

InterSystems IRIS: Vision & Roadmap

Tom Woodfin, Head of **Development**, Data Platforms
DACH Symposium
November 5, 2025



InterSystems IRIS: Vision & Roadmap

Gokhan Uluderya, Head of **Product**, Data Platforms
DACH Symposium
November 5, 2025



InterSystems IRIS: Vision & Roadmap

Gokhan Uluderya, Head of **Product**, Data Platforms
DACH Symposium
November 5, 2025



“Your Data & AI Operating System”



The Operating System Experience: 1985



The Operating System Experience: 1985



```
C:\DOS>chkdsk
Volume Serial Number is 3E76-4B58

2,146,467,840 bytes total disk space
    131,072 bytes in 2 hidden files
    32,768 bytes in 1 directories
    7,405,568 bytes in 124 user files
2,138,898,432 bytes available on disk

    32,768 bytes in each allocation unit
    65,505 total allocation units on disk
    65,274 available allocation units on disk

    655,360 total bytes memory
    602,704 bytes free

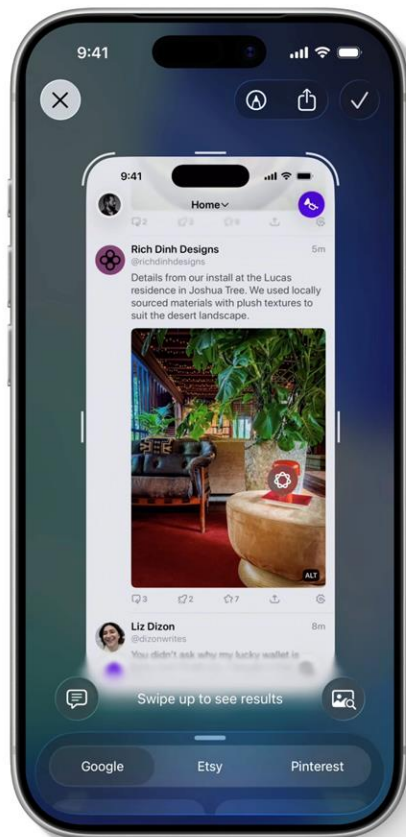
Instead of using CHKDSK, try using SCANDISK.  SCANDISK can reliably detect
and fix a much wider range of disk problems.  For more information,
type HELP SCANDISK from the command prompt.

C:\DOS>_
```

The Operating System Experience: 1995



The Operating System Experience: 2025

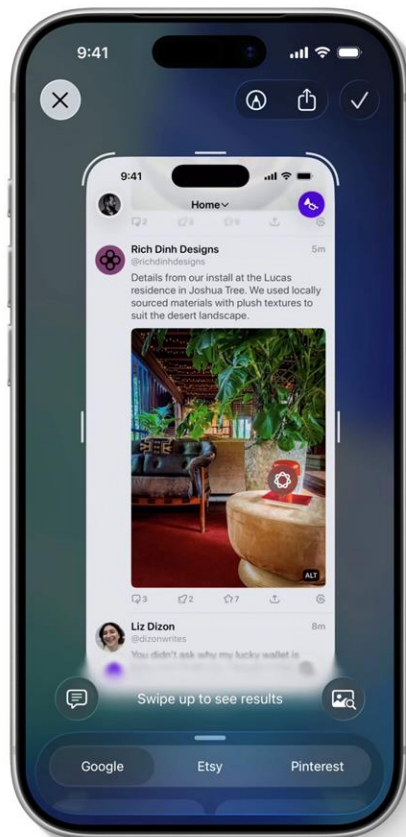


Visual intelligence

Let your iPhone take it all in.

Learn more about what's in front of you, whether you're out and about or using your iPhone. Instantly turn a poster into a Calendar event. Summarize what you're viewing. Ask questions, like recipe ideas for what's in your fridge. And search seamlessly across your favorite apps, like finding where to buy a cool outfit you see online.⁵

The Operating System Experience: 2025



Visual intelligence

Let your iPhone take it all in.

