



27 May, 2025

Leveraging Informatica CAI to Orchestrate LLM Calls

- Sreehari Desai, Senior Solutions Architect, CSA
- Aracely Salazar, Associate Solutions Architect, CSA
- Coomar Das, Senior Principal Customer Success Architect, CSA

Where data & AI come to **LIFE**

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 27, 2025

Leveraging Informatica CAI to Orchestrate LLM Calls

Aracely Salazar, Sreehari Desai, Coomar Das

Customer Success Architects

Where data & AI come to 

Agenda

- ❑ **Recap of Episode 2**
- ❑ **CAI Recipes**
- ❑ **Demo: Empowering Gen AI Journeys with IDMC**
 - ❑ **Building AI-powered Chatbots**
- ❑ **Closing Summary: Key Learnings from our Gen AI Tech Tuesday Series**
- ❑ **Our vision for Agentic AI**
- ❑ **Q&A**



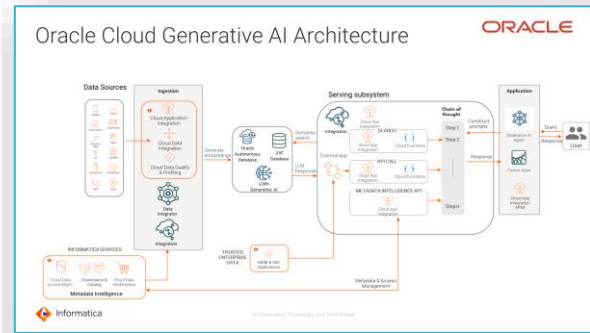
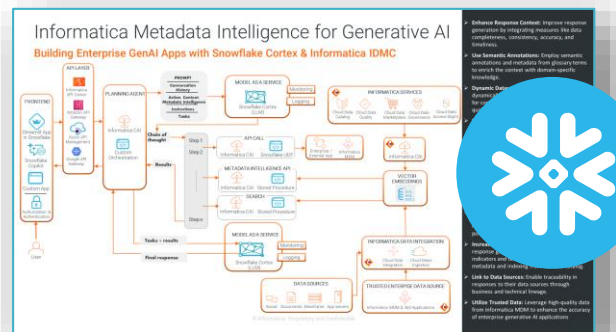
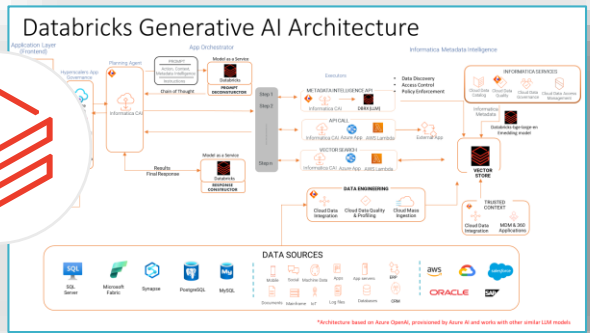
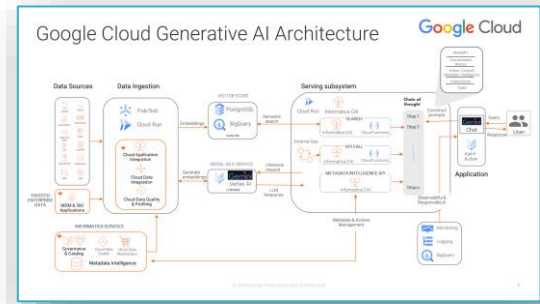
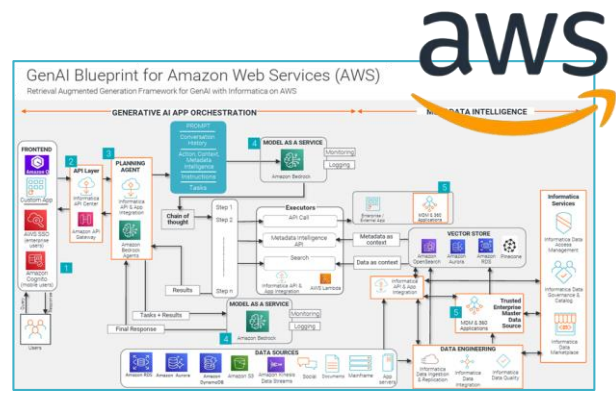
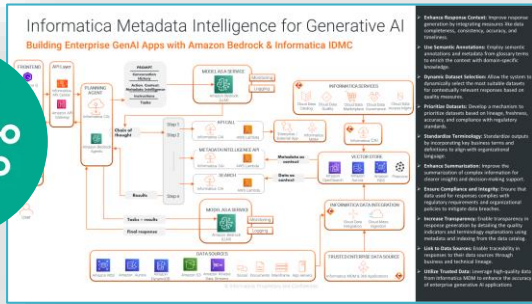
Episode 2 Recap

Where data & AI come to **LIFE**

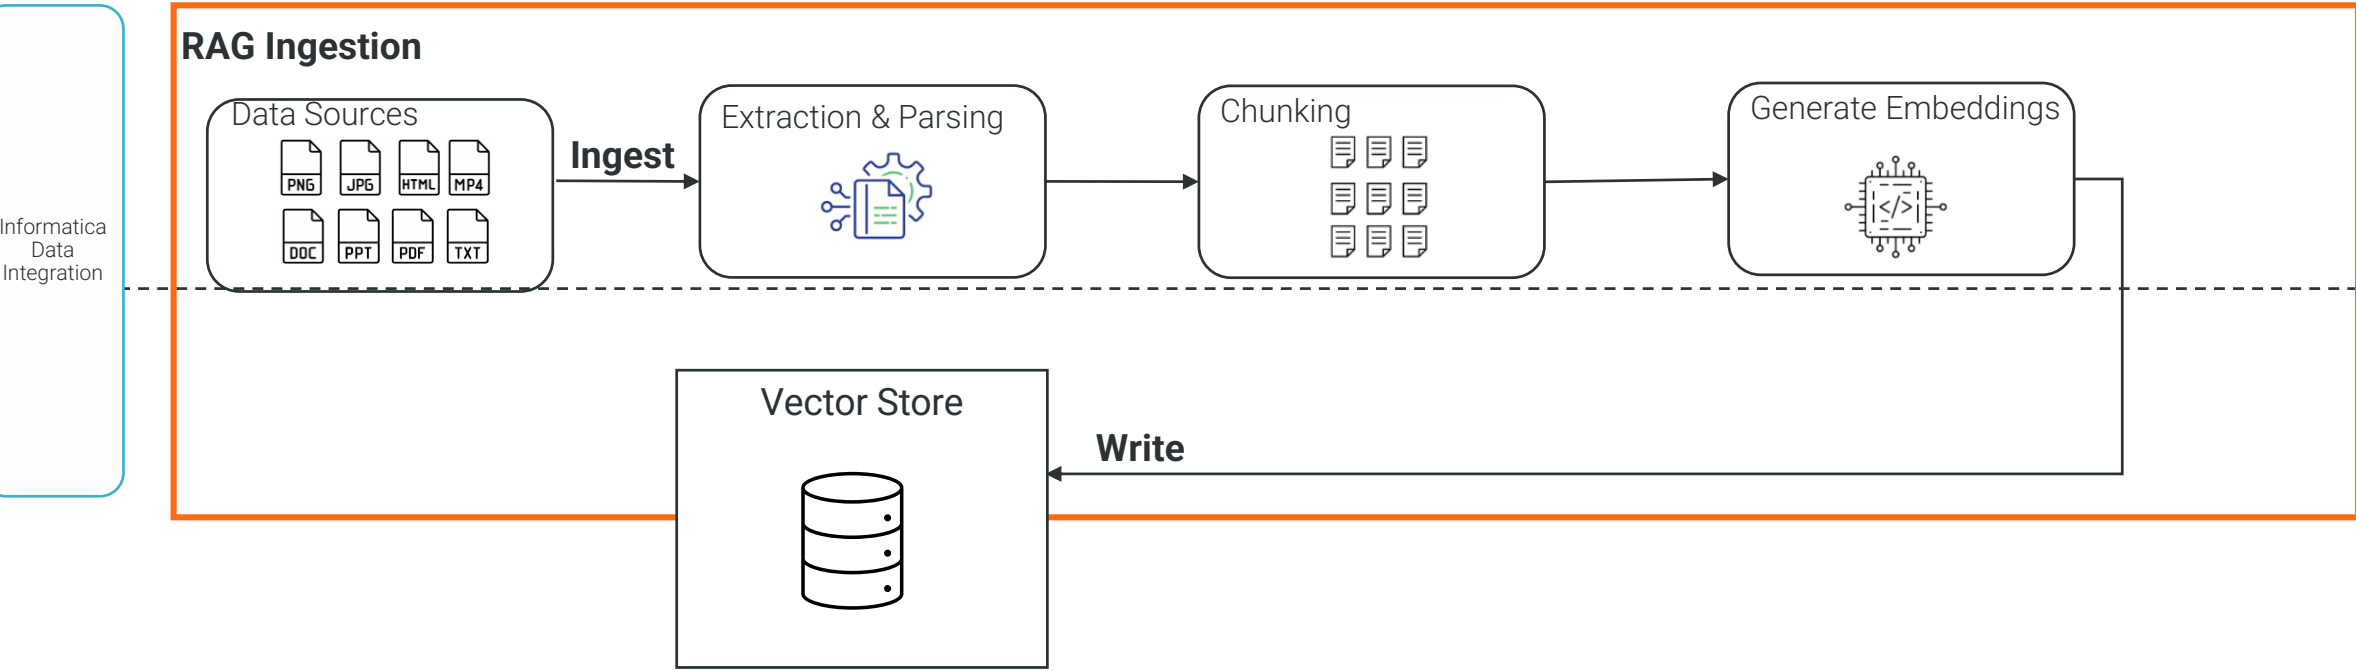
Stages of GenAI Implementation



Blueprints for Enterprise-Grade Gen AI

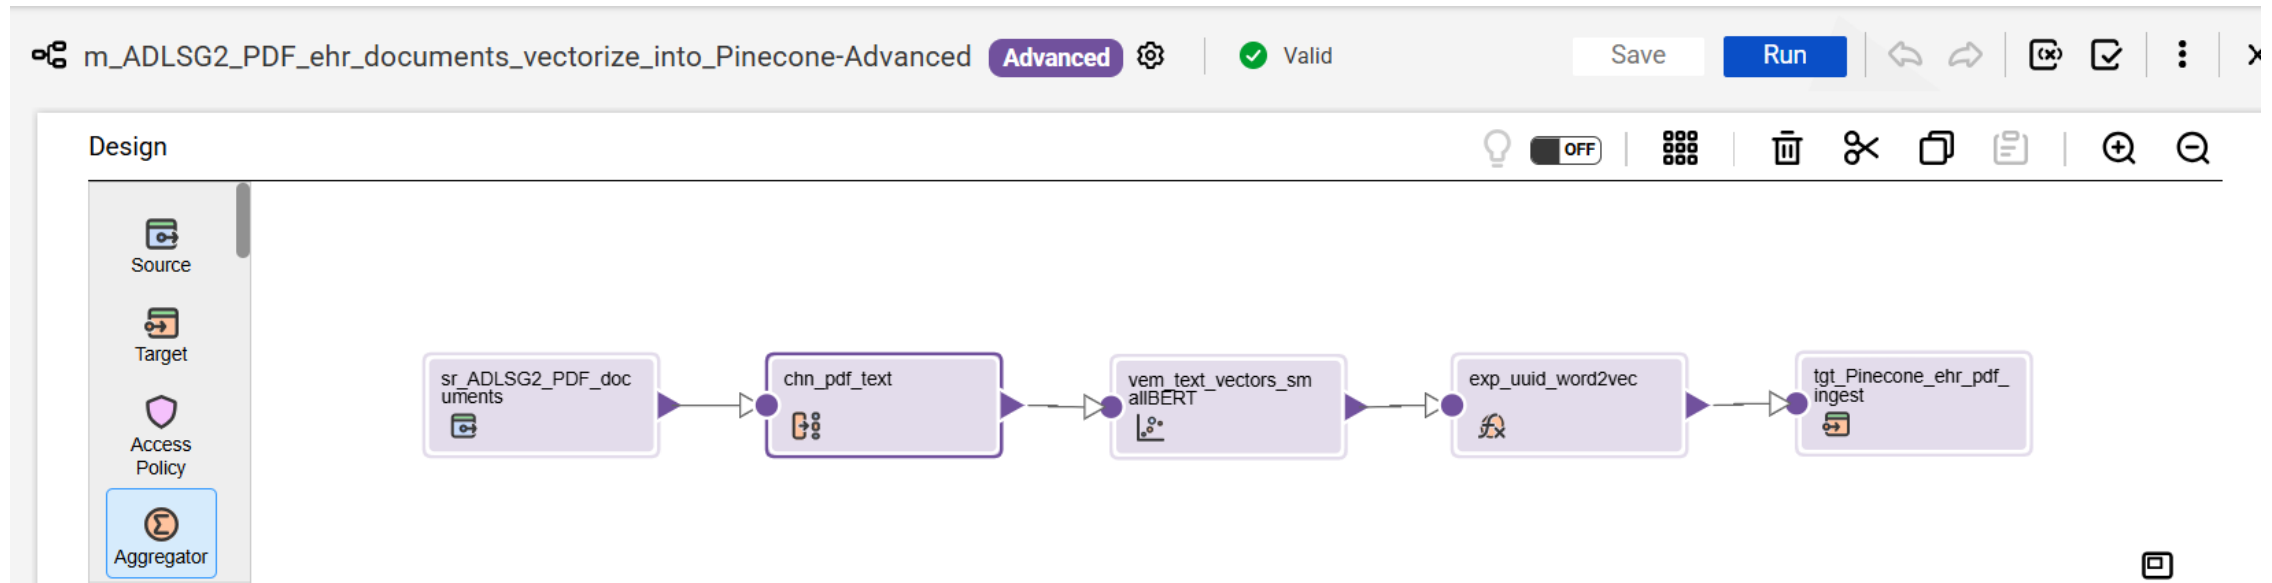


Informatica RAG Ingestion and AI Inference Architecture



Ingesting EHR files into Vector Database

CDI-Advanced Mapping



Querying Pinecone

Get your Answers Fast: Harness EHR Insights with Pinecone.

