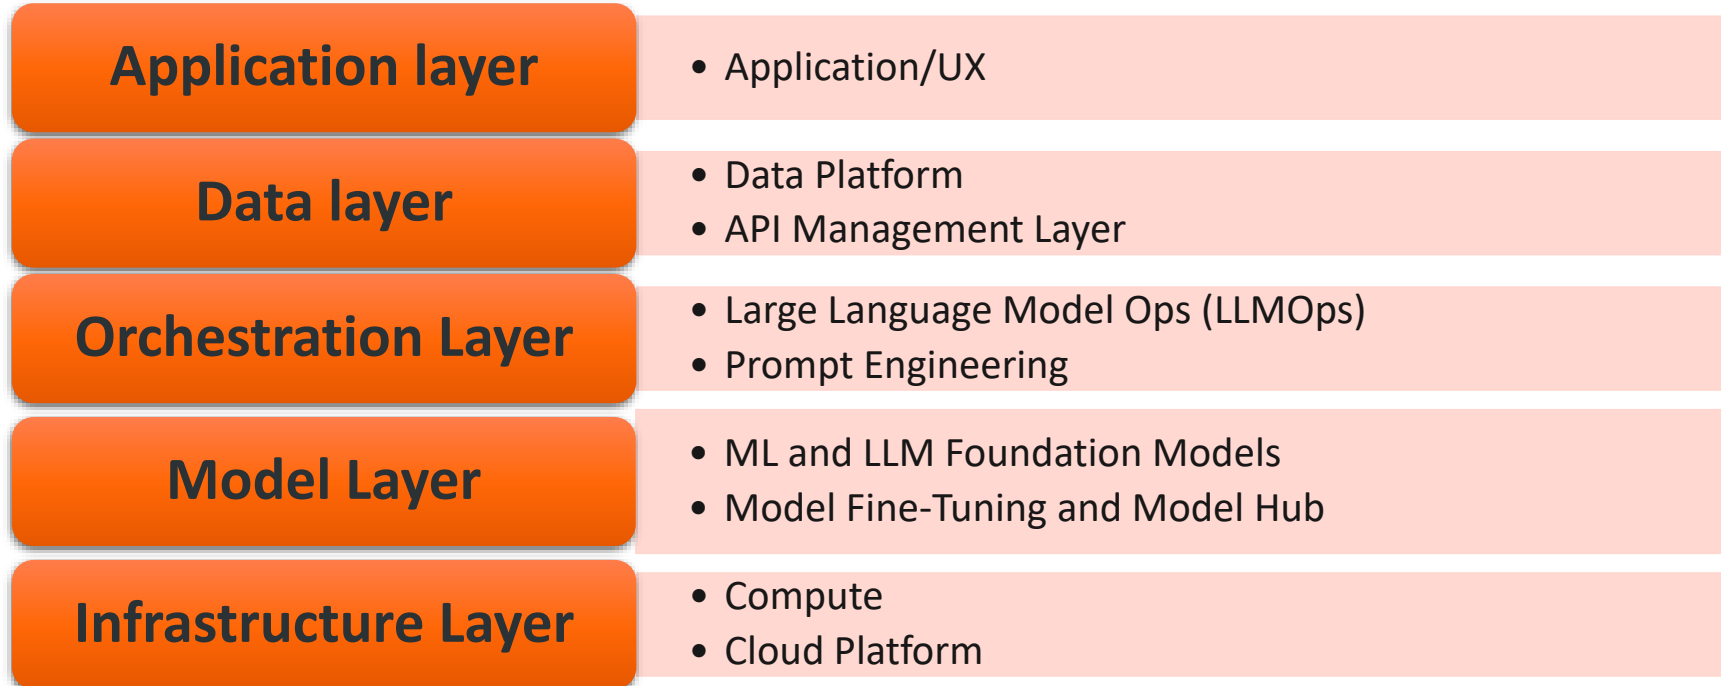
Episode 1 Recap

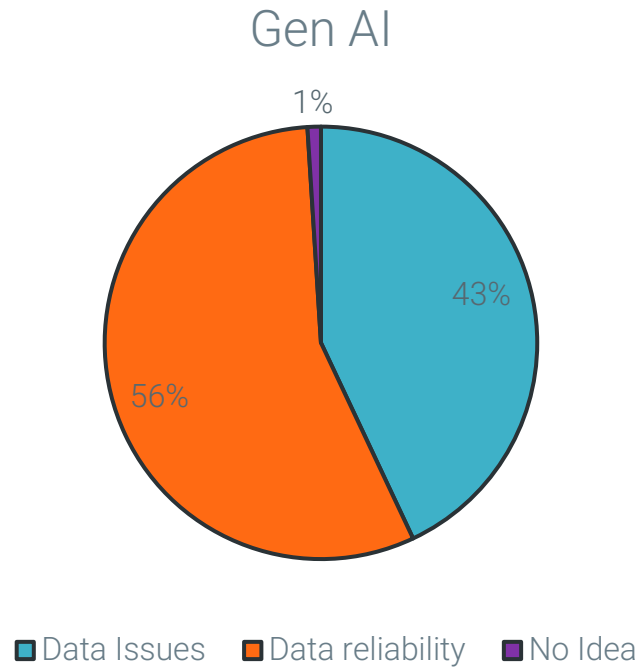
Where data & AI come to **LIFE**

Gen AI Terminology



Layers of Gen AI Architecture







Informatica®

Recipes

At a glance!

Episode 1 Recap

Recipes

All Recipes (49)

AI Agent for Dynamics 365 using Azure OpenAI

Use the AI Agent for Dynamics 365 using Azure OpenAI recipe to autonomously interact with Dynamics 365 and address user queries.

Tags: AI, agent

Published by Informatica

AI Agent for Salesforce using Google Gemini

Use the AI Agent for Salesforce using Google Gemini recipe to autonomously interact with Salesforce and address user queries.

Tags: caI, gemini, Salesforce, AI

Published by Informatica

AI Agent for Supply Chain Management using Amazon Bedrock

Use AI Agent for Supply Chain Management using Amazon Bedrock recipe to handle natural language queries across services, such as Informatica Master Data Management (MDM), Informatica Data...

Tags: GenAI, SupplyChain, AWS, Bedrock

Published by Informatica

Amazon Bedrock Chat with History

Use the Amazon Bedrock Chat with History recipe to maintain chat history in a file and use it as context for the next query you ask and the Large Language Model's (LLM) response.

Tags: Amazon, Bedrock, Chat, History

Published by Informatica

Amazon Bedrock Prompt Chaining

Use the Amazon Bedrock Prompt Chaining recipe to design prompt chains and resolve them in sequence so that it provides the desired responses from a Large Language Model (LLM).

Tags: Amazon, Bedrock, Prompt, Chaining

Published by Informatica

Automobile Insurance Claim Processing with Amazon Bedrock

Use the Automobile Insurance Claim Processing with Amazon Bedrock recipe to evaluate a claim request, assess the vehicle damage, and estimate the insurance payout based on the uploaded images.

Tags: GenAI, insurance, human_task, images, amazon, bedrock

Published by Informatica

Automobile Insurance Claim Processing with Azure OpenAI

Use the Automobile Insurance Claim Processing with Azure OpenAI recipe to evaluate a claim request, assess the vehicle damage, and estimate the insurance payout based on the uploaded images.

Tags: GenAI, insurance, human_task, images, azure, openAI

Published by Informatica

Automobile Insurance Claim Processing with Google Gemini

Use the Automobile Insurance Claim Processing with Google Gemini recipe to evaluate a claim request, assess the vehicle damage, and estimate the insurance payout based on the uploaded images.

Tags: GenAI, insurance, human_task, images, gemini

Published by Informatica

Azure AI Search Ingestion Automation with ADLS

Use Azure AI Search Ingestion Automation with ADLS to learn how to create an index, data source, and indexer for the existing files in Amazon Data Lake Storage (ADLS), populate the index with the data, and...

Tags: Azure, AI, Search

Published by Informatica

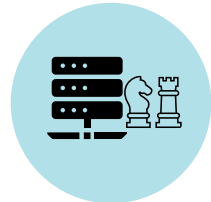
Where data & AI come to

Understanding Gen AI Project Flow – Breaking down the stages of a Gen AI implementation

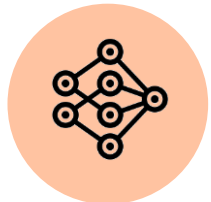
Overview of the Stages



Problem Definition



Data Strategy



Model Strategy



Architecture &
Infrastructure



Training / Tuning



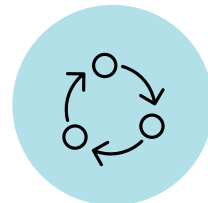
Testing & Evaluation



Deployment



Governance, Risk & Compliance



Continuous Improvement

Governance, Risk & Compliance

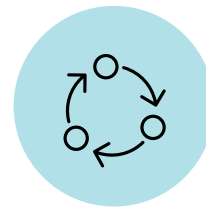


Address ethics, data security, regulatory requirements.

1. Problem Definition:
2. Data Strategy:
3. Model Strategy:
4. Architecture and Infrastructure:
5. Training / Tuning:
6. Evaluation and Testing:
7. Deployment:
8. Monitoring and Maintenance:
9. Iteration and Improvement:



Continuous Improvement



Monitor, retrain, and evolve the model with real-world feedback.