Learn more about what's in front of you, whether you're out and about or using your iPhone. Instantly turn a poster into a Calendar event. Summarize what you're viewing. Ask questions, like recipe ideas for what's in your fridge. And search seamlessly across your favorite apps, like finding where to buy a cool outfit you see online.⁵

The Data Experience: 2025



```
sql> SELECT * FROM information_schema.tables;
```

table_catalog	table_schema	table_name	table_type	owner
accounting_db	public	accounts	BASE TABLE	dbadmin
accounting_db	public	customers	BASE TABLE	dbadmin
accounting_db	public	invoices	BASE TABLE	dbadmin
accounting_db	public	invoice_items	BASE TABLE	dbadmin
accounting_db	public	payments	BASE TABLE	dbadmin
accounting_db	public	transactions	BASE TABLE	dbadmin
accounting_db	public	journal_entries	BASE TABLE	dbadmin
accounting_db	public	tax_rates	BASE TABLE	dbadmin

The Data Experience: 2025



```
sql>
table
-----
acco
acco
acco
acco
acco
acco
acco
acco
```

- > rocinante
 - razorback
 - tycho
 - Tables
 - Volumes
 - Functions
 - Models
 - unity
 - default
 - Tables
 - marksheet
 - marksheet_uniform
 - numbers
 - user_countries
 - Volumes
 - Functions
 - fx lowercase
 - fx sum

Description [✎](#)
Default schema

Tables

Search

Tables Volumes Functions Models

Name	Created At
marksheet	07/17/2024, 18:40:05
marksheet_uniform	07/17/2024, 18:40:05
numbers	07/17/2024, 18:40:05
user_countries	07/17/2024, 18:40:05

Schema details

Created at: 07/17/2024, 18:40:05

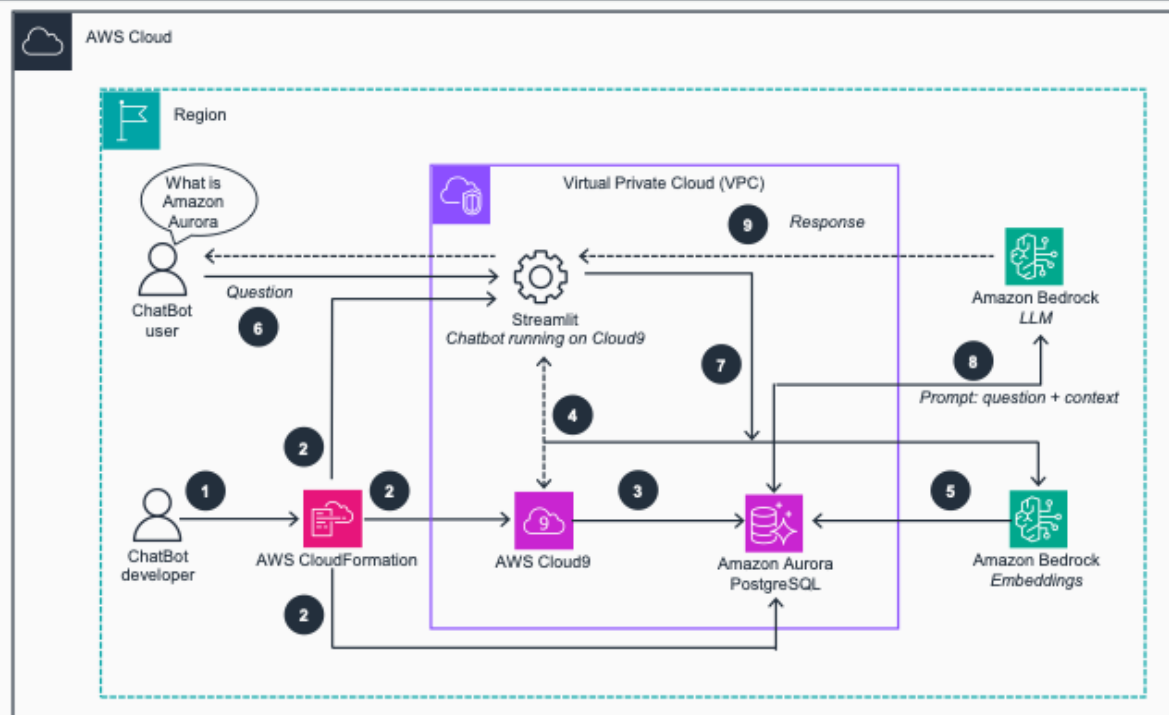
```
r
-----
min
min
min
min
min
min
min
min
```

The Data Experience: 2025



Guidance for High-Speed RAG Chatbots on AWS

This architecture diagram shows how to build an artificial intelligence (AI)-powered chatbot that lets you ask questions based on content in your PDF files in natural language. Once you upload PDF files to the application, type in questions in simple English. The AI-powered application will process the questions and use the Retrieval-Augmented Generation (RAG) technique to generate a response based on the relevant content from the PDFs.