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

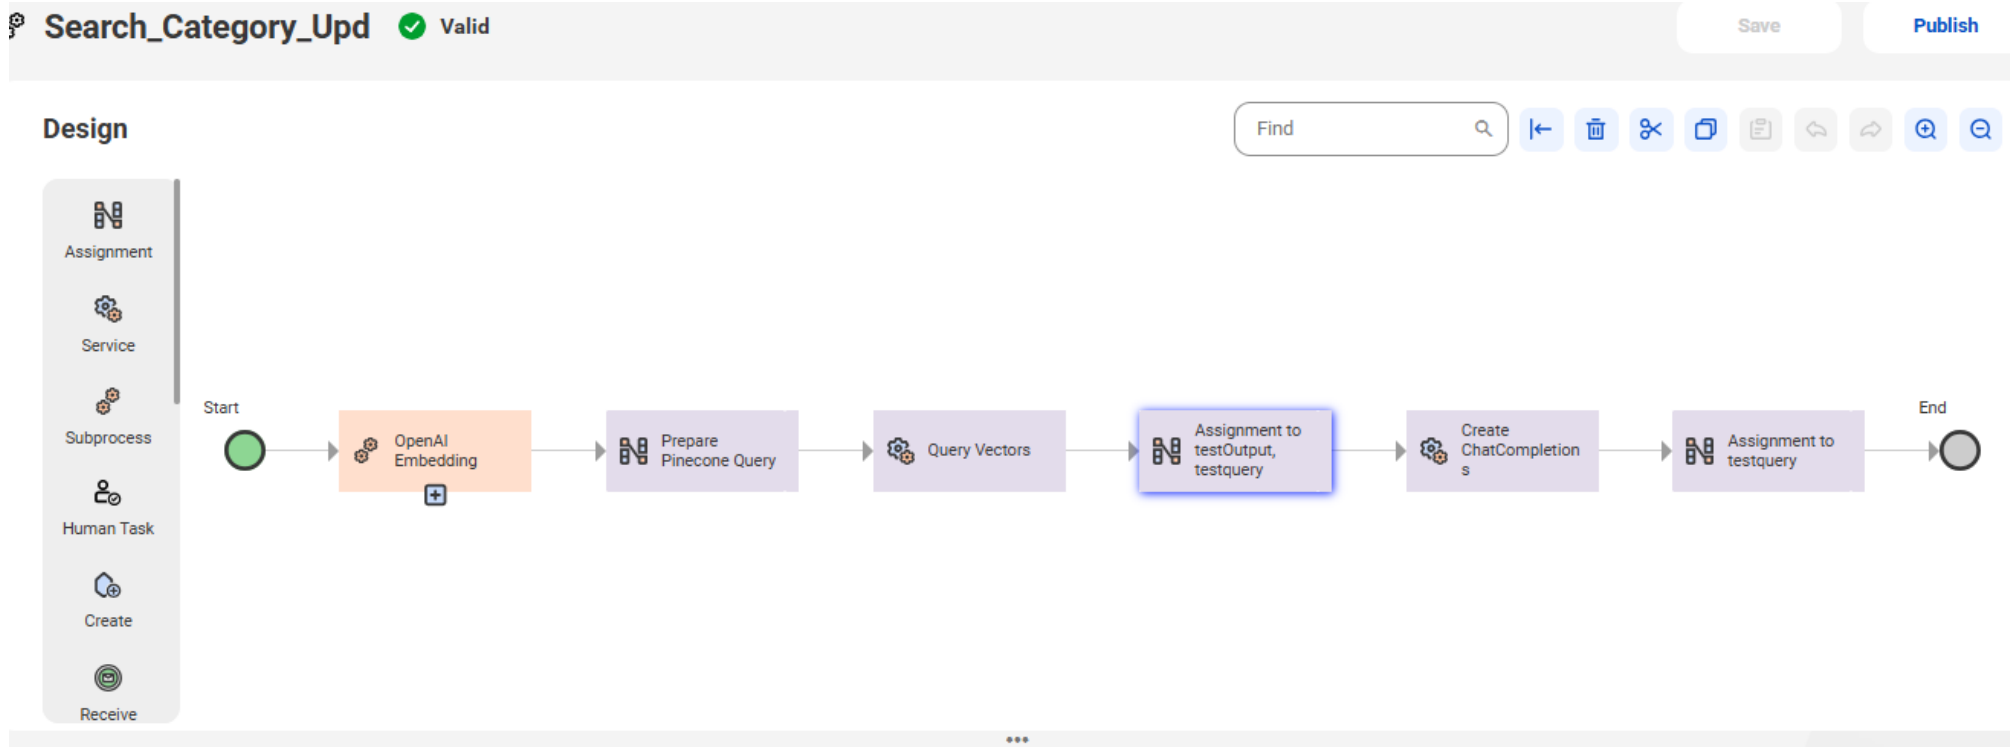
- Semantic Search: This is a powerful approach where queries are matched with vectors based on their semantic relevance rather than just keywords. It's perfect for understanding context and meaning within documents.
- Lexical Search: On the other hand, lexical search focuses on direct keyword matching,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

Querying Pinecone

CAI Process Flow





CAI Recipes

Where data & AI come to **LIFE**

Overview

- Introduction to CAI
- What Are CAI Recipes?
- Key Features of Informatica Cloud Recipes
- Benefits of Using Recipes
- Common Use Cases
- Example Recipe Overview
- Best Practices for Using Recipes

Introduction to CAI

Cloud Application Integration (CAI) ensures seamless connectivity between cloud and on-premise applications, enabling data flow across systems.

Top Use Cases:

- Process Orchestration
- App-to-App Integration
- Legacy App Modernization
- No-Code AI Applications
- API Implementation and Management

What Are CAI Recipes?

- CAI recipes are pre-built, reusable templates designed to simplify and accelerate common integration tasks.
- These recipes are ready-to-use, customizable, and tailored for specific use cases, making them highly efficient for businesses.
- They help organizations overcome integration challenges by providing a structured and tested approach.

Key Features of CAI Recipes

- Drag-and-drop functionality for ease of use
- Pre-configured workflows for common scenarios
- Support for diverse applications (e.g., Salesforce, Gemini, OpenAI)
- Scalability for enterprise needs
- Real-time and batch integration capabilities