1. Performance Monitoring
2. Error and Incident Tracking
3. User Feedback Collection
4. Regular Updates and Improvements
5. Scalability Adjustments
6. Security and Compliance Audits
7. Backup and Recovery Testing
8. Version Management and Documentation
9. User Training and Support Updates
10. Sustainability and Efficiency Evaluations



Overview of the stages - Final Thoughts

Start small. Prove value. Scale responsibly.



May 20th 2025

Understanding Gen AI Project Flow – Breaking down the stages of a Gen AI implementation

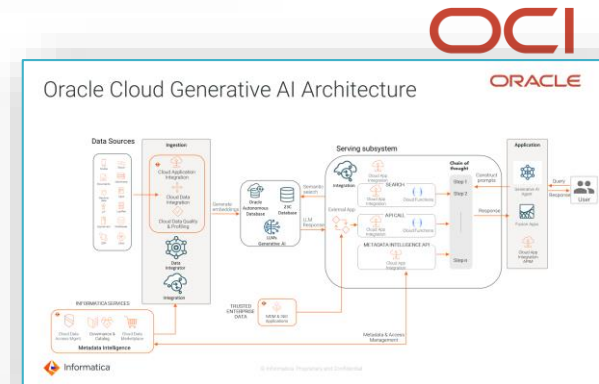
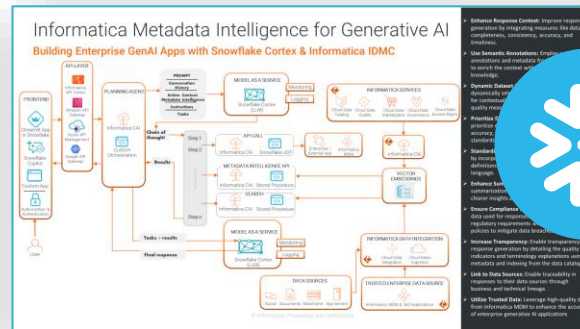
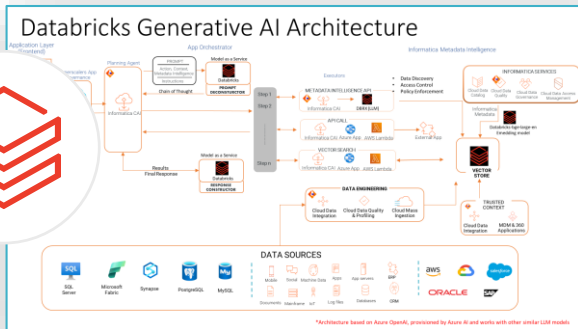
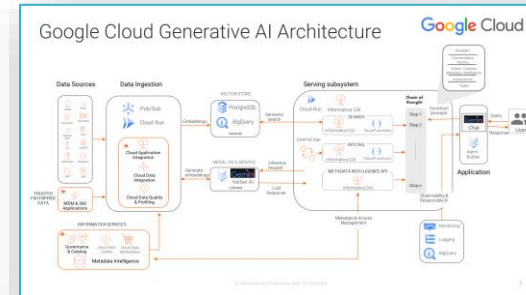
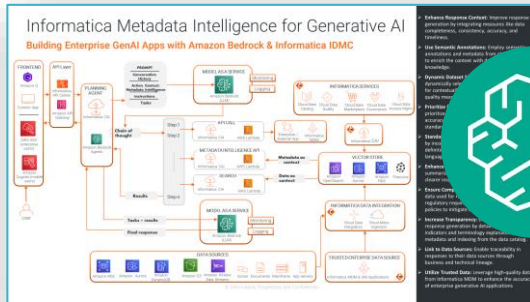
Ravibabu Tubati

Customer Success Architect(CSA)

Where data & AI come to 

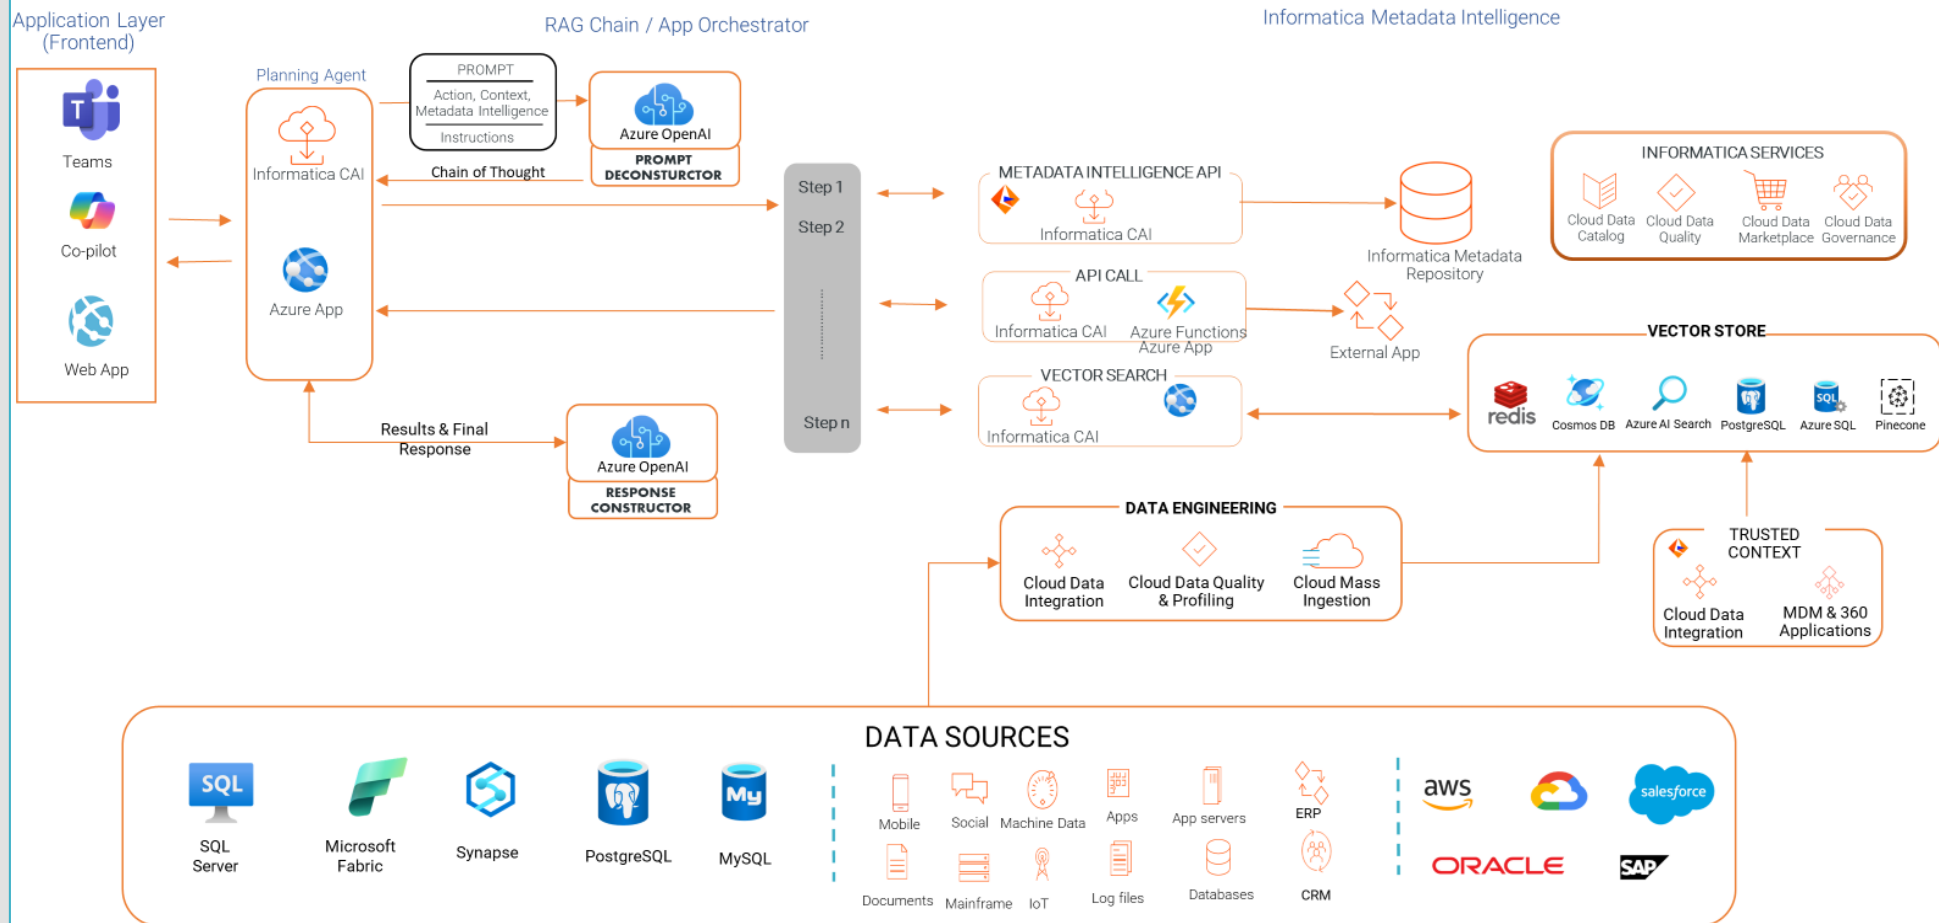
Deep Dive into Informatica's blueprints of Gen AI Architectures