Reviewed for technical accuracy May 30, 2024
© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

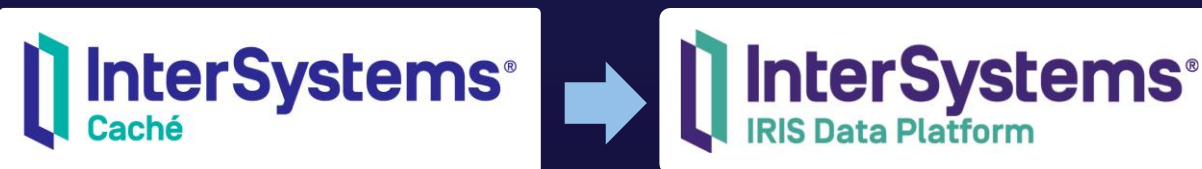
AWS Reference Architecture

The Data Platform Experience: 1997 - 2025



Architecture	Data Center	+ Cloud
Filesystem	Globals	
Virtual Memory	Enterprise Cache Protocol (ECP)	
Data Models	Objects, Tables, Streams	+ Document, Vector
Runtime	ObjectScript, Java, .NET, C++, ...	+ Python
API Model	Client/Server	+ Managed REST
IDE	Caché Studio	VSCode

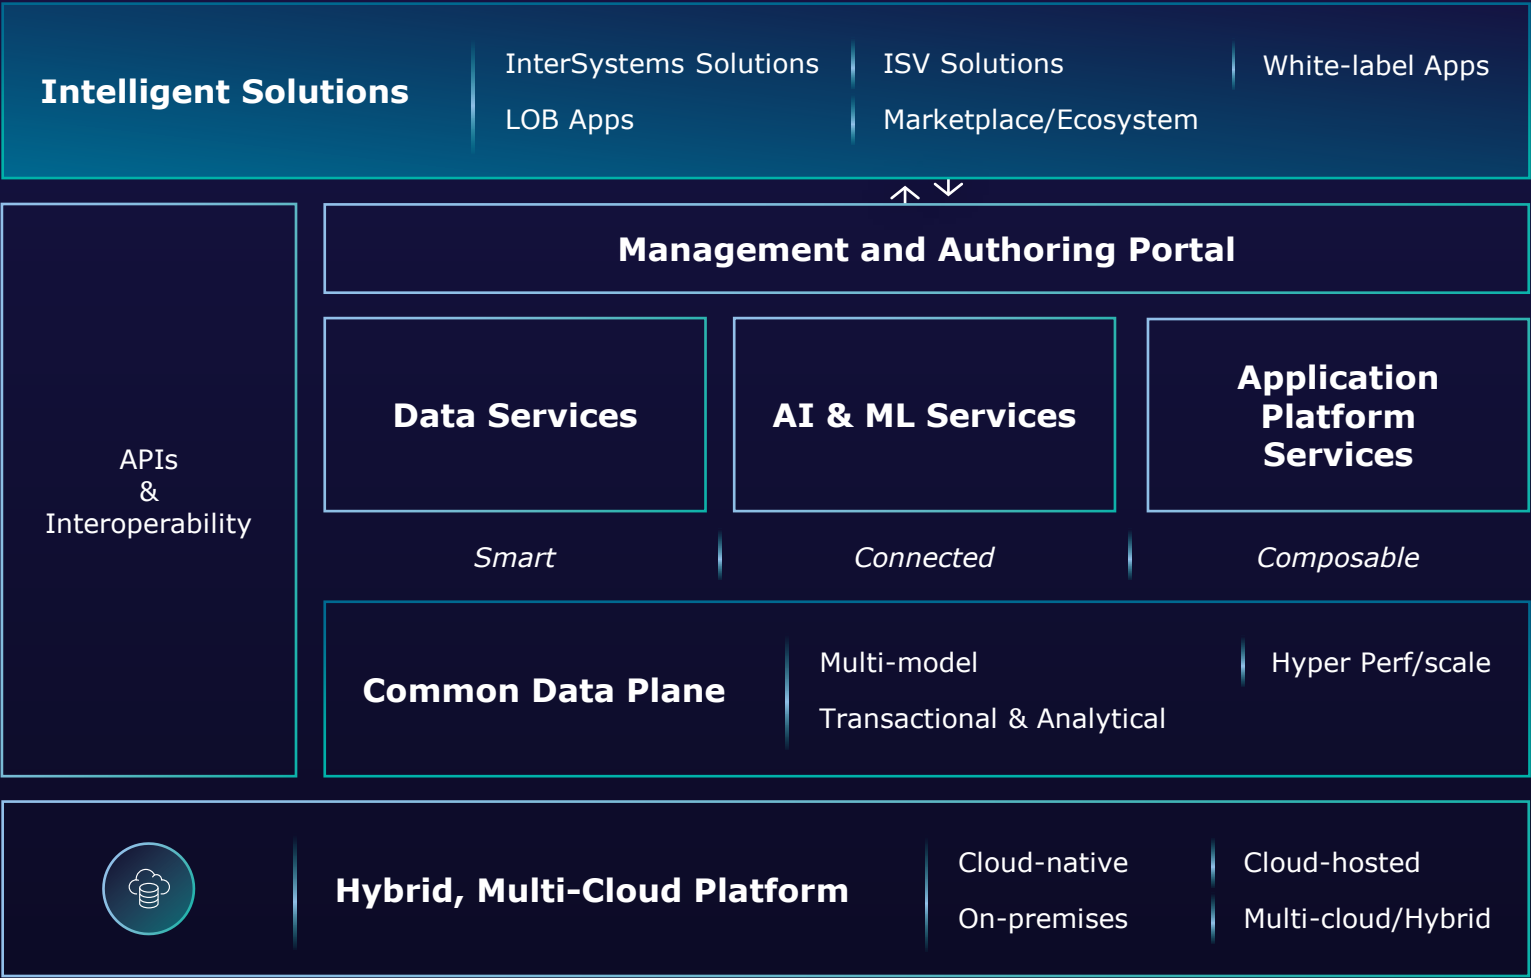
The Data Platform Experience: 1997 - 2025