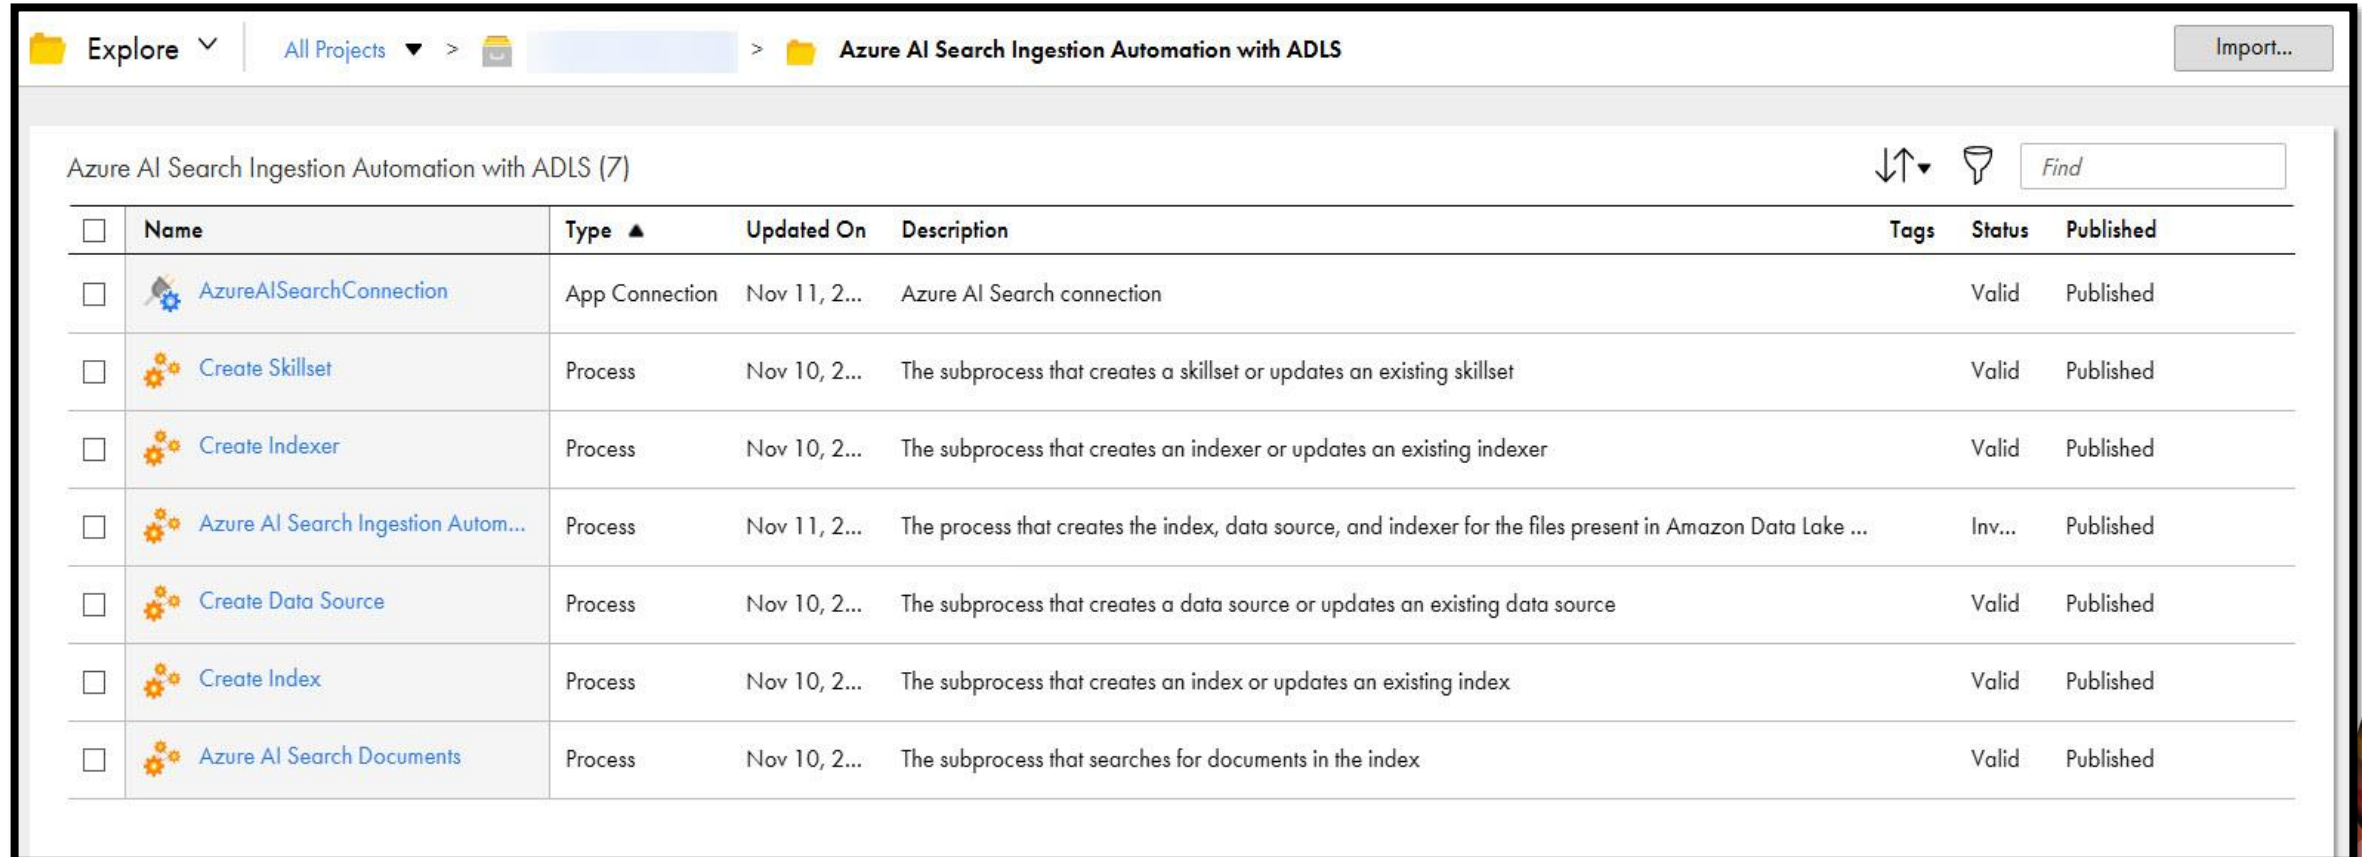
Benefits of Using Recipes

- Time-saving: Reduces development effort and time
- Ease of use: No extensive coding required
- Error reduction: Pre-tested templates minimize errors
- Flexibility: Easily customizable to fit specific business needs
- Cost-efficiency: Speeds up deployment, lowering costs








Common Use Cases

- Loan Processing
- Insurance Claim Processing
- Synchronizing customer data across CRM and ERP systems
- Real-time order processing and inventory updates
- Employee onboarding workflows (e.g., HR and payroll integration)
- Marketing campaign automation and analytics
- RAG consumption
- Prompt Engineering

Recipe Overview – Azure AI Search Ingestion Automation



The screenshot displays the Informatica Data Integration interface. At the top, there is a navigation bar with 'Explore', 'All Projects', and a breadcrumb path leading to 'Azure AI Search Ingestion Automation with ADLS'. An 'Import...' button is visible in the top right corner. Below the navigation bar, the main area shows a list of recipes under the heading 'Azure AI Search Ingestion Automation with ADLS (7)'. The list includes a search bar and sorting options. The recipes are listed in a table with columns for Name, Type, Updated On, Description, Tags, Status, and Published.

<input type="checkbox"/>	Name	Type ▲	Updated On	Description	Tags	Status	Published
<input type="checkbox"/>	 AzureAISearchConnection	App Connection	Nov 11, 2...	Azure AI Search connection		Valid	Published
<input type="checkbox"/>	 Create Skillset	Process	Nov 10, 2...	The subprocess that creates a skillset or updates an existing skillset		Valid	Published
<input type="checkbox"/>	 Create Indexer	Process	Nov 10, 2...	The subprocess that creates an indexer or updates an existing indexer		Valid	Published
<input type="checkbox"/>	 Azure AI Search Ingestion Autom...	Process	Nov 11, 2...	The process that creates the index, data source, and indexer for the files present in Amazon Data Lake ...		Inv...	Published
<input type="checkbox"/>	 Create Data Source	Process	Nov 10, 2...	The subprocess that creates a data source or updates an existing data source		Valid	Published
<input type="checkbox"/>	 Create Index	Process	Nov 10, 2...	The subprocess that creates an index or updates an existing index		Valid	Published
<input type="checkbox"/>	 Azure AI Search Documents	Process	Nov 10, 2...	The subprocess that searches for documents in the index		Valid	Published

Recipe Overview – Azure AI Search Ingestion Automation

- Purpose: Automate the creation of index, data source, skillset, and indexer for files ingested from ADLSG2
- Key functionality: Populate Azure AI Search index and retrieve documents meeting user search criteria
- Technology: REST API driven orchestration
- Prerequisite: Basic understanding of Azure AI Search connector concepts

Recipe Overview – Azure AI Search Ingestion Automation

- Components included
 - App Connection
 - Processes for creating index, data source, skillset, and indexer
- Automation flow for ingestion from ADLSG2

Azure AI Search Ingestion Automation with ADLS Recipe Assets

Asset types in the recipe package:

- Azure AI Search Connection (App Connection)
- Create Indexer (Process)
- Create Skillset (Process)
- Create Index (Process)
- Create Data Source (Process)
- Azure AI Search Documents (Process)
- Azure AI search ingestion Automation with ADLS (Process)

Azure AI Search Connector – Introduction

- Facilitates ingestion of data into Azure Search indexes
- Enables seamless integration with various data sources (Blob Storage, SQL DB, Cosmos DB, other applications)
- Supports use cases like e-commerce, content management, customer support knowledge bases

Understanding Indexes in Azure AI Search

- Indexes are data structures used to store and query searchable content
- Defined by fields with specific data types (e.g., String, Double, Complex Types)
- Field attributes include searchable, filterable, sortable, facetable
- Index keys uniquely identify each document within an index

Index Field Configuration Examples

Sample fields:

- HotelId (String, key, filterable)
- HotelName (String, searchable, sortable)
- Description (String, searchable, with language analyzer)
- Category (String, facetable, filterable)
- Tags (Collection of Strings, searchable, facetable, uses custom analyzer)
- Rating (Double, filterable, sortable, facetable)
- Address (Complex Type with subfields like City, searchable and filterable)

Summary of Azure AI Search

- Azure AI Search streamlines index creation and management
- Indexes enable rich and efficient search capabilities with flexible field configurations
- Integration with cognitive skills enhances searchable content quality
- API and SDK support allows full control over indexing and querying processes
- Informatica connectors in both CDI and CAI

Best Practices for Using Recipes

- Understand your integration requirements thoroughly
- Start with a pre-built recipe and customize incrementally
- Test integrations in a sandbox environment before deployment
- Leverage Informatica's support and community resources
- Regularly update recipes to align with evolving business needs

Resources and Further Reading

- Informatica Cloud Documentation
 - Getting Started with Recipes: <https://docs.informatica.com/integration-cloud/application-integration/current-version/getting-started-with-recipes/preface.html>
 - Azure AI Search CAI connector guide: <https://docs.informatica.com/integration-cloud/application-integration/current-version/azure-ai-search-connector-guide/preface.html>
 - Azure AI Search Ingestion Automation with ADLS: <https://docs.informatica.com/integration-cloud/application-integration/current-version/azure-ai-search-ingestion-automation-with-adls/preface.html>

Empowering Gen AI Journeys with IDMC

Where data & AI come to 

Querying your accumulated knowledge base – a classic Gen AI use case

- Knowledge Assets
 - Structured
 - OLTP Databases, OLAP Warehouses, Frontend CRM systems, Backend ERP systems,
 - Master Data Management Systems
 - Unstructured
 - Data Lakes, Lake houses, Data Swamps, SharePoint, Confluence Wiki
 - Internal, Proprietary, External, Public-facing
 - Website, Google Search, Perplexity, ChatGPT
- Search Service
 - Indexing
 - Querying

Azure AI Search

- Search Service
 - Indexing
 - Tokenize content (multimodal – text, images)
 - Store token in an inverted index in non-vector fields
 - Chunk and Vectorize content
 - Store vectors in vector fields in the same index
 - Querying
 - Hybrid searches
 - Lucene full text search
 - Vector search
 - Semantic ranking

Here's a question

“What are the authentication options with Azure Data Lake Storage connector?”

Should be easy with Informatica's Doc Portal, right?

The screenshot shows the Informatica Documentation search interface. The search query is "What are the authentication options with Azure Data Lake Storage connector?". The results are filtered by relevance and date, showing 1-10 of 2,066 results. The top result is "Understanding Technical Assets" dated 2025-05-14. Below it are three search results from Microsoft Azure, Oracle, and Google BigQuery, each with a snippet of text.

Product

<input type="checkbox"/> Address Verification (On-Premises)	24
<input type="checkbox"/> Administrator	101
<input type="checkbox"/> Application Integration	57
<input type="checkbox"/> Axon Data Governance	75
<input type="checkbox"/> B2B Data Exchange	450

Guide Title

<input type="checkbox"/> Command Reference	1,795
<input type="checkbox"/> User Guide	1,387
<input type="checkbox"/> Administrator Guide	661
<input type="checkbox"/> Developer Tool Guide	425
<input type="checkbox"/> What's New and Changed 1057	390

Version

<input type="checkbox"/> 10.5	1,407
<input type="checkbox"/> 10.5.3	1,189

Search Term: What are the authentication options with Azure Data Lake Storage connector? x Save Search Clear

Results 1-10 of 2,066 for What are the authentication options with Azure Data Lake Storage connector? in 4.13 seconds RELEVANCE DATE v

Understanding Technical Assets 2025-05-14
[Data Governance and Catalog Current Version](#)

Showing 1-5 of 5 results

Microsoft Azure Data Lake Storage Gen2
Data Governance and Catalog Current Version ... Data Lake Storage Gen2 Microsoft Azure Data Lake Storage Gen2 Microsoft Azure Data Lake ... The same agent can also process delimited files.

Microsoft Azure Synapse
Data Governance and Catalog Current Version ... Microsoft Azure Synapse Microsoft Azure Synapse Microsoft Azure Synapse Data ... Stored Procedures linked to tables or views

Oracle
Data Governance and Catalog Current Version ... created in Amazon S3, Microsoft Azure Data Lake, and Oracle Cloud Object Storage source systems. ... Stored Procedures linked to tables or views

Tableau
Data Governance and Catalog Current Version ... Site Stored Procedure Data Lineage ... Amazon S3 Microsoft Azure Data Lake Storage Gen2 Microsoft Azure SQL Server Oracle

Google BigQuery
Data Governance and Catalog Current Version ... The data profiling task runs profiles on the following data types: ... TIMESTAMP Data Lineage The following lineage data is available for

Show Only Top Results ^

What does Gemini say? Pretty good, but...

Google

What are the authentication options with Informatica's Azure Data Lake Storage connector? X [voice icon] [refresh icon] [search icon]

All Videos Images Short videos Forums Shopping Web More Tools

Did you mean: What are the authentication options with **Informatica** Azure Data Lake Storage connector?

◆ AI Overview Learn more

Informatica's Azure Data Lake Storage (ADLS) connector supports several authentication options: **Account Key**, **Service Principal**, **Managed Identity**, and **Shared Access Signature (SAS)**. These methods provide varying levels of security and control, allowing users to choose the most appropriate option for their needs.

Here's a more detailed look at each method:

- 1. Account Key:**
This method uses the primary or secondary access key for your ADLS Gen2 account.
- 2. Service Principal:**
This approach involves using a registered application in Azure Active Directory (Azure AD) to authenticate.
- 3. Managed Identity:**
This allows your Informatica Cloud Service (ICS) to access ADLS Gen2 using a managed identity associated with the service, which simplifies the authentication process.
- 4. Shared Access Signature (SAS):**
This method uses a unique signature (token) that grants specific permissions and access to a container within your ADLS Gen2 account.

AI responses may include mistakes. [Export] [Save] [Like] [Share]

Set up your Azure Data Lake Storage Gen2 connection
Sep 16, 2024 — Step 2: Select and set your authentication. Under...
Learn Microsoft

Access control model for Azure Data Lake Storage - Learn Microsoft
Data Lake Storage supports the following authorization mechanisms...
Learn Microsoft

Okay.. How about ChatGPT?

ChatGPT ▾

What are the authentication options with Informatica's Azure Data Lake Storage connector?