Blueprints for Enterprise-Grade Gen AI





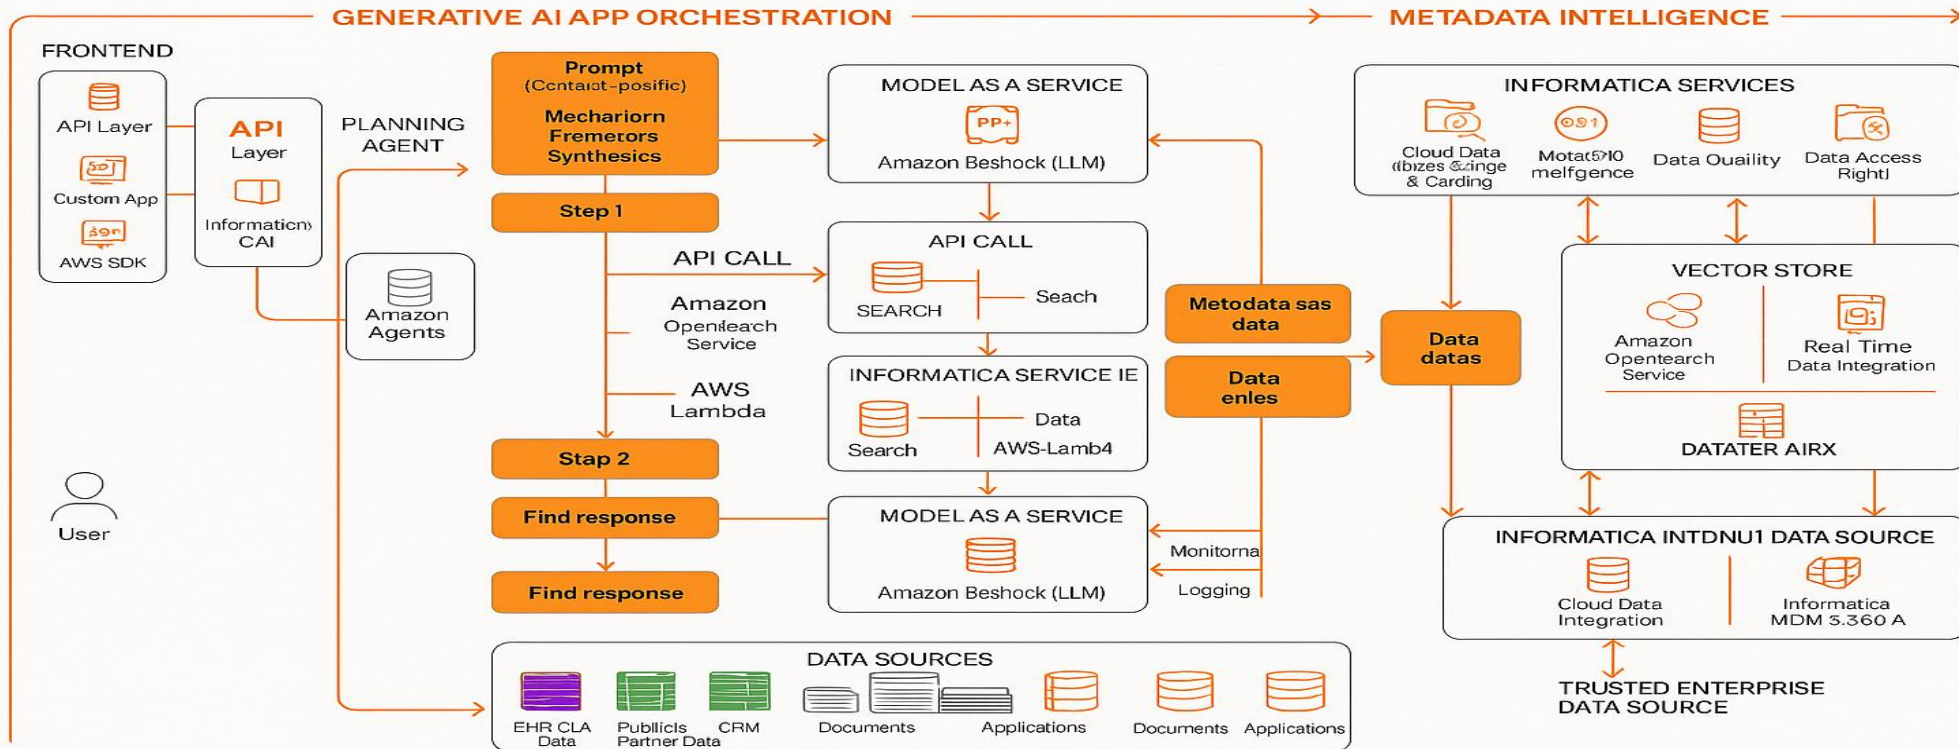
Informatica Enabled GenAI on Azure



GenAI Blueprint for Amazon Web Services (AWS)

GenAI Blueprint for Amazon Web Services (AWS)

Retrieval-Augmented Generation (RAG) framework in Informatica LEMs



GenAI Blueprint for AWS

- Powered by Informatica and AWS
- A reference architecture for building GenAI applications on AWS.

Core Pillars

1. Generative AI App Orchestration
2. Metadata Intelligence
3. Data Engineering
4. Model as a Service
5. Vector Store Integration

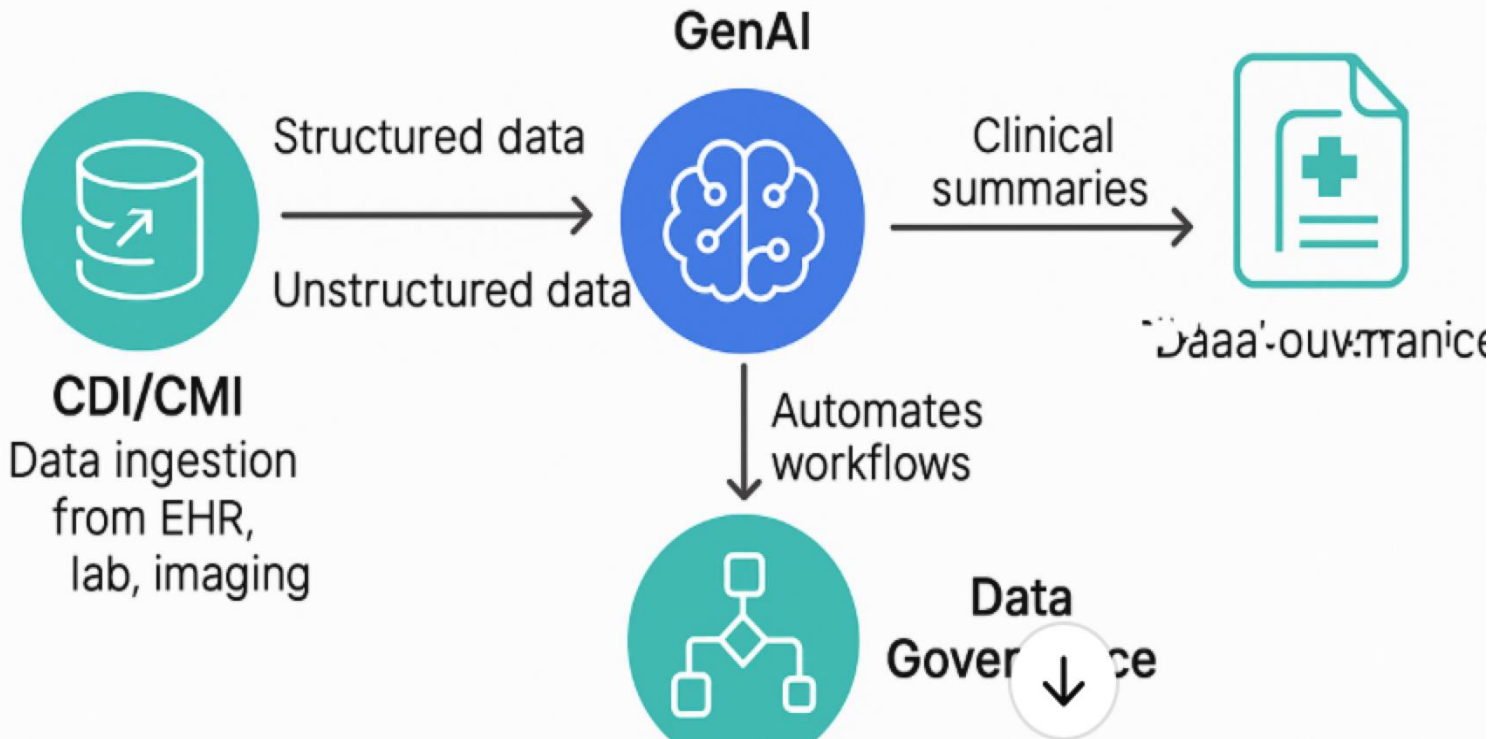
Gen AI Blueprint for AWS

- Powered by Informatica and AWS
- A reference architecture for building GenAI applications on AWS.