Architecture	Data Center	+ Cloud
Filesystem	Globals	
Virtual Memory	Enterprise Cache Protocol (ECP)	
Data Models	Objects, Tables, Streams	+ Document, Vector
Runtime	ObjectScript, Java, .NET, C++, ...	+ Python
API Model	Client/Server	+ Managed REST
IDE	Caché Studio	VSCode

Our Vision

IRIS: The Data & AI Operating System



- ✓ **Solutions** - not just building blocks
- ✓ Integrated Cloud Experiences
- ✓ Develop data & AI powered applications
- ✓ Smart Data Fabric in your data ecosystem
- ✓ Data **360-view** natively part of the platform
- ✓ Solve the complexities of cloud infrastructure and operations

Extreme Multi-Model: **Optimized Data Encodings**



List



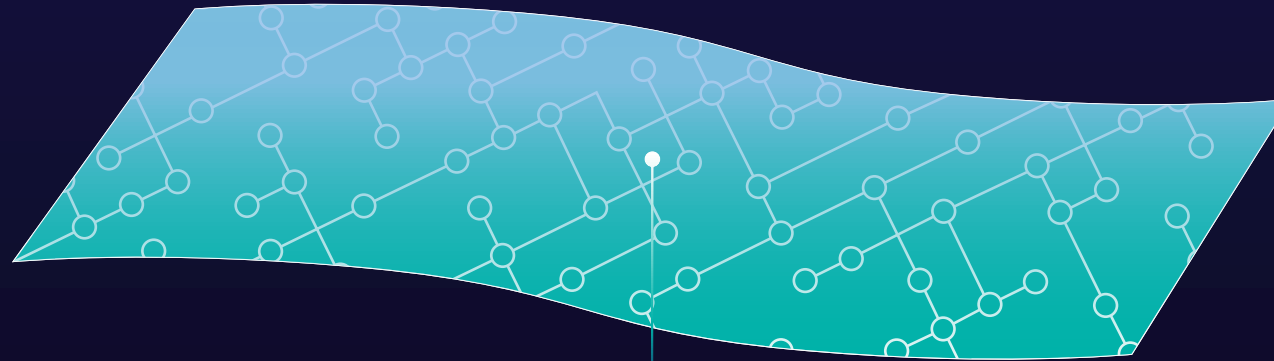
Vector



Bitmap



JSON



\wedge global(<key1>, <key2>, ...) = \$encoding(<val1>, <val2>, ...)