Informatica's Azure Data Lake Storage (ADLS) connector supports several authentication options for securely connecting to Azure Data Lake Storage Gen1 and Gen2. The available options may depend on the version of Informatica (e.g., Cloud Data Integration vs. PowerCenter) and the connector version, but generally include the following:

For Azure Data Lake Storage Gen2 (ADLS Gen2):

1. Azure Active Directory (Azure AD) Authentication (Recommended)

This is the most secure and commonly used method, especially for enterprise setups.

- **Service Principal with Client Secret:**
 - Authenticate using an Azure AD application (service principal) with:
 - Tenant ID
 - Client ID (Application ID)
 - Client Secret
 - Requires the service principal to have the necessary permissions on the storage account or containers via Azure RBAC.
- **Service Principal with Certificate:**
 - Similar to the above, but uses a certificate instead of a client secret for added security.

2. Managed Identity (for Azure-hosted Informatica instances)

- If you are running Informatica Intelligent Cloud Services (IICS) on Azure:

Ask anything

+ Tools

🔊

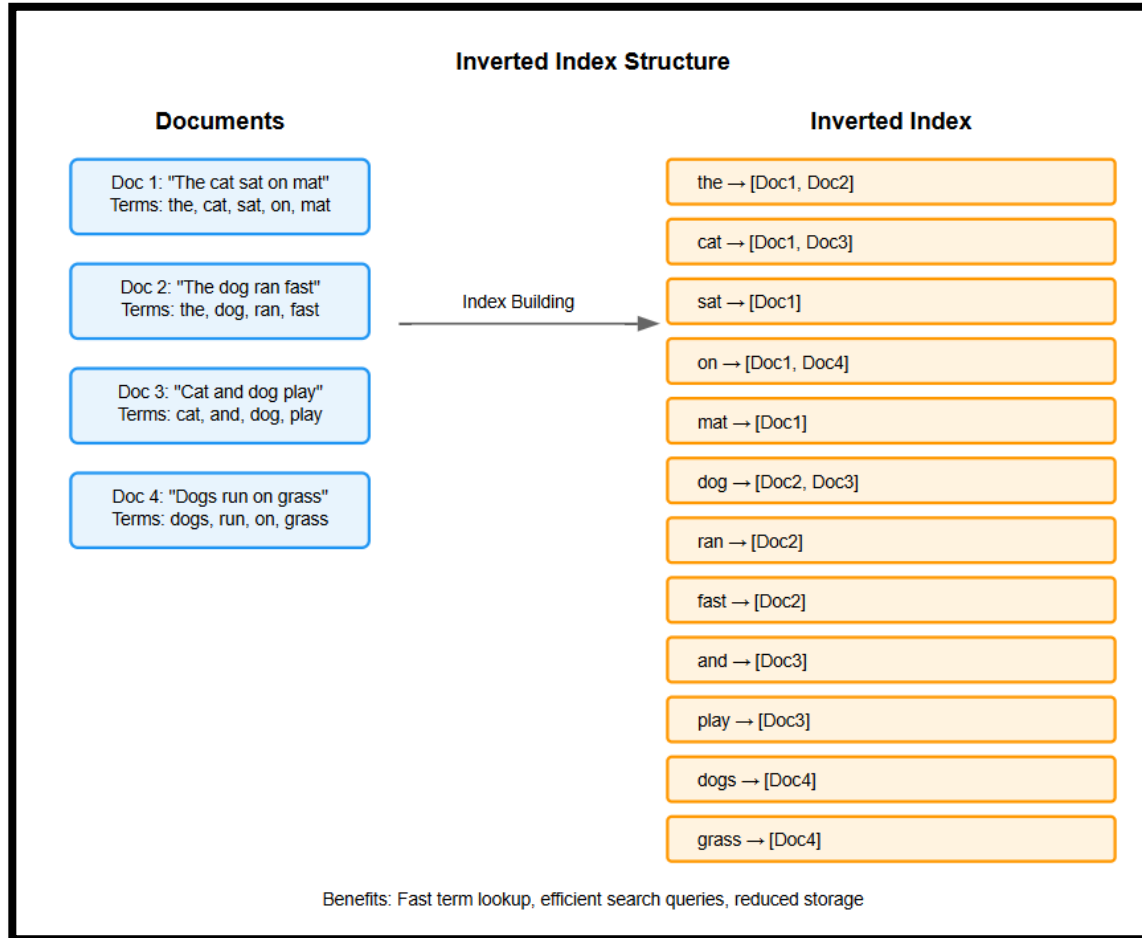
They are almost there...

But can we do better?

From data to dialogue – let's build a knowledge base bot

- Components
 - Search Service
 - Azure AI Search
 - Data Source, Index, Skillset, Indexer
 - Semantic Ranker, Vectorizer
 - Agentic Search
 - Search Index
 - Knowledge Agent (LLM)
 - Agentic Retrieval Pipeline
 - LLM (Response generation, grounded in retrieved, perhaps proprietary, data)
 - User chat interface

Lucene Inverted Index



INDEX

<i>A Memoir of Jane Austen</i> (see Austen Leigh, James Edward)	Bage, Robert, 58
Aristotelianism, 16, 18	<i>Hermesprong, or Man as He Is Not</i> , 58
Armstrong, Nancy, 114	Bal, Mieke, 9
attitudes to <i>Pride and Prejudice</i>	Barbault, Anna Laetitia, 1
Austen family, 52, 66, 138	Bautz, Annika, 126
contemporary, 52, 53, 54, 70, 96, 125, 138	Behn, Aphra, 2
Jane Austen's view, 96, 137	Bellos, David, 128
later, 129, 140	Bennett, Arnold, 175
Auerbach, Emily, 175, 181	Bermingham, Ann, 98
Austen, Anna, 44, 112	Besterman, Theodore, 47
Austen, Caroline, 4, 46	'Biographical Notice of the Author' (see Austen, Henry)

So many skills

Custom
Azure Machine Learning (AML)
Custom Web API Skill
Custom Web API Skill - Azure Functions
Vision
Image Analysis Skill
OCR Skill

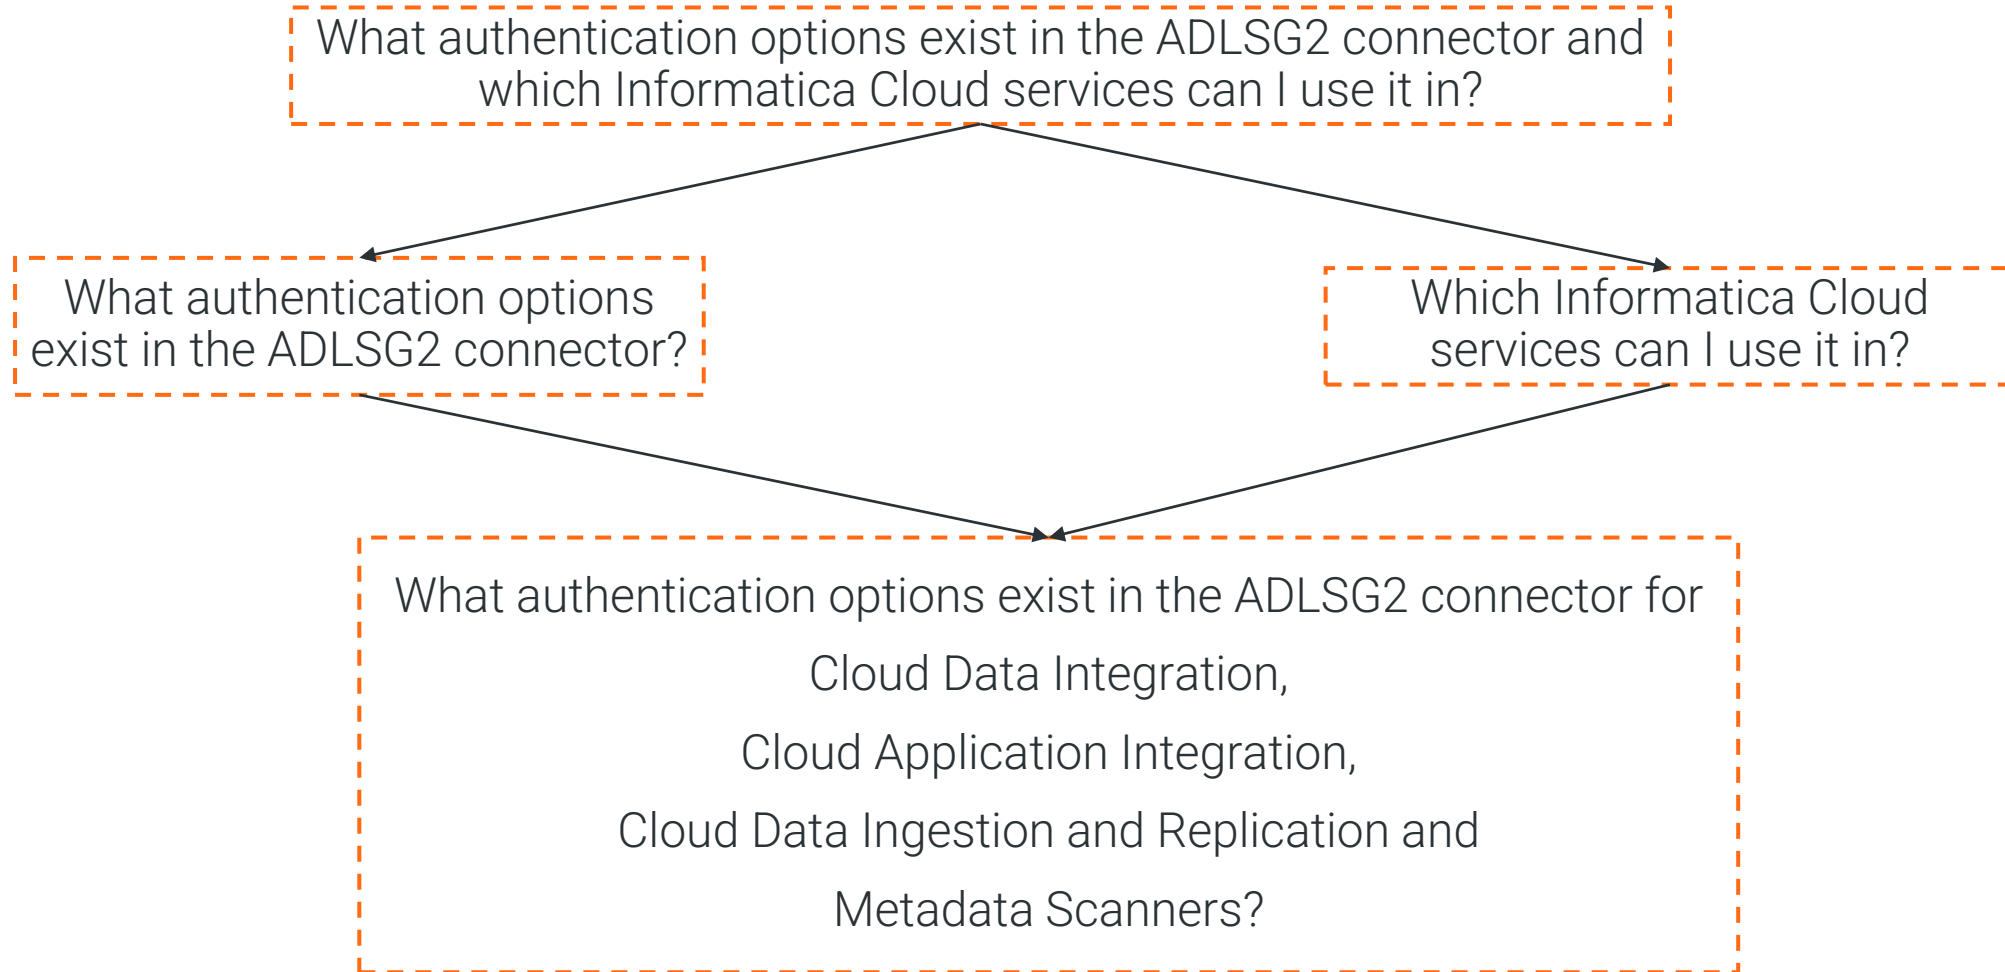
Text
Entity Recognition Skill (V3)
Entity Linking Skill (V3)
Key Phrase Extraction Skill
Language Detection Skill
Merge Skill
Split Skill
Sentiment Skill (V3)
Translation Skill
Custom Entity Lookup Skill
PII Detection Skill
Azure OpenAI Embedding Skill

Util
Shaper Skill
Conditional Skill
Document Extraction Skill
Document Intelligence Layout Skill (preview)
Vectorize
Text Vectorize Skill (preview)
Image Vectorize Skill (preview)

Agentic Retrieval Pipeline? Wait, isn't that just... RAG?