Healthcare Use Case 1: Intelligent Clinical Documentation

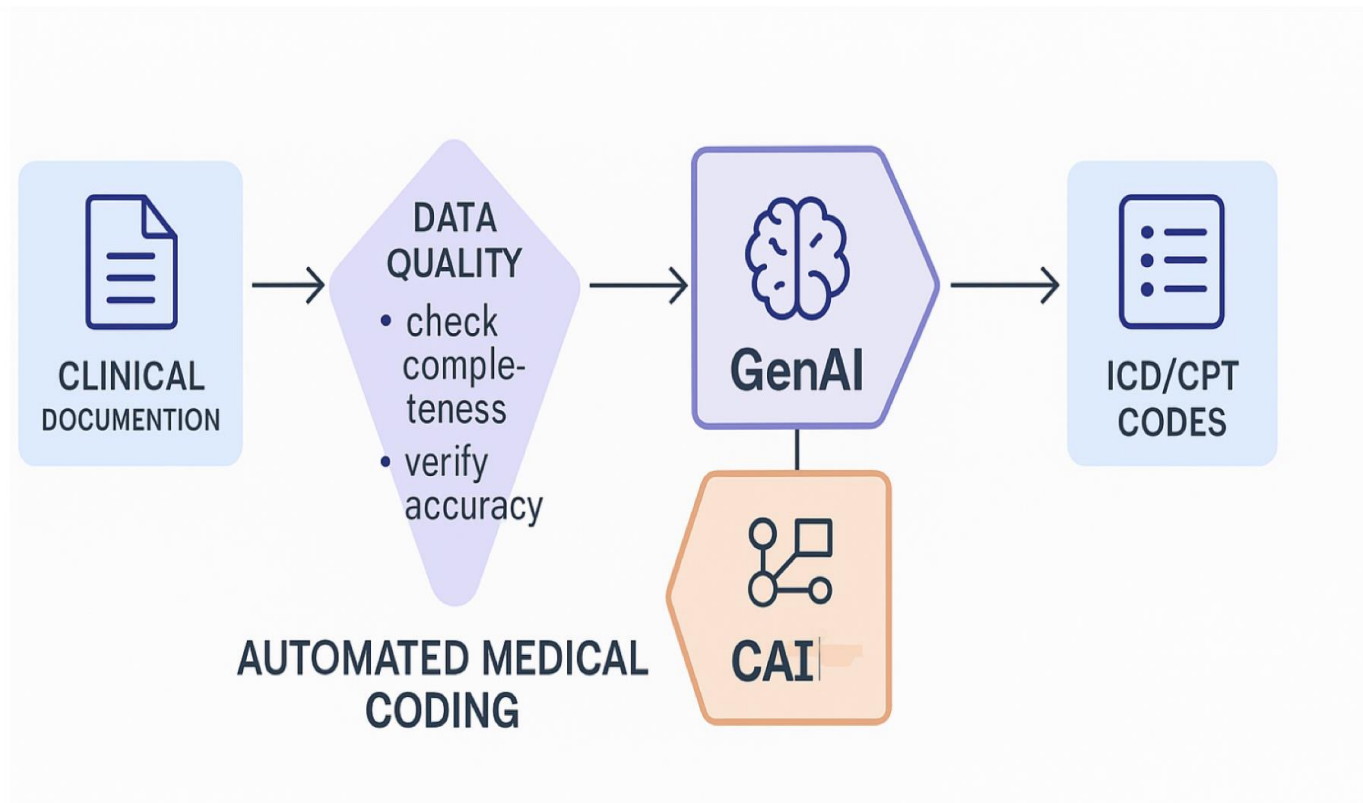
Problem: Time-consuming manual note-taking.

Solution: GenAI generates clinical summaries from structured data



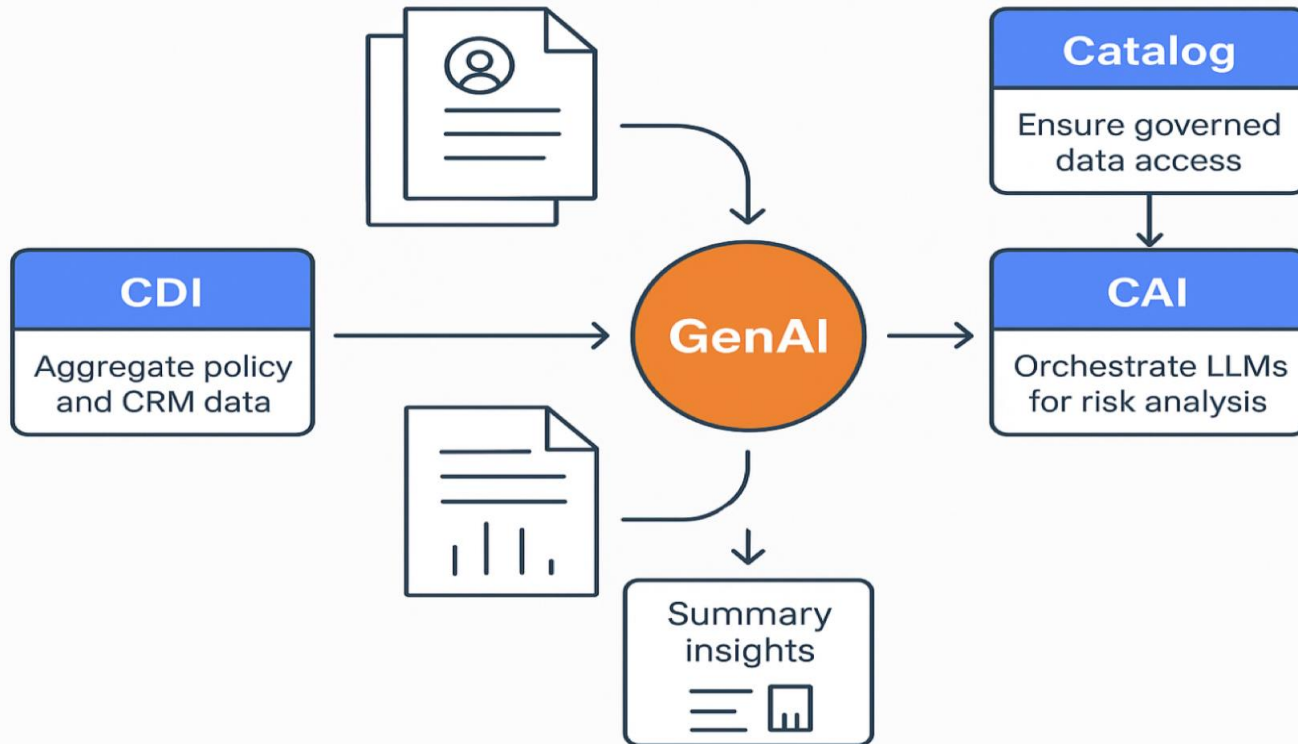
- Problem: Time-consuming manual note-taking.
- Solution: GenAI generates clinical summaries from structured & unstructured data.
- IDMC Role:
 - * CDI/CMI: Data ingestion from EHR, lab, imaging.
 - * CAI: Automates GenAI workflows.
 - * Data Governance: Ensures HIPAA compliance.

Healthcare Use Case 2: Automated Medical Coding & Claims



- Problem: Delayed claims due to manual coding.
- Solution: GenAI maps documentation to ICD/CPT codes.
- IDMC Role:
 - * Data Quality: Maintains accuracy.
 - * CAI Flows: Triggers real-time model use.

Insurance Use Case 1: Underwriting Automation



- Problem: Manual review of extensive data.
- Solution: GenAI summarizes applications and risk data.
- IDMC Role:
 - * CDI: Aggregates policy and CRM data.
 - * CAI: Orchestrates LLMs for risk analysis.
 - * Catalog: Ensures governed data access.

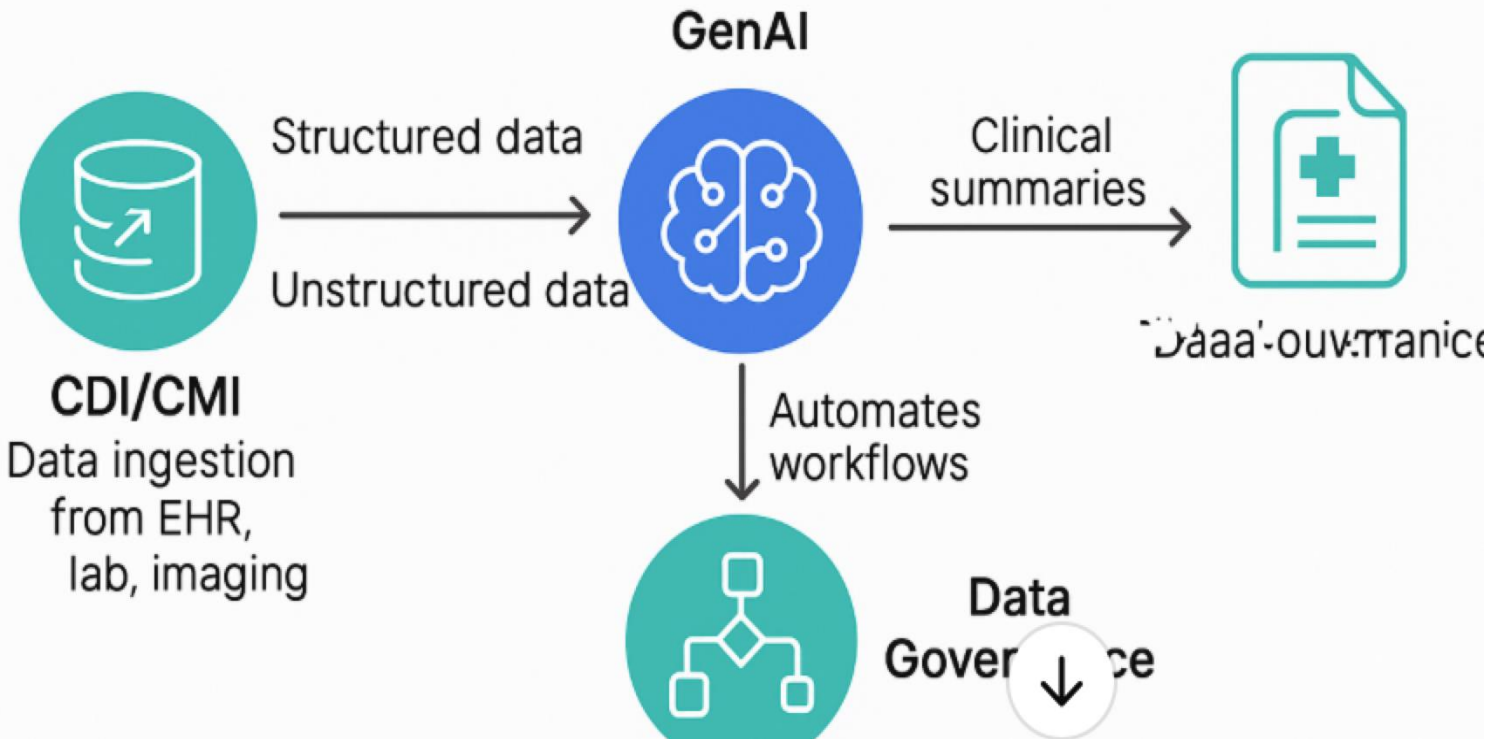
Integrating EHR files with Vector Database using CDI