- Agentic Retrieval Pipeline
 - User Query → Knowledge Agent (LLM)
 - → Azure AI Search Service Index
 - (Lucene search, Vector search, Hybrid search)
 - Agentic Search (LLM)
 - Compound query breakdown into subqueries
 - Query rewrite, Query expansion
 - All in parallel
 - Results
 - Semantic ranking
 - Merge

Query breakdown and rewrite



Semantic Ranking

What authentication options exist in the ADLSG2 connector and what **metadata** can it **extract**?

Rank 1

ADLSG2 CDGC scanner

Can **extract** workbooks, worksheets, columns from Excel

Can **extract** elements and attributes from XML

Rank 2

CDI ADLSG2 connector

Can connect to ADLSG2 to read/write files and objects

Rank 3

CAI ADLSG2 connector

Can write to ADLSG2 container

But first, let's get our data loaded into a data source

- Objective
 - Search chatbot
 - Could be on proprietary or public-facing datasets
 - Structured or Unstructured
 - We will use Informatica Cloud Services product documents as our sample
 - Unstructured PDF files
 - We will ask it the same question about the authentication options in the ADLSG2 connector
- Steps
 - Use Cloud Data Integration to bulk load our documents to a Data Lake container
 - Use Cloud Application Integration for all subsequent programmatic actions on Azure AI Search, Knowledge Agent Retrieval actions and overall LLM orchestration
 - Test our chat interface

Let's cover some more terminology as the data load runs

- Chunks
- Vectors
- Embedding
- Cosine similarity
- Vectorizer

Chunking

The Azure Data Lake Storage Connector has three authentication options.
These are Managed Identity, SAS key and SPN.

The Azure Data Lake Storage Connector has three authentication options.

These are Managed Identity, SAS key and SPN.

...Azure Data Lake Storage Connector...

Managed Identity

SAS key

...three authentication options...

SPN

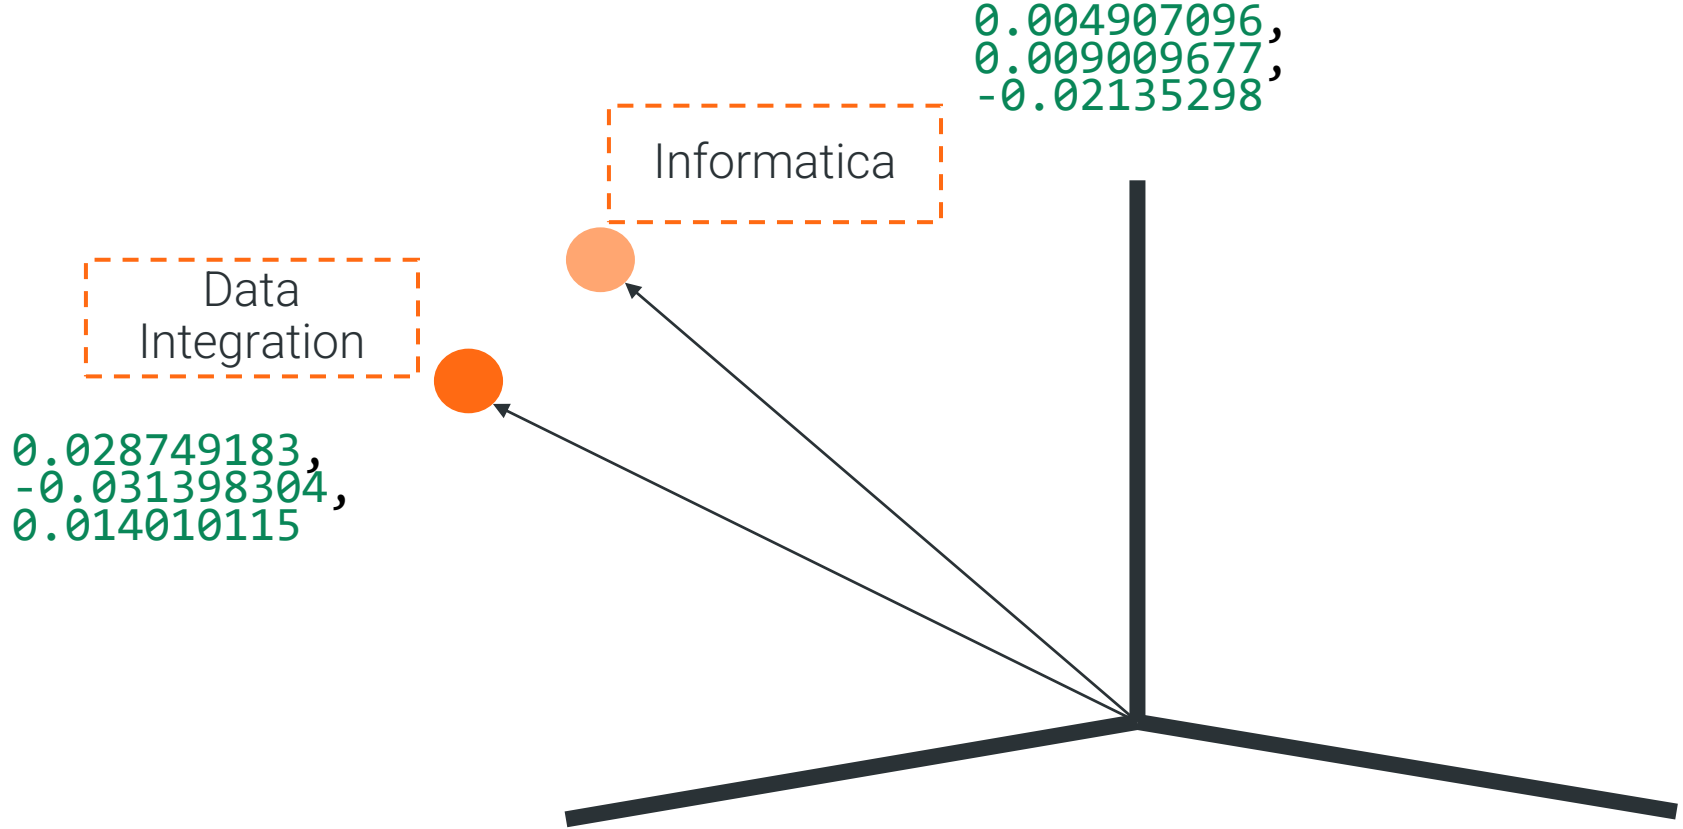
auth

ent

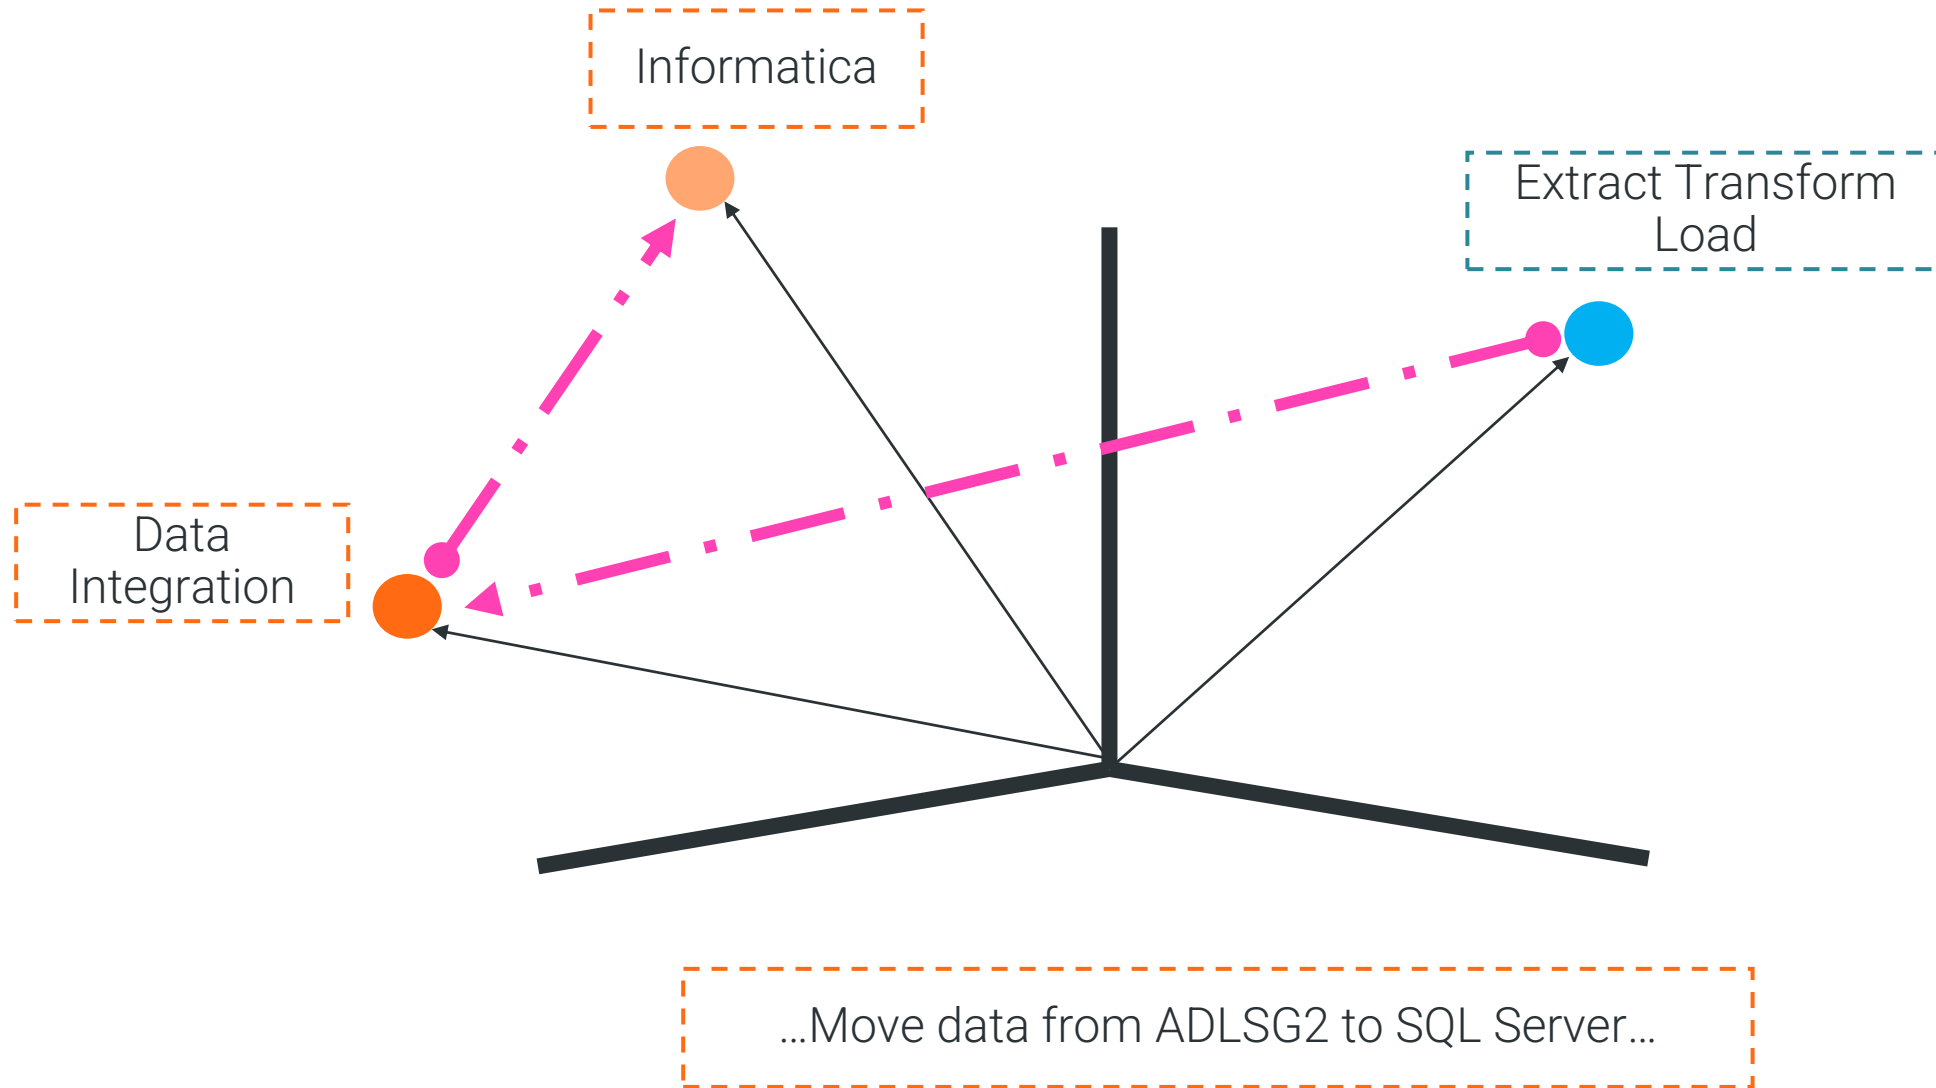
tion

ca

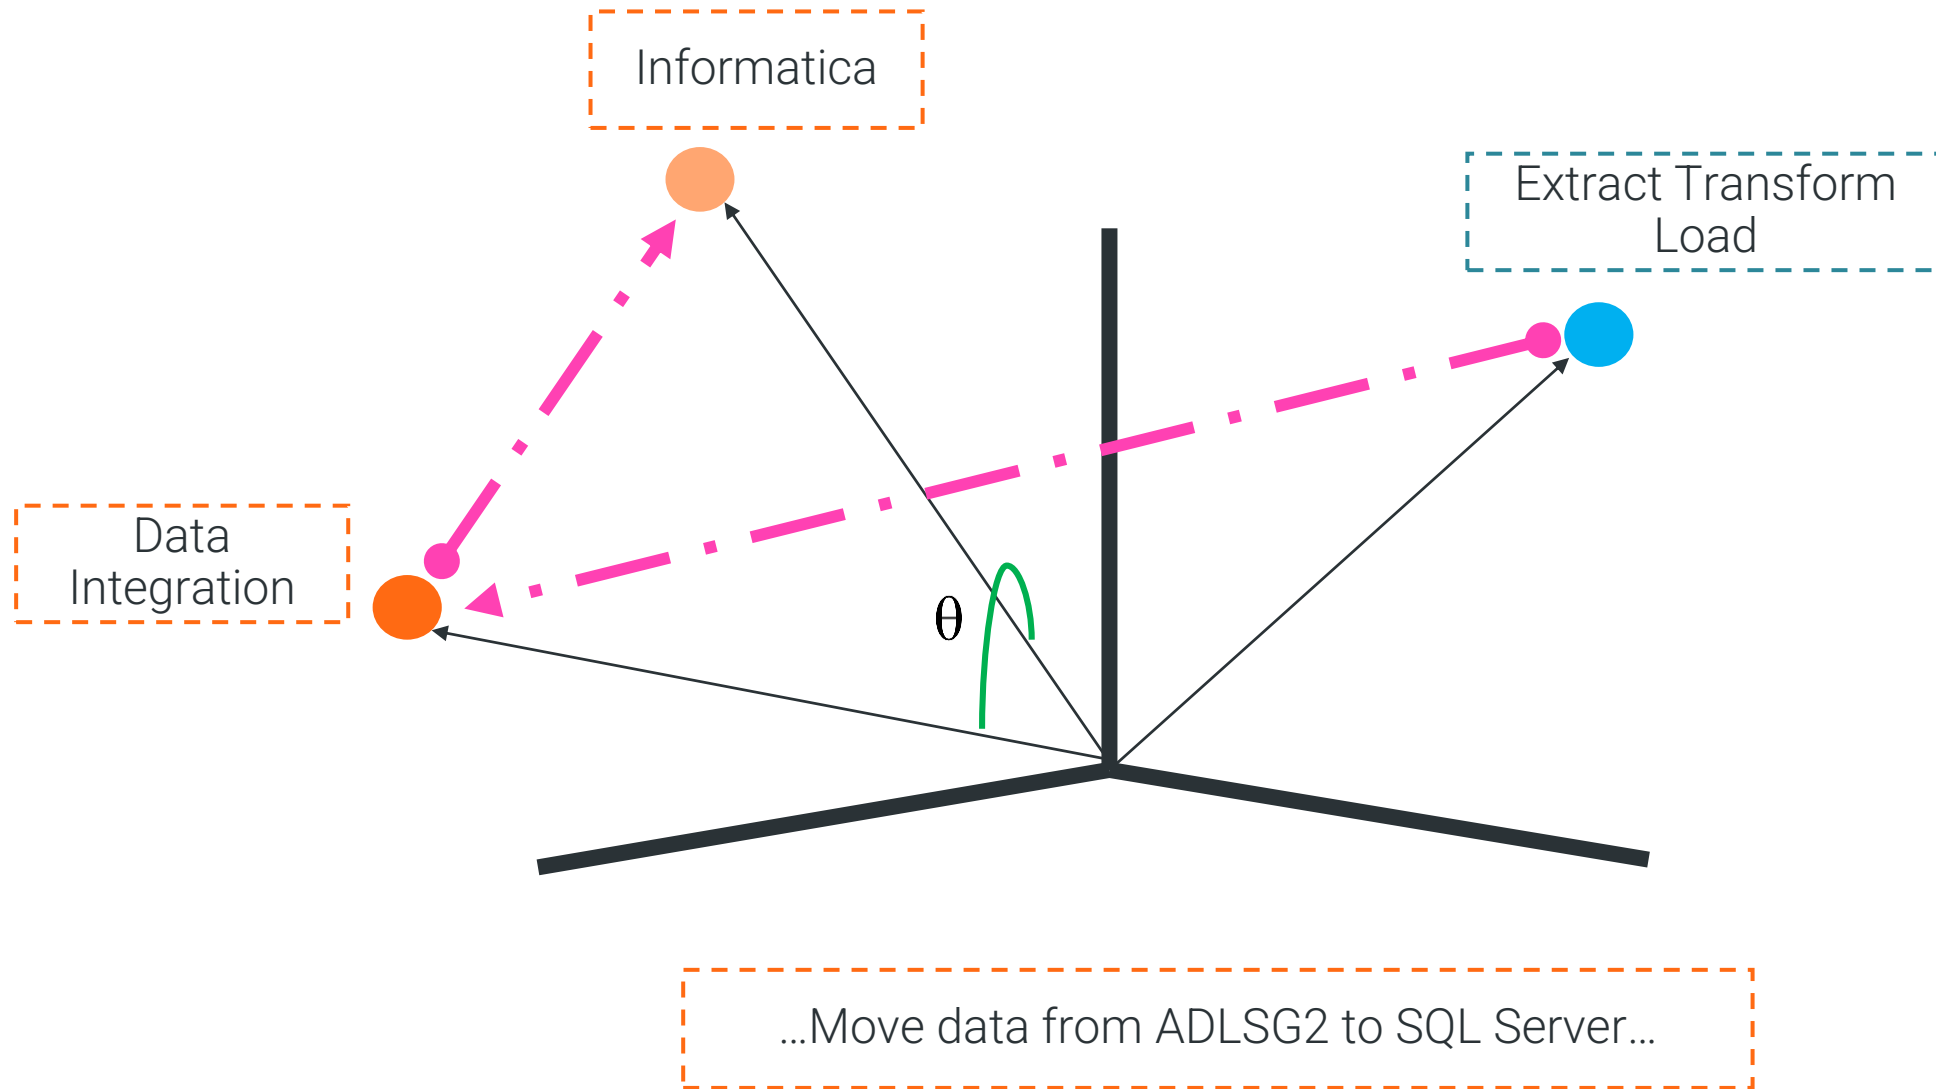
Vectors



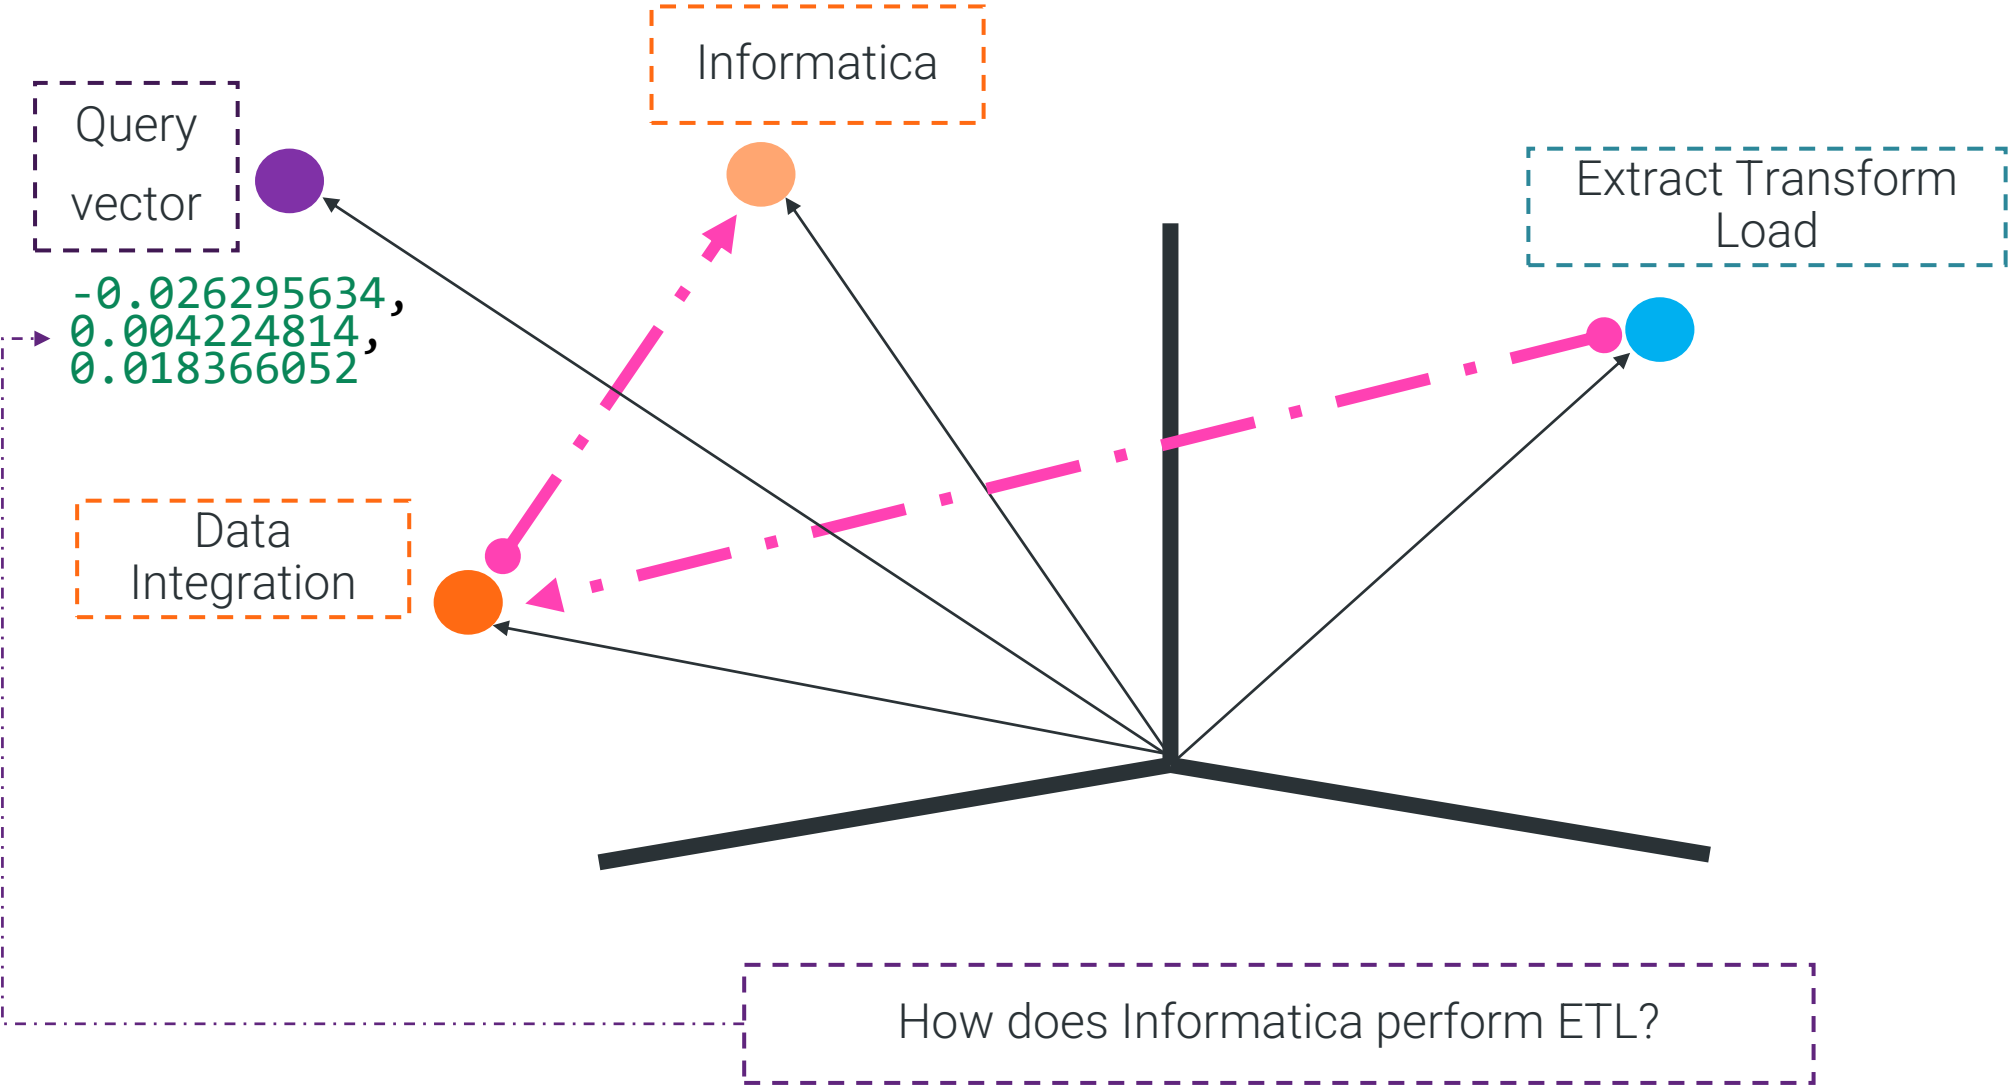
Embedding



Cosine Similarity



Vectorizer



Okay.. Let's keep building...