DEMO

Healthcare Use Case 1: Intelligent Clinical Documentation

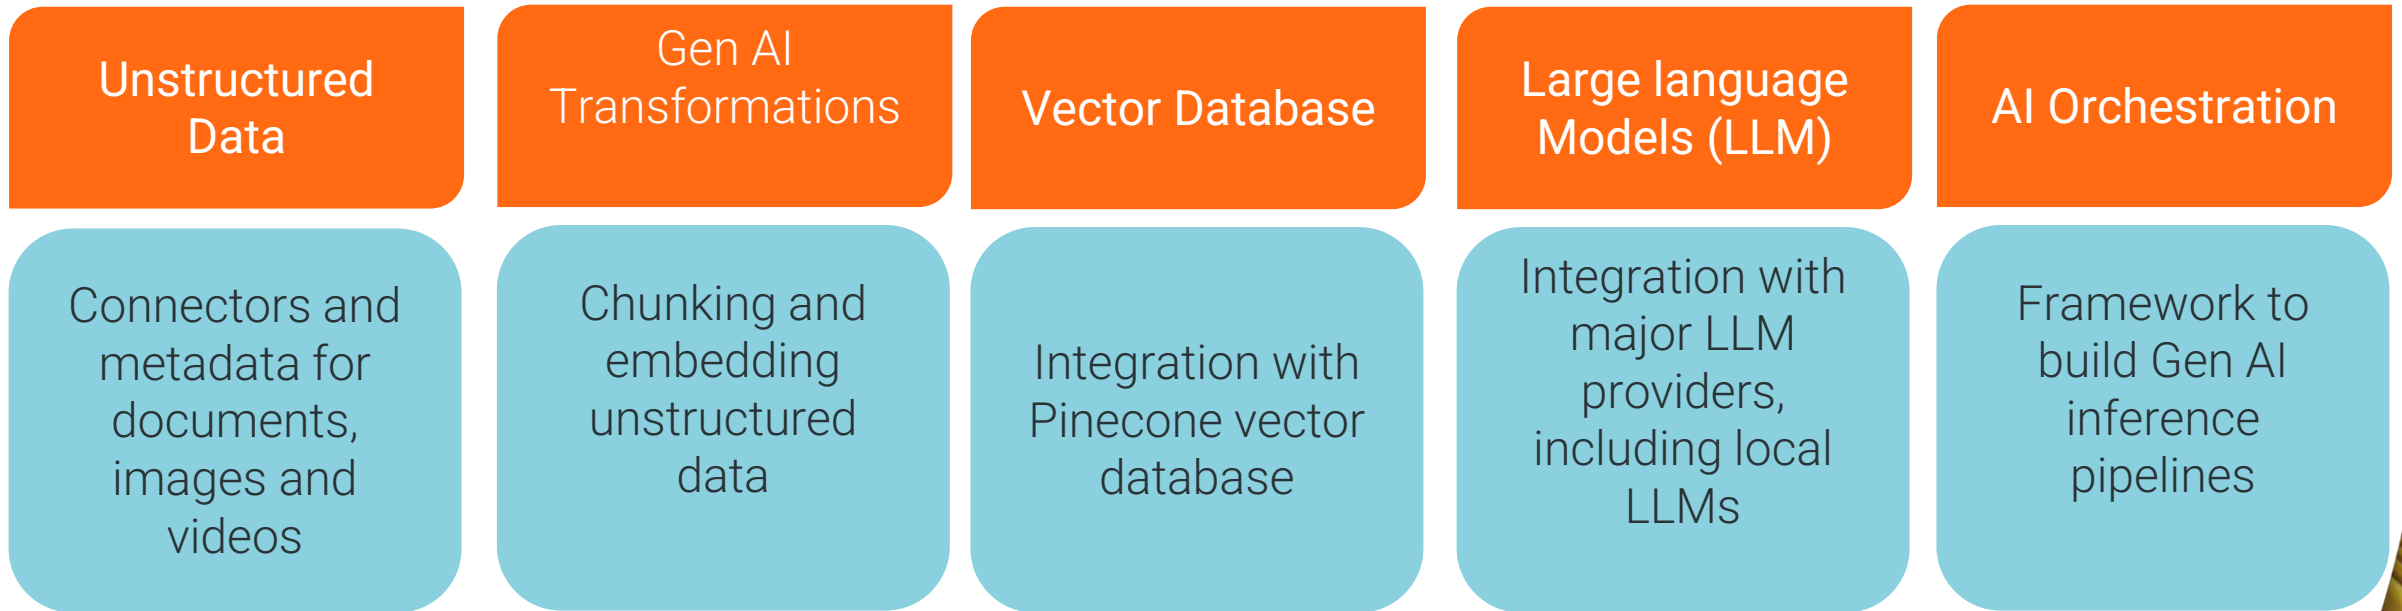
Problem: Time-consuming manual note-taking.