- Azure AI Search – Data source
 - Let's inform Azure AI Search where the raw data resides, so we can index it
- Azure AI Search – Index
 - This is where all our searchable data will be indexed, ready for vector searches (or hybrid ones)
- Azure AI Search – Skillset
 - Want to ingest images for search? Extract text? Split? Merge?
 - How about Chunking? Embedding?
 - Document Layout? Markdown? OCR to recognize text in images?
 - Language translation? Key phrase detection? Entity detection?
- Azure AI Search – Indexer
 - Let's connect the dots – load up our index from the data source per the strategies (skills) in the skillset

Looking good so far...

- Our Azure AI Search index is ready to search
 - We could stop here and
 - Do text searches
 - We could do similarity searches on the vectors, but that would take converting our search text into vectors first
- Can we automate that?
 - Yes. Yes, we can! This is where our Knowledge Agent comes in
 - We will use OpenAI GPT-4o-mini (hosted on Azure) for our Agentic Retrieval Knowledge Agent

Knowledge Agent built.. Now what?

- Let's point the Agent to our AI Search Index and ask it our "test" question
 - No, this is NOT our chatbot – hang on, that's coming!
 - We will review the Knowledge Agent's response, but more importantly
 - We will review what activities it performed
 - Query planning
 - Azure Search Query
 - Semantic Ranking
 - And critically, its References!!

Those references sure seem “grounded” in truth

- ... why do we need a Chatbot again? Isn't this good enough?
- It is pretty good –
 - in fact, we could persist this data,
 - some of which was previously not searchable in their raw form (think text on images, as an example),
 - which has now become structured and searchable
 - through enrichment, analysis and inference (Skills in Skillsets)
 - in a Knowledge Store (Blob container, Table store)
 - for further exploration
- But we could also ask an LLM to use it as a reference and provide a more accurate answer to our original question, without retraining it on our data set
- Let's do that...

Let's RAG and ROLL

- We will once again use GPT-4o-mini as our Chatbot's LLM except
 - We will put some guardrails and ask it to ONLY refer to the data retrieved by the Knowledge Agent
 - We will provide it the content the Knowledge Agent retrieved
 - We will point it to our Azure Search AI index and inform it of the vector fields to refer to
 - And some other parameters on top N documents and strictness
- No data will leave our data source
- We will also put a Flask proxy app on top
 - that connects the Chatbot's LLM to the Chatbot's frontend, which we will host on a web server locally
- And then, we will ask our question again!

Better, right?

The screenshot shows a web browser window with the URL `localhost:3080/Chatbot.html`. The chatbot interface has a purple header with the text "Let's RAG and ROLL" and two key icons. Below the header is the question "How may I help you today?". A text input field contains the question: "What are the authentication options with Azure Data Lake Storage connector?". A purple "Send" button is below the input field. The chatbot's response is displayed in a light gray box and reads: "What are the authentication options with Azure Data Lake Storage connector?: The Azure Data Lake Storage connector offers three authentication options: 1. **Service Principal Authentication:** This method uses the client ID, client secret, and tenant ID of an application registered in Azure Active Directory to connect to Azure Data Lake Storage Gen2. The required properties include: - **Client ID:** The application ID for your registered application. - **Client Secret:** The secret key generated for the client ID. - **Tenant ID:** The directory ID of the Azure Active Directory [doc1](#). 2. **Shared Key Authentication:** This method uses the storage account name and account key to connect to Azure Data Lake Storage Gen2. The required properties include: - **ADLS Gen2 Storage Account Name:** The name of the storage account. - **ADLS Gen2 Account Key:** The access key for the storage account. - **File System Name:** The name of the file system in the storage account [doc2](#). 3. **Managed Identity Authentication:** This method authenticates using identities assigned to applications in Azure. It can be either system-assigned or user-assigned. The required properties include: - **Client ID:** The client ID for the user-assigned managed identity (leave empty for system-assigned). - **ADLS Gen2 Storage Account Name:** The name of the storage account. - **File System Name:** The name of the file system in the storage account [doc3](#).

Imagine the possibilities

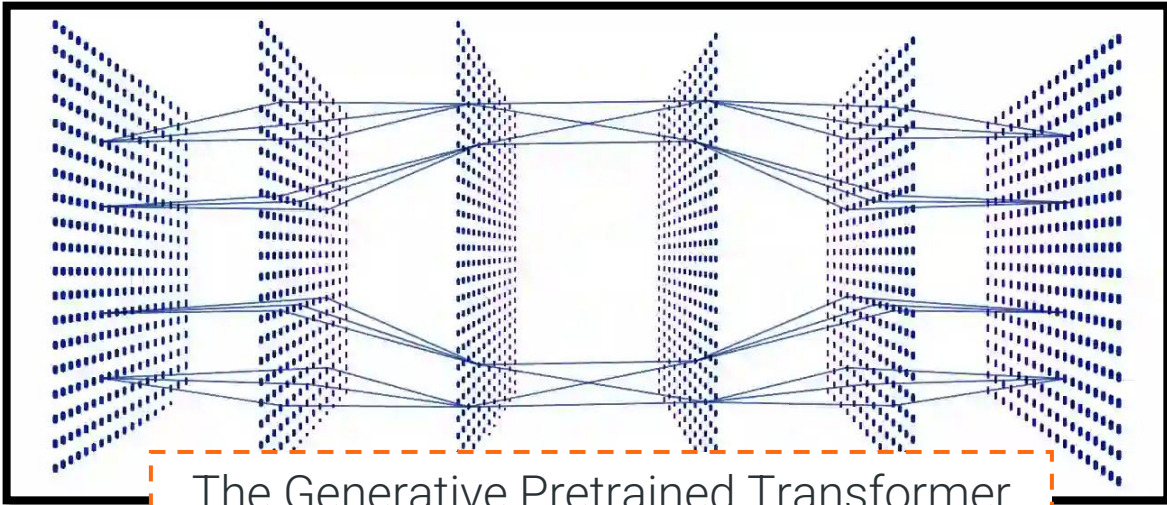
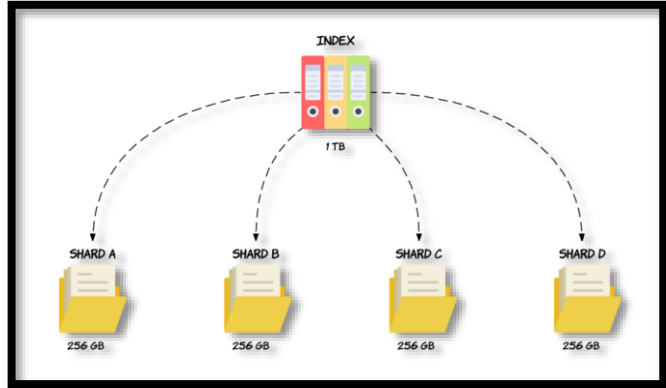
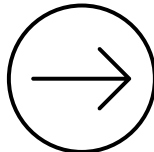
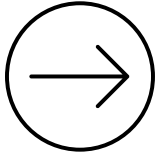
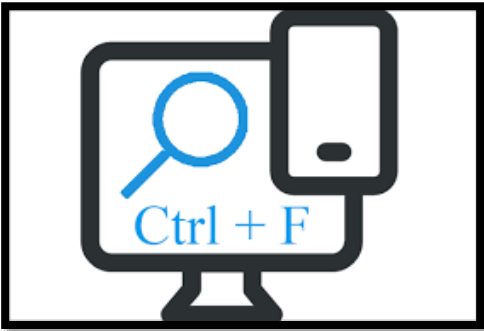
- This is what we could do with unstructured stuff (images, PDF, clickstreams, tweets...) in just a few minutes
- Now imagine what we could do with structured assets, where half the work of structuring and assigning meaning to the data is already done!!
- That's the power of Azure AI Search
 - Connect it to all your data sources – Azure or non-Azure
 - Use its multitude of out-of-box skills – or write your own custom ones
 - Hook it up to hosted Embedding or Inference models from OpenAI or Facebook – or your in-house SLM implementation
- Not on Azure?
 - No worries – Kendra, Bedrock have similar facilities on AWS
 - As does Google Cloud, with Gemini and Vertex AI



Key Learnings from Gen AI Tech Tuesday Webinar

Where data & AI come to **LIFE**

What's new with Enterprise Data?



The Generative Pretrained Transformer

What should you take away from our webinar series?

- The Informatica Data Management Cloud for Generative AI
 - Connectors
 - Azure AI Search, OpenAI
 - Bedrock, Gemini, Vertex AI
 - Pinecone and Ollama
 - Transformations
 - Chunking and Vector Embedding transformations on elastic compute
 - And all the associated Metadata and Data Management Services
 - Data Quality, Governance, Catalog
 - Observability, Privacy, Access Control, Marketplace
 - Data, API and Application Integration – on ephemeral, managed, native compute

What should you take away from our webinar series?

- Generative AI from Informatica Data Management Cloud
 - CLAIRE AI
 - Insights, Inferences, Suggestions
 - CLAIRE GPT
 - Data Discovery
 - Data and Metadata Exploration
 - Data Transformation
 - CLAIRE Copilot
 - Summarization
 - Code Generation
 - ...and more on the way

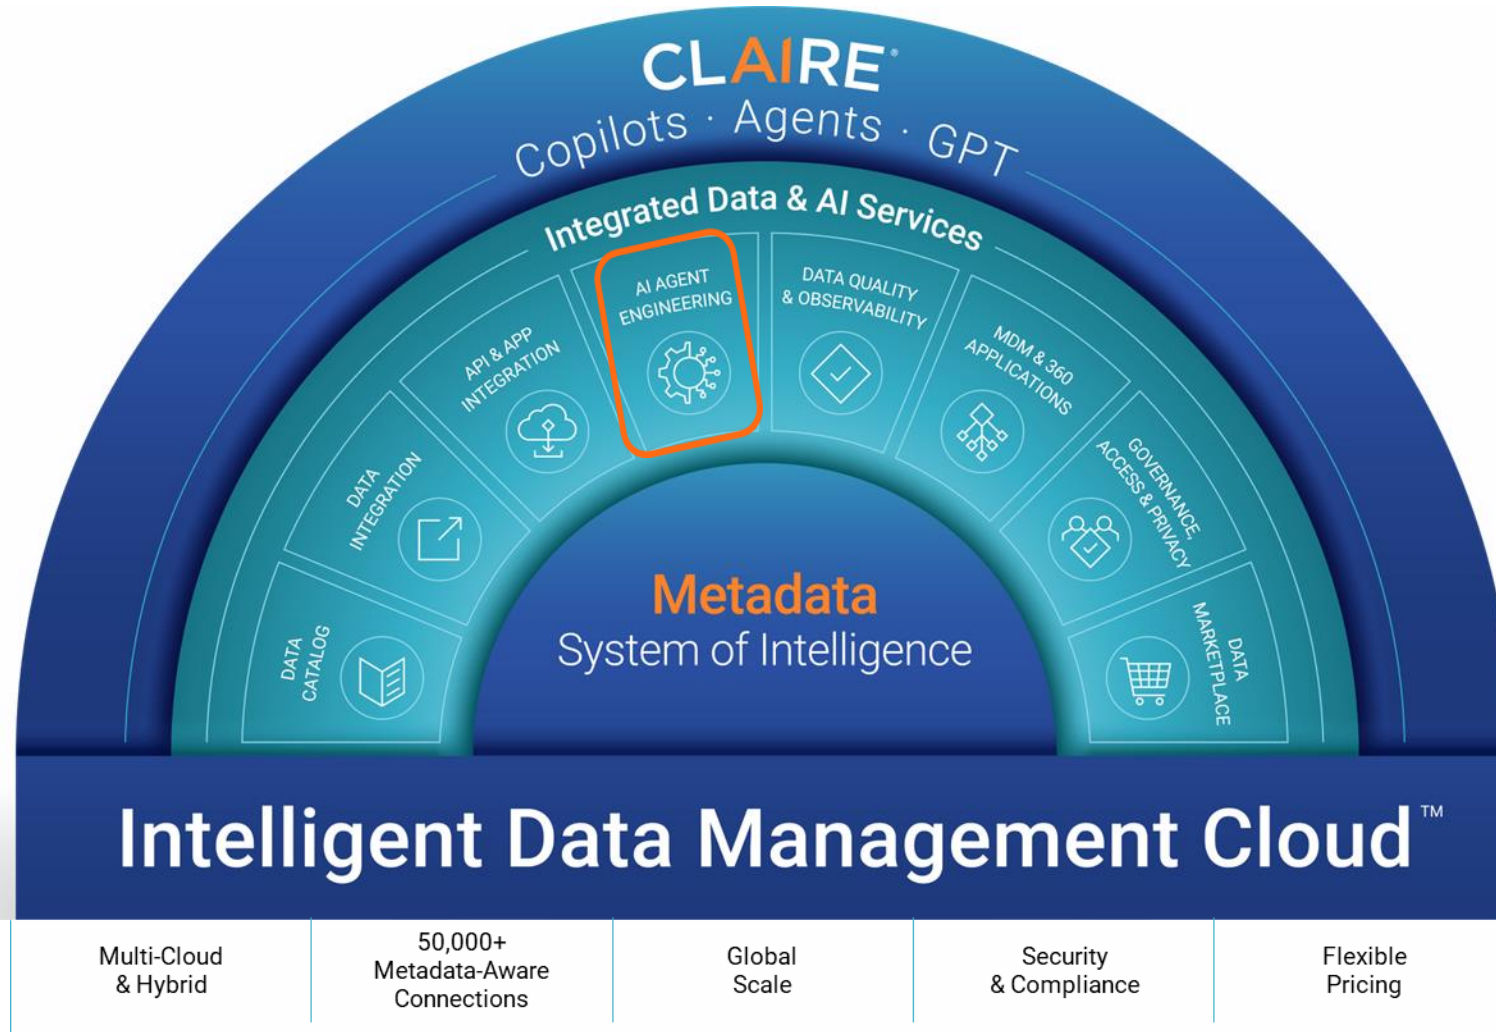
**READY,
SET, AI**



Our Vision for Agentic AI

Where data & AI come to **LIFE**

New AI Engineering Services



AI Agent Engineering

Build accurate agents

- AI-native, no-code Agent Builder & Orchestrator speeds design with drag-and-drop interface, MCP client/server creation, and embeddability into Slack, Teams, and chatbots
- Agent Hub to discover, share, and federate Informatica and 3rd party agents
- Built-in test frameworks that safeguard accuracy, reliability, and compliance before agents reach production

Connect agents across enterprise

- High accuracy through metadata-aware data to contextualize reasoning engine with trusted enterprise data and existing processes
- Built-in integration to popular LLMs, Vector stores, frameworks, and libraries including Langchain, Llamaindex, Python libraries
- Reuse IDMC assets (CAI, CDI, CDQ) turning proven processes into agent skills

Manage agents with confidence

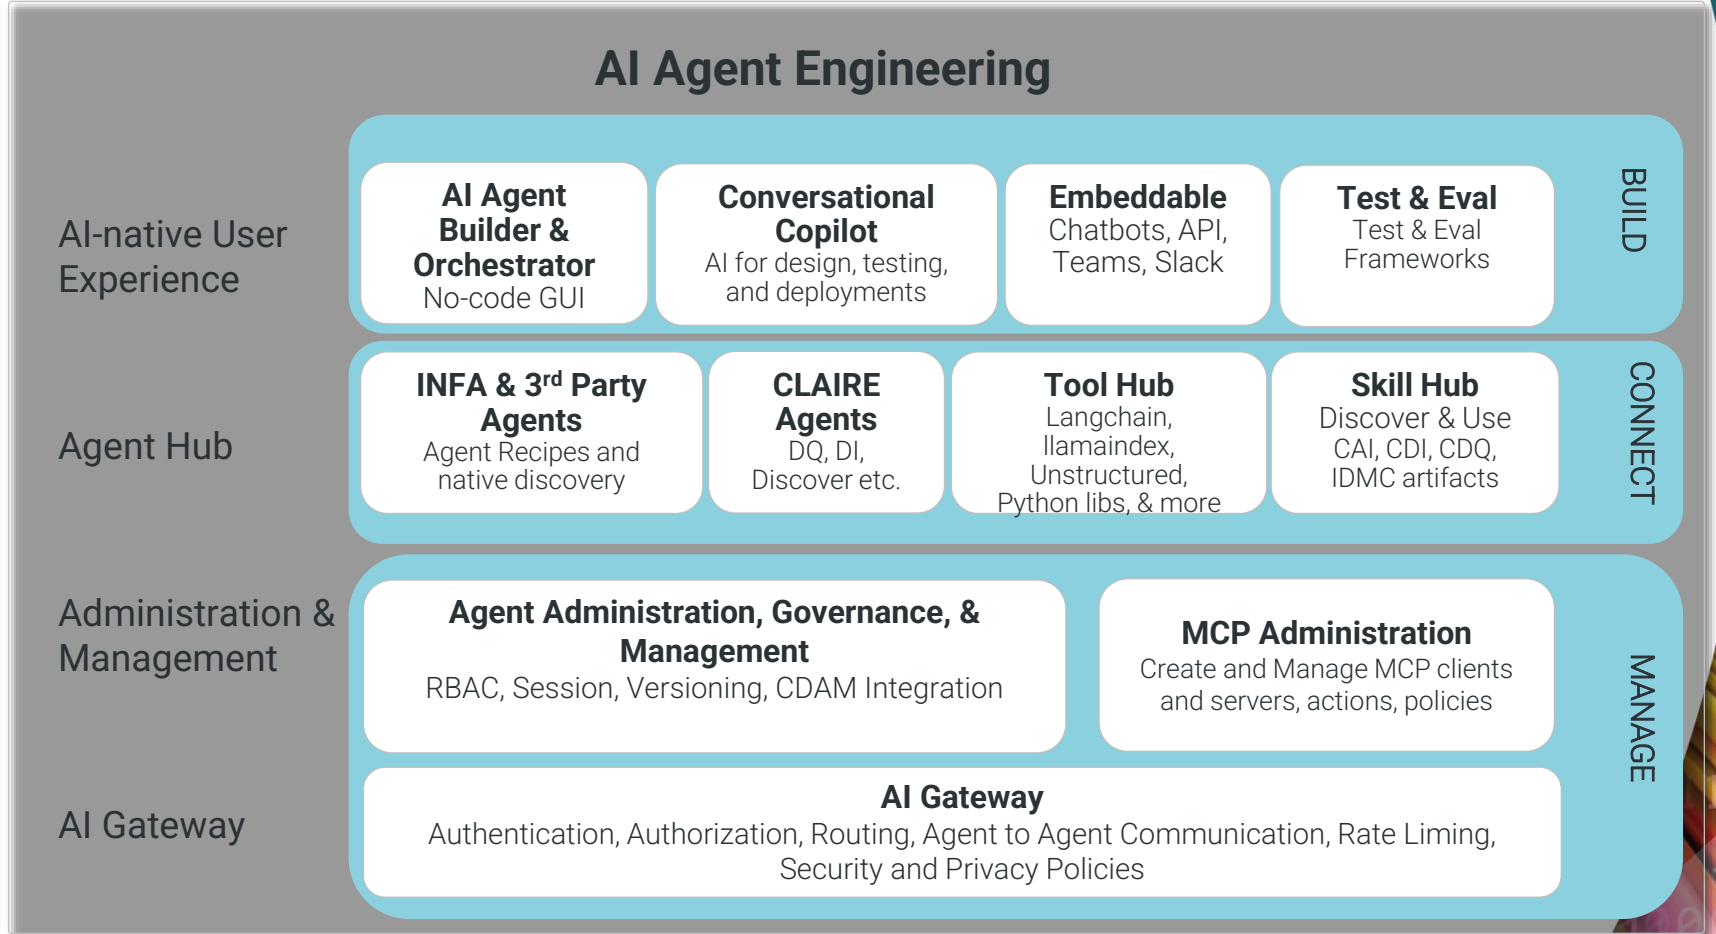
- AI Gateway for centralized enforcement of authentication, routing, and security policies
- Deploy securely to serverless or to your self-hosted
- Enterprise-grade AI governance with RBAC, session management, versioning, rollback, CI/CD

AI Agent Engineering – Product Vision

Build, Connect, and Manage Custom AI Agents

AI Agent Engineering

Build accurate agents, Connect agents across enterprise, and Manage agents with confidence



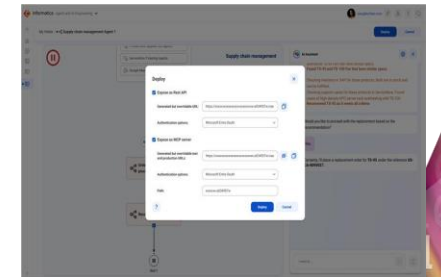
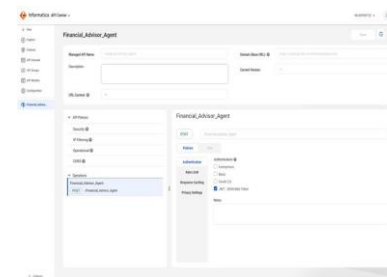
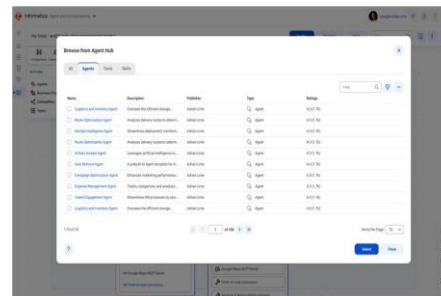
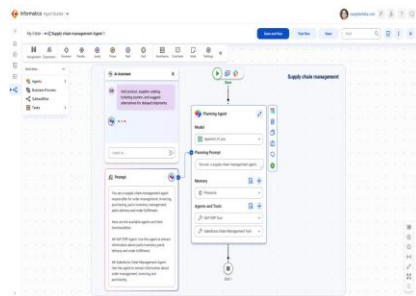
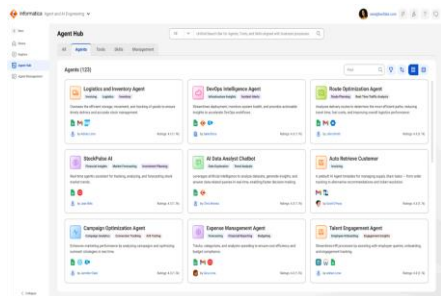
Innovation for AI Agent Engineering

AI agent hub with pre-built recipes and skills

Democratizing AI agents with AI-native no-code experience for building and operationalizing AI agents

Support for third-party agents and agent ecosystems

Support for deployment into cloud ecosystems and as MCP servers



Thank you

Q&A

Reference Links

- Informatica World 2025 Access Sessions and Content: www.informaticaworld.com
- Informatica World 2025 Day 1 General Session:
<https://video.informatica.com/detail/videos/informatica-world/video/6372806636112/informatica-world-2025-day-1-general-session>
- Informatica World 2025 Day 2 General Session:
<https://video.informatica.com/detail/videos/informatica-world/video/6372852816112/informatica-world-2025-day-2-general-session>