Solution: GenAI generates clinical summaries from structured data

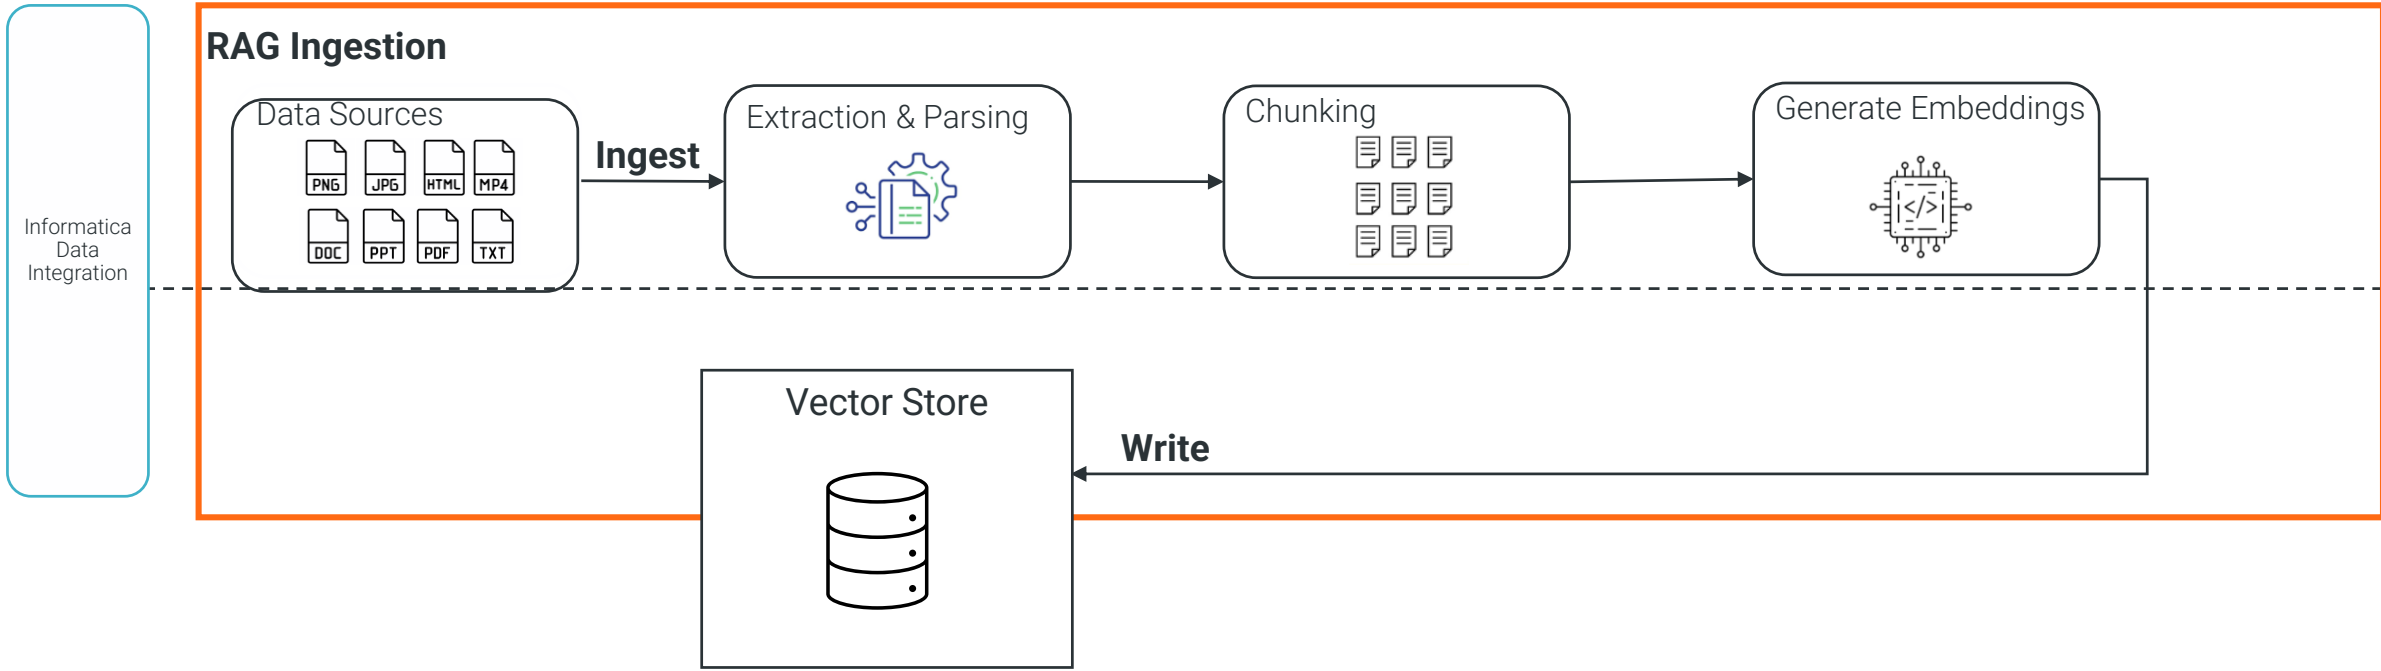


- Problem: Time-consuming manual note-taking.
- Solution: GenAI generates clinical summaries from structured & unstructured data.
- IDMC Role:
 - * CDI/CMI: Data ingestion from EHR, lab, imaging.
 - * CAI: Automates GenAI workflows.
 - * Data Governance: Ensures HIPAA compliance.

Gen AI Requirements in Modern Data Architecture



Informatica RAG Ingestion and AI Inference Architecture



Gen AI Transformations

❑ Chunking transformation



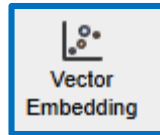
Chunking transformation splits large pieces of text into smaller segments, or chunks.

Sentence splitting

Word splitting

Fixed length character splitting

❑ Vector Embedding transformation



Vector Embedding transformation to generate vector embeddings for input text, capturing the semantic meaning of the text in a vector format.

Word embedding

BERT embedding

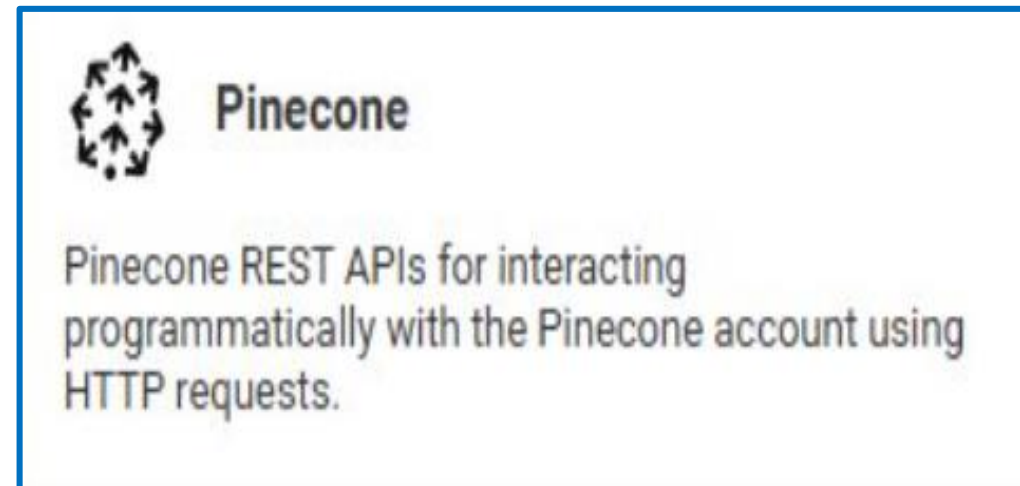
Vector DB Connector (Pinecone)

❑ Pinecone Connector

Pinecone Connector streamlines the process of ingesting data into Pinecone, performing vector searches, and integrating search results into your application workflows.

Connection Details

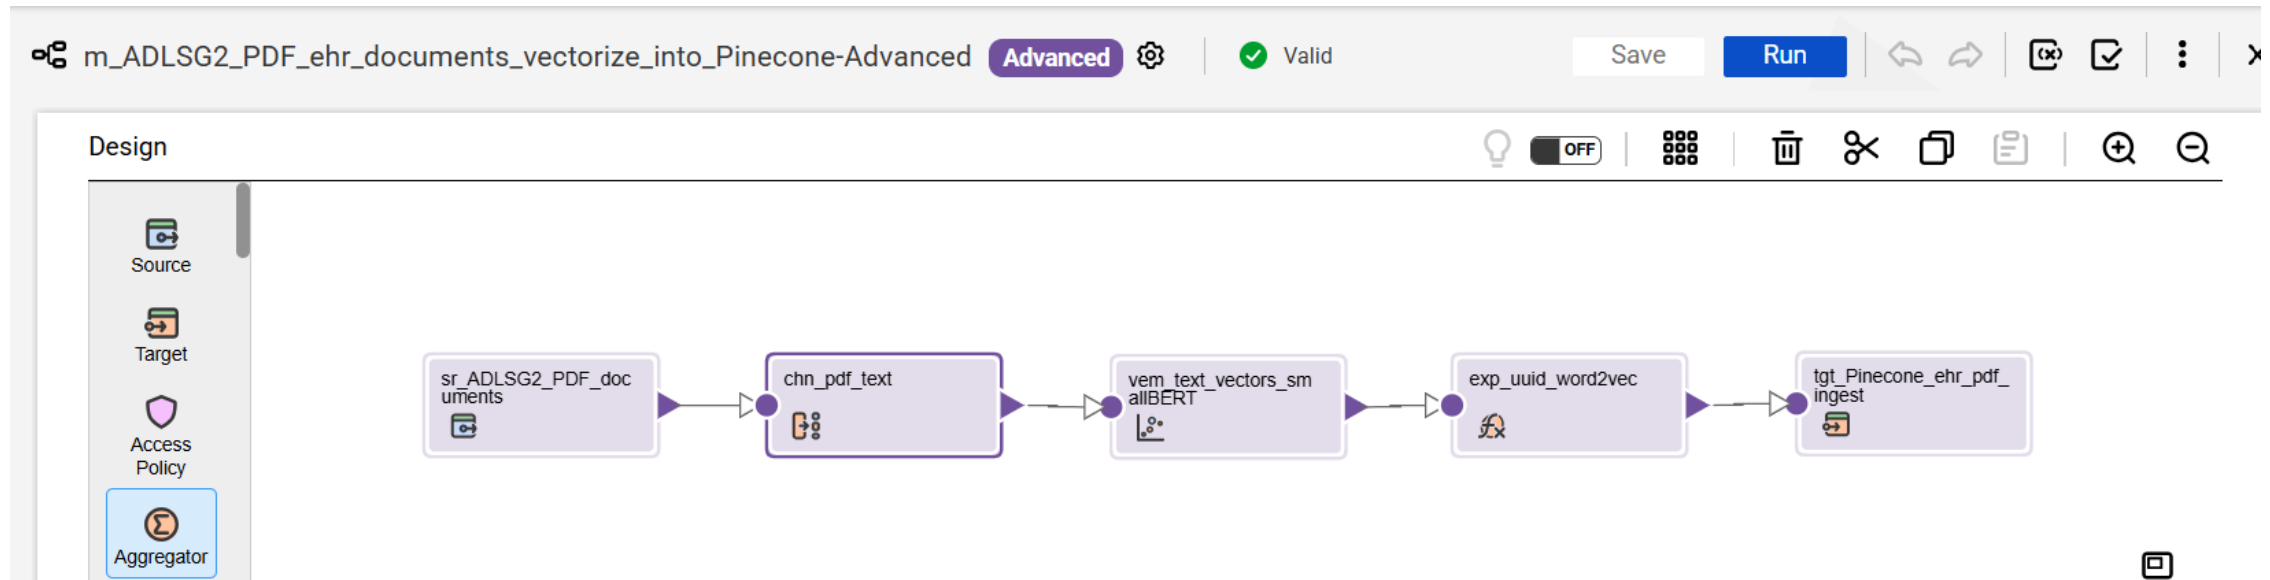
Connection Name:*	<input type="text" value="Pinecone"/>
Description:	<input type="text"/>
Type:* ?	<input type="text" value="Pinecone"/>
<input type="checkbox"/> Use Secret Vault ?	
Runtime Environment:* ?	<input type="text" value="Select..."/>
API Key:* ?	<input type="text"/>



Note: The API key of your Pinecone account to authenticate access to the Pinecone APIs.

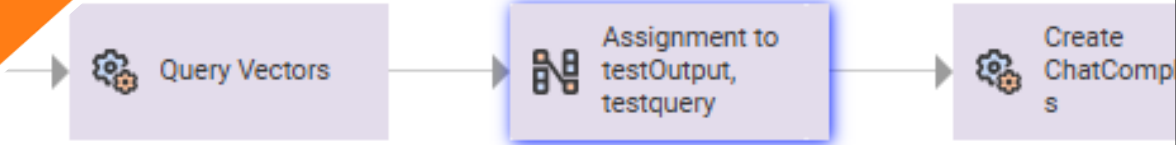
Integrating EHR files with Vector Database

CDI-Advanced Mapping

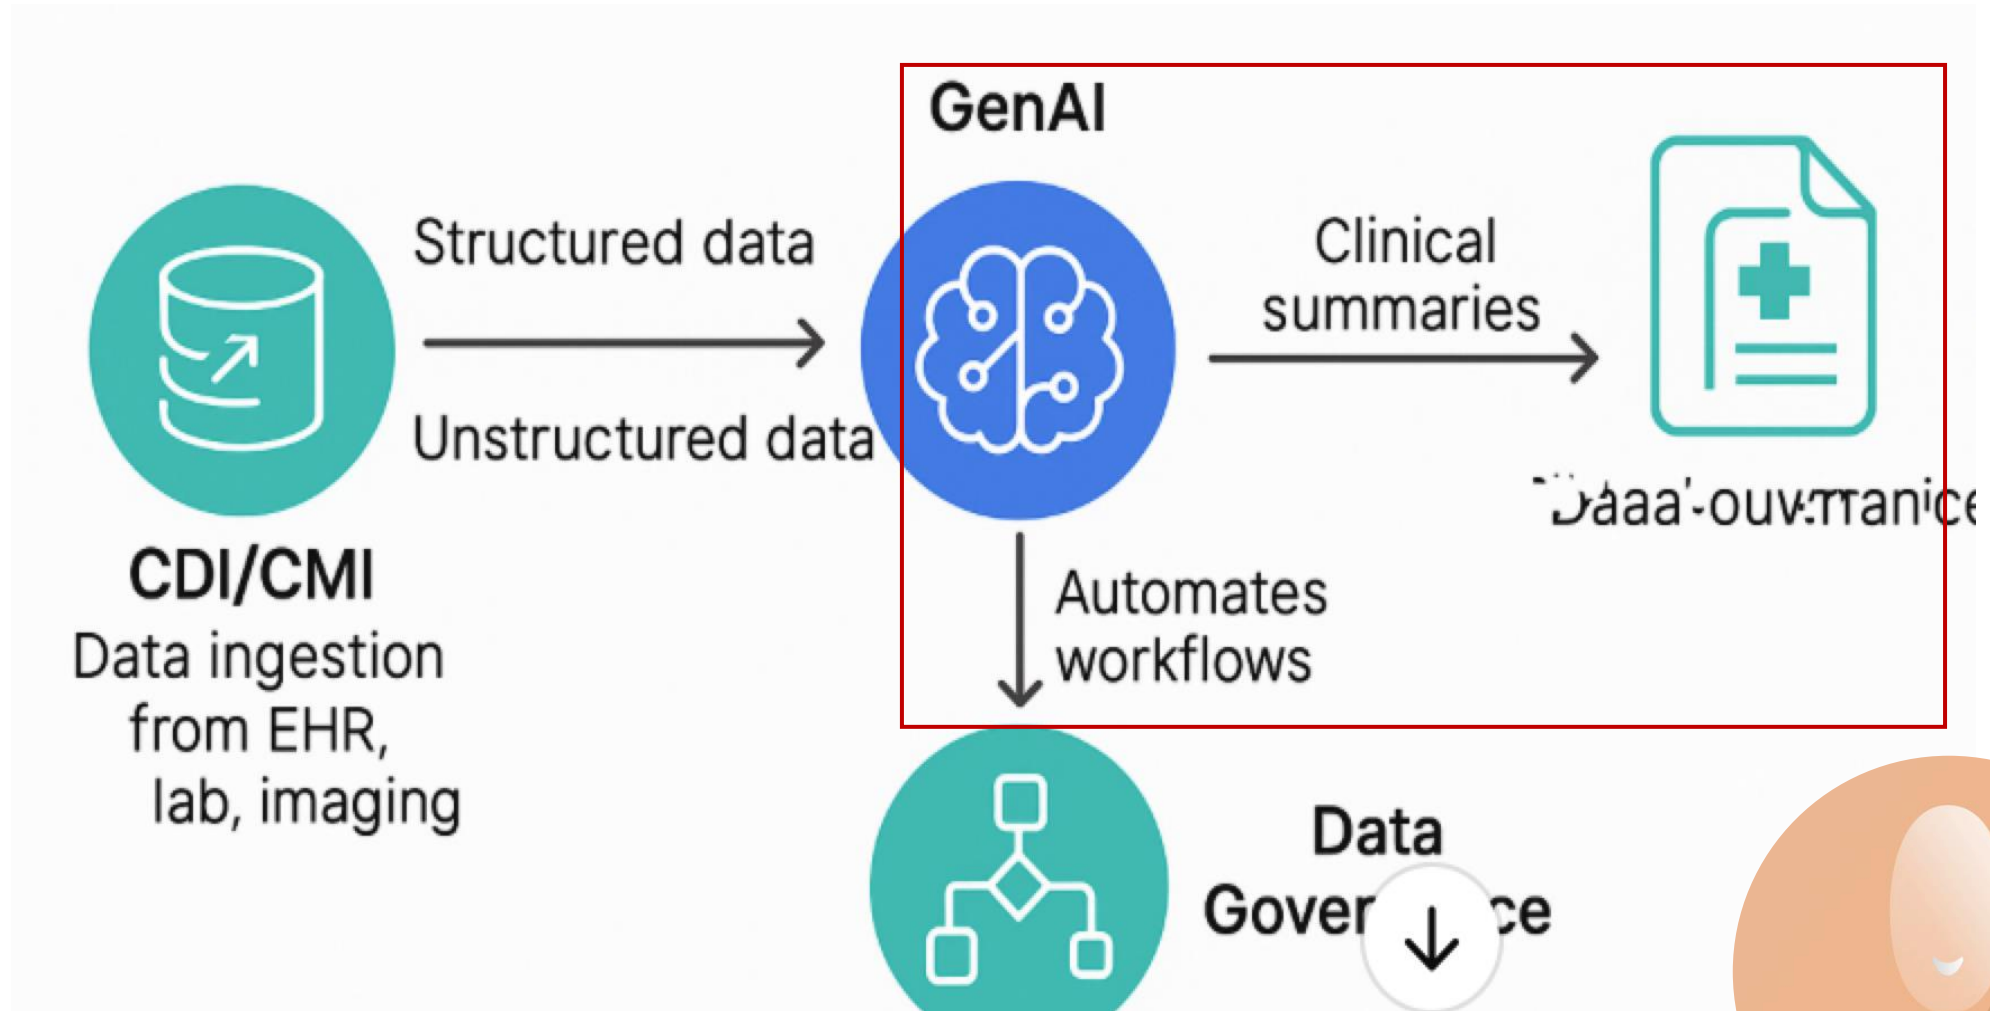


Query Pinecone using CAI

DEMO

Healthcare Use Case 1: Intelligent Clinical Documentation



Querying Pinecone

Get your Answers Fast: Harness EHR Insights with Pinecone.

Pinecone enhances your data retrieval through both semantic and lexical searches:

- Semantic Search: This is a powerful approach where queries are matched with vectors based on their semantic relevance rather than just keywords. It's perfect for understanding context and meaning within documents.
- Lexical Search: On the other hand, lexical search focuses on direct keyword matching, often valuable when you need precision on specific terms present within your dataset.

Steps involved:

- Context search using Vector Embeddings
- Generate Vector Embeddings using Open AI
- Formulate Pinecone Query
- Query Pinecone
- Use Open AI chat completion for the response

Utilizing Ingested Data with OpenAI

Get your Answers Fast: Harness EHR Insights with Pinecone.

Formulating Pinecone Queries:

We'll execute Pinecone queries utilizing both vector and lexical search strategies. This duality ensures comprehensive retrieval of relevant data based on semantic context as well as keyword precision.

- **Getting Answers using OpenAI Chat Completion**

The retrieved dataset is processed using OpenAI's chat completion feature, where AI interprets the data contextually, generating insightful answers to our queries. This involves harnessing AI capabilities to not just retrieve raw data points, but to synthesize them into meaningful insights.

Sample EHR-PDF File

Patient's clinical history:

Please also state the source of the information (e.g. from medical records, from the patient, from the applicant etc.).

Mr Tan Ah Kow was accompanied by his son, Mr Tan Ah Beng, for the examination. Mr Tan is a 55 year old man, who is divorced, and unemployed. Mr Tan is currently living with his son, Ah Beng, in Ah Beng's flat. Mr Tan Ah Beng informed me that Mr Tan Ah Kow used to work as a cleaner.

Mr Tan Ah Kow has a history of medical conditions. He has had hypertension and hyperlipidemia since 1990 and suffered several strokes in 2005. He subsequently developed heart problems (cardiomyopathy), cardiac failure and chronic renal disease and was treated in ABC Hospital.

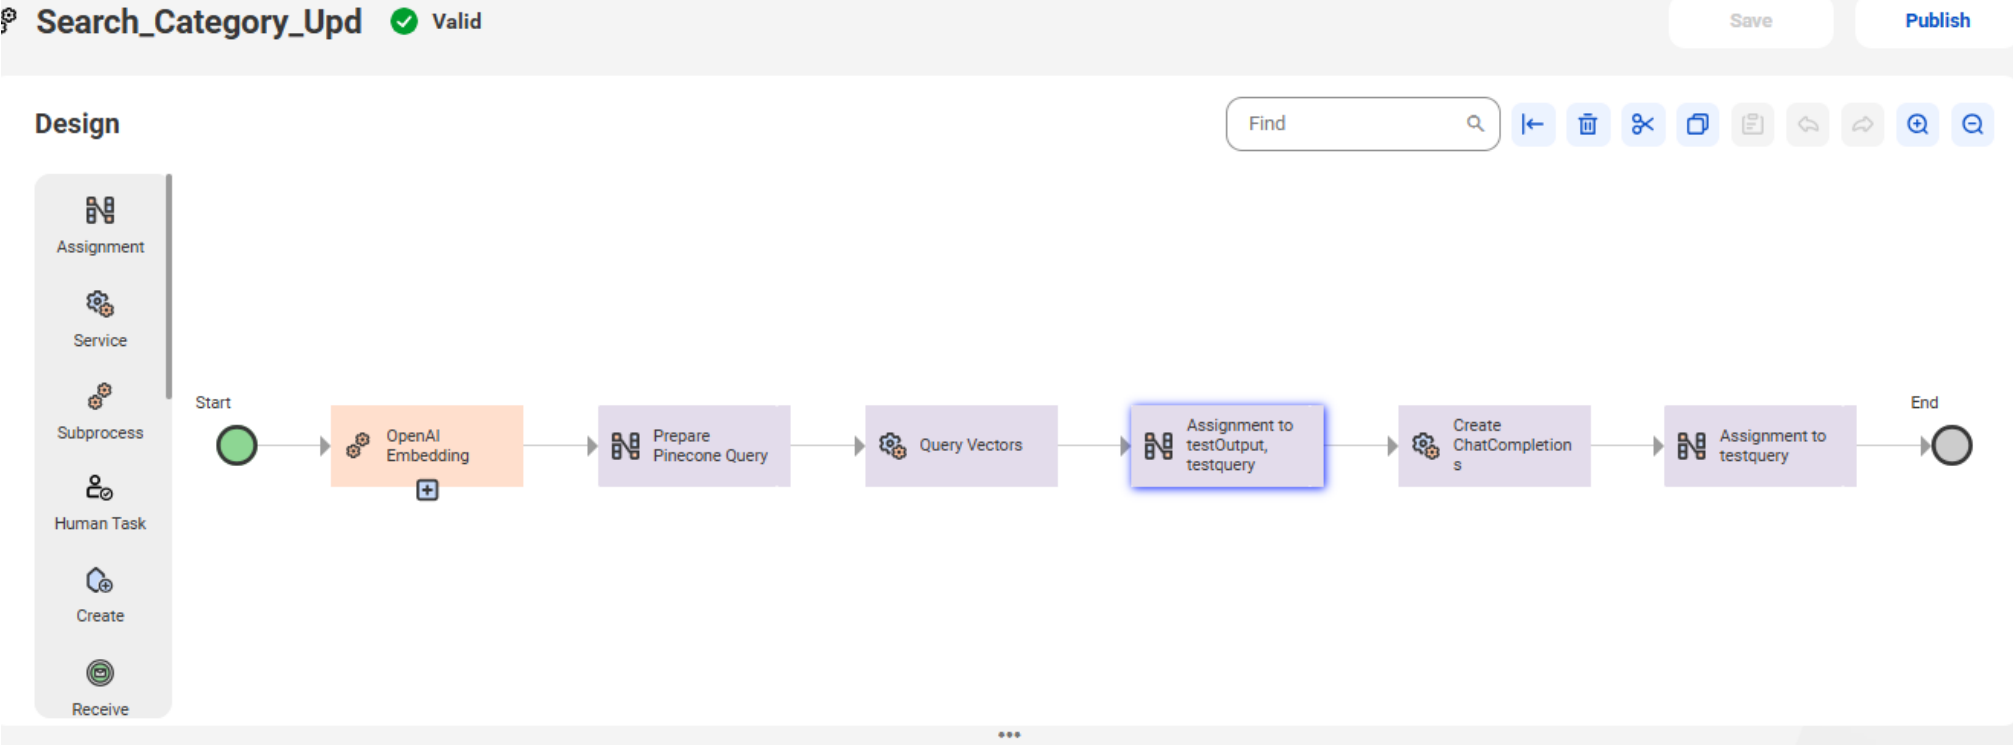
He was last admitted to the ABC Hospital on 1 April 2010 till 15 April 2010, during which he was diagnosed to have suffered from a stroke. This was confirmed by CT and MRI brain scans.

Thereafter, he was transferred to XYZ Hospital for stroke rehabilitation on 15 April 2010.

After that, Mr Tan was referred to Blackacre Hospital for follow-up treatment from in November 2010. The clinical impression was that he was manifesting behavioural and psychological symptoms secondary to Dementia.

Querying Pinecone

CAI Process Flow



Query 1

Get your Answers Fast: Harness EHR Insights with Pinecone.

The screenshot shows a REST client interface with the following components:

- Request Tab:** Shows the request body as raw JSON. The JSON is:

```
{  "context": "Give Patient Details that we are trying to Analyze",  "max_candidates": "100",}
```
- Response Tab:** Shows the response body as JSON. The JSON is:

```
{  "employment_status": "unemployed",  "medical_history": [    "Hypertension since 1990",    "Hyperlipidemia since 1990",    "Suffered several strokes in 2005",    "Heart problems (cardiomyopathy)",    "Cardiac failure",    "Chronic renal disease"  ],  "name": "Mr Tan Ah Kow",  "living_arrangement": "Living with son, Mr Tan Ah Beng",  "address": "Flat with son (incorrectly stated as Bedok, actual location Jurong)",  "current_condition": {    "cognitive_ability": "Gradual deterioration observed",  }
```
- UI Elements:** Includes tabs for 'Parameters', 'Authorization', 'Headers (9)', 'Body', 'Scripts', and 'Settings'. The 'Body' tab is active. Below the request, there are radio buttons for 'none', 'form-data', 'x-www-form-urlencoded', 'raw' (selected), 'binary', 'GraphQL', and a 'JSON' dropdown. The response tab also has a 'JSON' dropdown, 'Preview', and 'Visualize' options.

Query 2

Get your Answers Fast: Harness EHR Insights with Pinecone.

SECTION 2: DOCTOR'S PARTICULARS

Full name of doctor: Tan Ah Moi

NRIC/FIN/Passport no. of doctor: S2222222Z

MCR no. of doctor: 333333

Hospital / Clinic name and address: 1 Blackacre Hospital, Singapore 01010101

```
"context": "Give the doctor and clinic details",  
"max_candidates": "100",
```

cookies Headers (15) Test Results | ↻

Preview Visualize

```
],  
"testquery": [  
  {  
    "mcr_number": "333333",  
    "full_name": "Dr Tan Ah Moi",  
    "nric": "S2222222Z"  
  },  
  {  
    "name": "Blackacre Hospital",  
    "address": "1 Blackacre Hospital, Singapore 01010101"  
  }  
]
```

Query 3

Get your Answers Fast: Harness EHR Insights with Pinecone.

MEDICAL REPORT

SECTION 1: PATIENT'S PARTICULARS

Full name of patient: Mr Tan Ah Kow

NRIC/FIN/Passport no. of patient: S1111111X

Age of patient: 55 years old

The screenshot shows a REST client interface with the following details:

- Method:** POST
- URL:** `https://na1.ai.dm-us.informaticacloud.com/active-bpel/public/rt/0uYwTZkQeXajPU1GO0gUyF/S`
- Body:**

```
{
  "context": "Give Patient name and his age and particulars",
  "max_candidates": "100",
  "max_results": "50",
  "index_host": "ehr-pdf-ingest-uwavz9d.svc.aped-4627-b74a.pinecone.io",
  "namespace1": "ehr-pdf-ingest"
}
```
- Response:**

```
{
  "name": "Mr Tan Ah Kow",
  "NRIC": "S1111111X",
  "age": "55",
  "doctor": {
    "MCR_no": "333333",
    "name": "Dr Tan Ah Moi",
    "hospital": "1 Blackacre Hospital, Singapore 01010101",
    "NRIC": "S2222222Z"
  }
}
```

Query 4

```
"context": "Give Patient name and summary of the condition",  
"max_candidates": "100"
```

ies Headers (15) Test Results | ↻

200 OK • 8.42 s • 5.

▶ Preview Visualize | ⌵

```
"Summary": {  
  "Prognosis": "Unlikely to regain mental capacity; condition expected to worsen over time.",  
  "CurrentCondition": {  
    "MentalCapacity": {  
      "PropertyandAffairs": "Lacks capacity to make decisions regarding property and financial affairs.",  
      "PersonalWelfare": "Lacks capacity to make decisions regarding personal welfare."  
    },  
    "PhysicalCondition": "Incontinent, unable to bathe or use the toilet independently, but able to feed himself.",  
    "CognitiveFunction": "Gradual deterioration observed over the years."  
  },  
  "MedicalHistory": [  
    "Dementia",  
    "Stroke",  
    "Hypertension",  
  ]  
}
```



Anita Ayyagari

Thank You

Where data & AI come to **LIFE**





Shoot...

Q&A

Where data & AI come to **LIFE**