Extreme Multi-Model: Core Structures



**Vector
Embedding**



Relational



Document



Object



Cube



BLOB



List



Vector



Bitmap



JSON

Secret Sauce: Fast, Flexible Data Encodings



Knowledge
Graph



Time Series



Task Queue



Message



Business
Process



Vector
Embedding



Relation



Document



Object



Cube



BLOB



List



Vector/Tensor

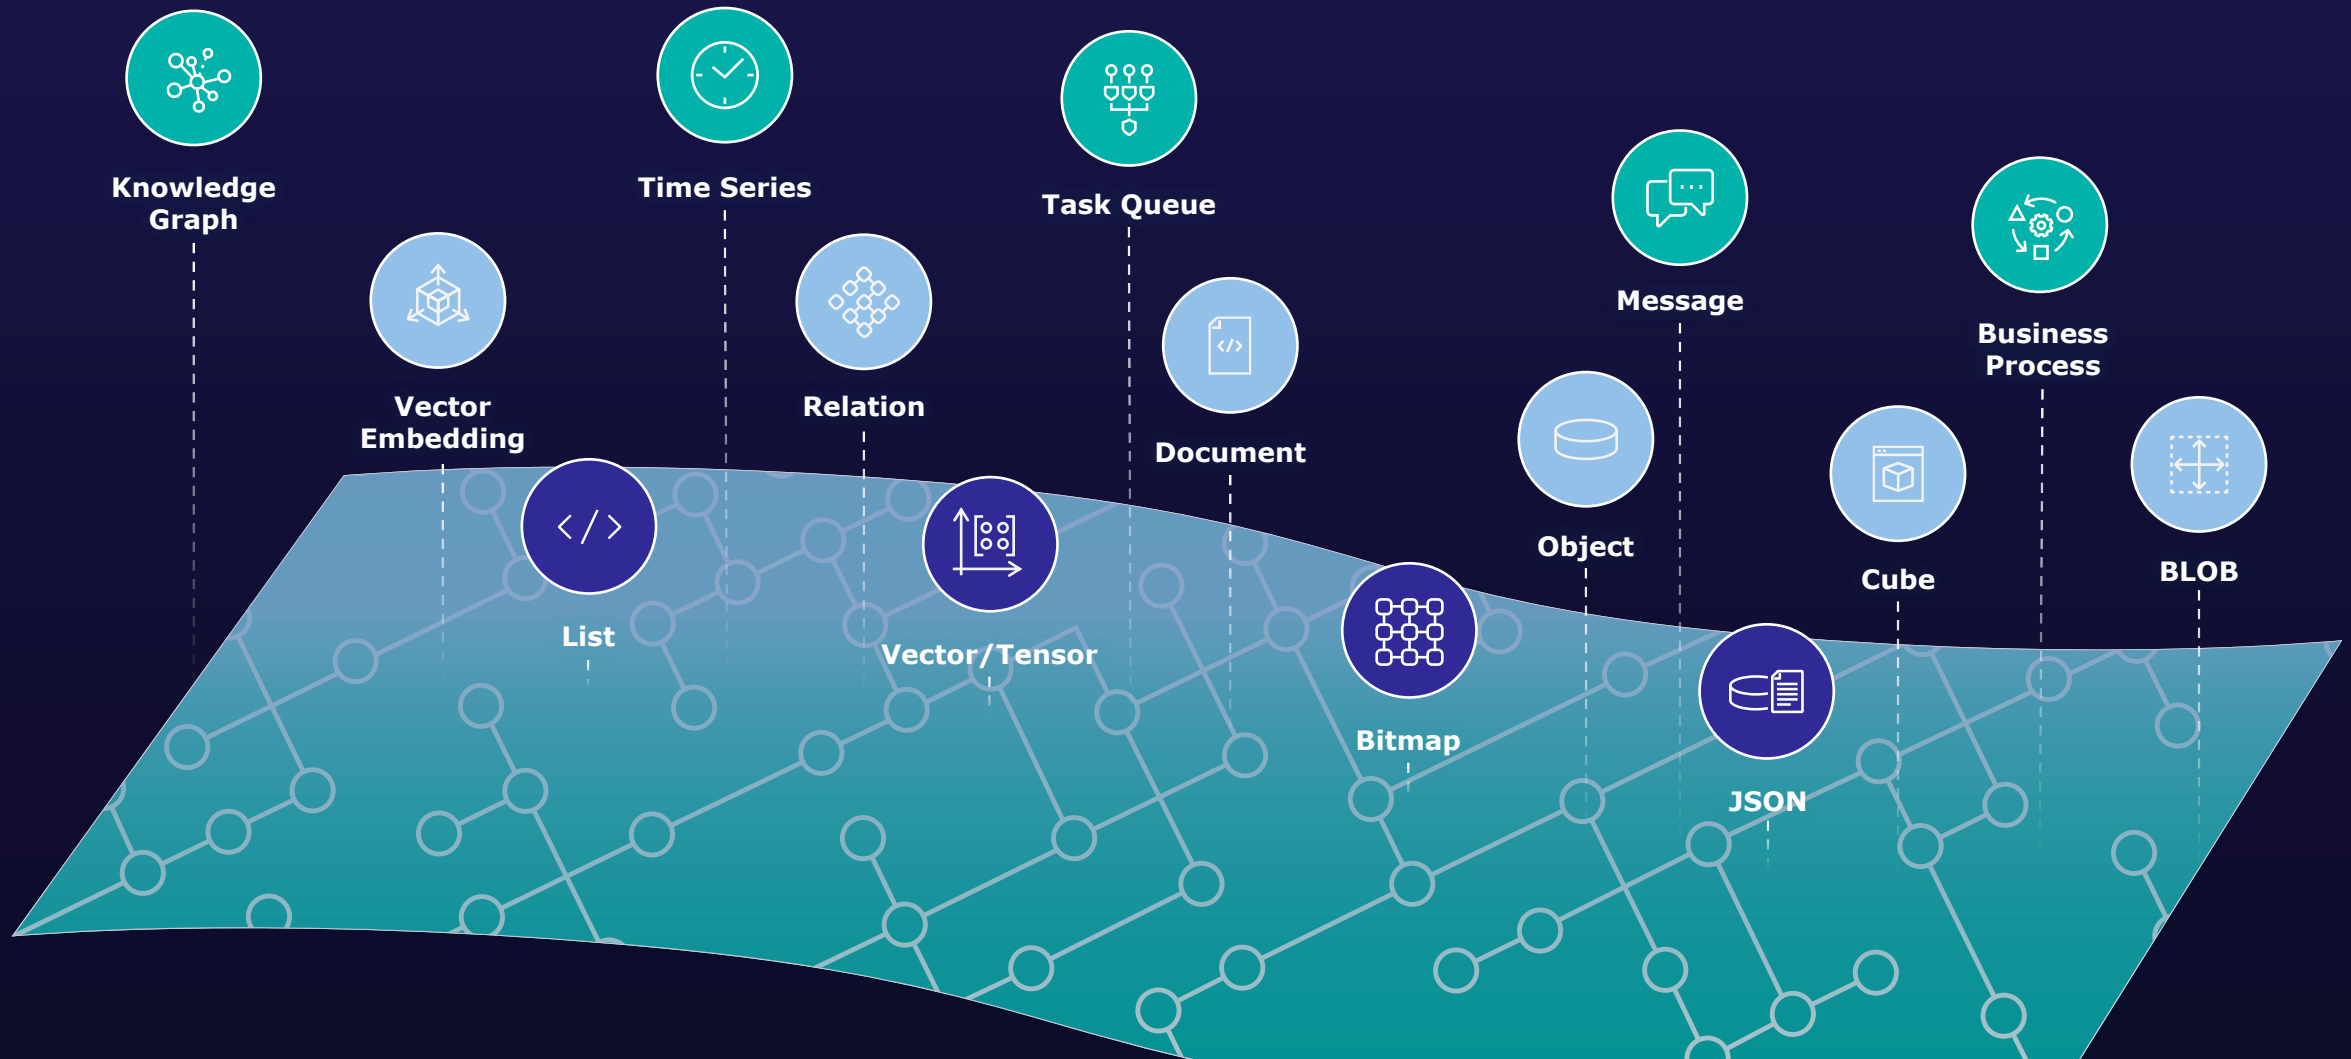


Bitmap



JSON

Secret Sauce: Fast, Flexible Data Encodings





Your Data & AI Operating System: User Experiences



Your Data & AI Operating System: User Experiences



For the **Business User**:
Build in the **solution space**,
with InterSystems **Data Studio**

Your Data & AI Operating System: User Experiences



For the **Business User**:
Build in the **solution space**,
with InterSystems **Data Studio**



For the **Developer**:
Embrace the **developer ecosystem**,
with tools developers already love

Your Data & AI Operating System: User Experiences



For the **Business User**:
Build in the **solution space**,
with InterSystems **Data Studio**



For the **Developer**:
Embrace the **developer ecosystem**,
with tools developers already love



For the **AI Agent**:
Trusted data & **governed** action,
with InterSystems **AI Hub**

InterSystems Data Studio

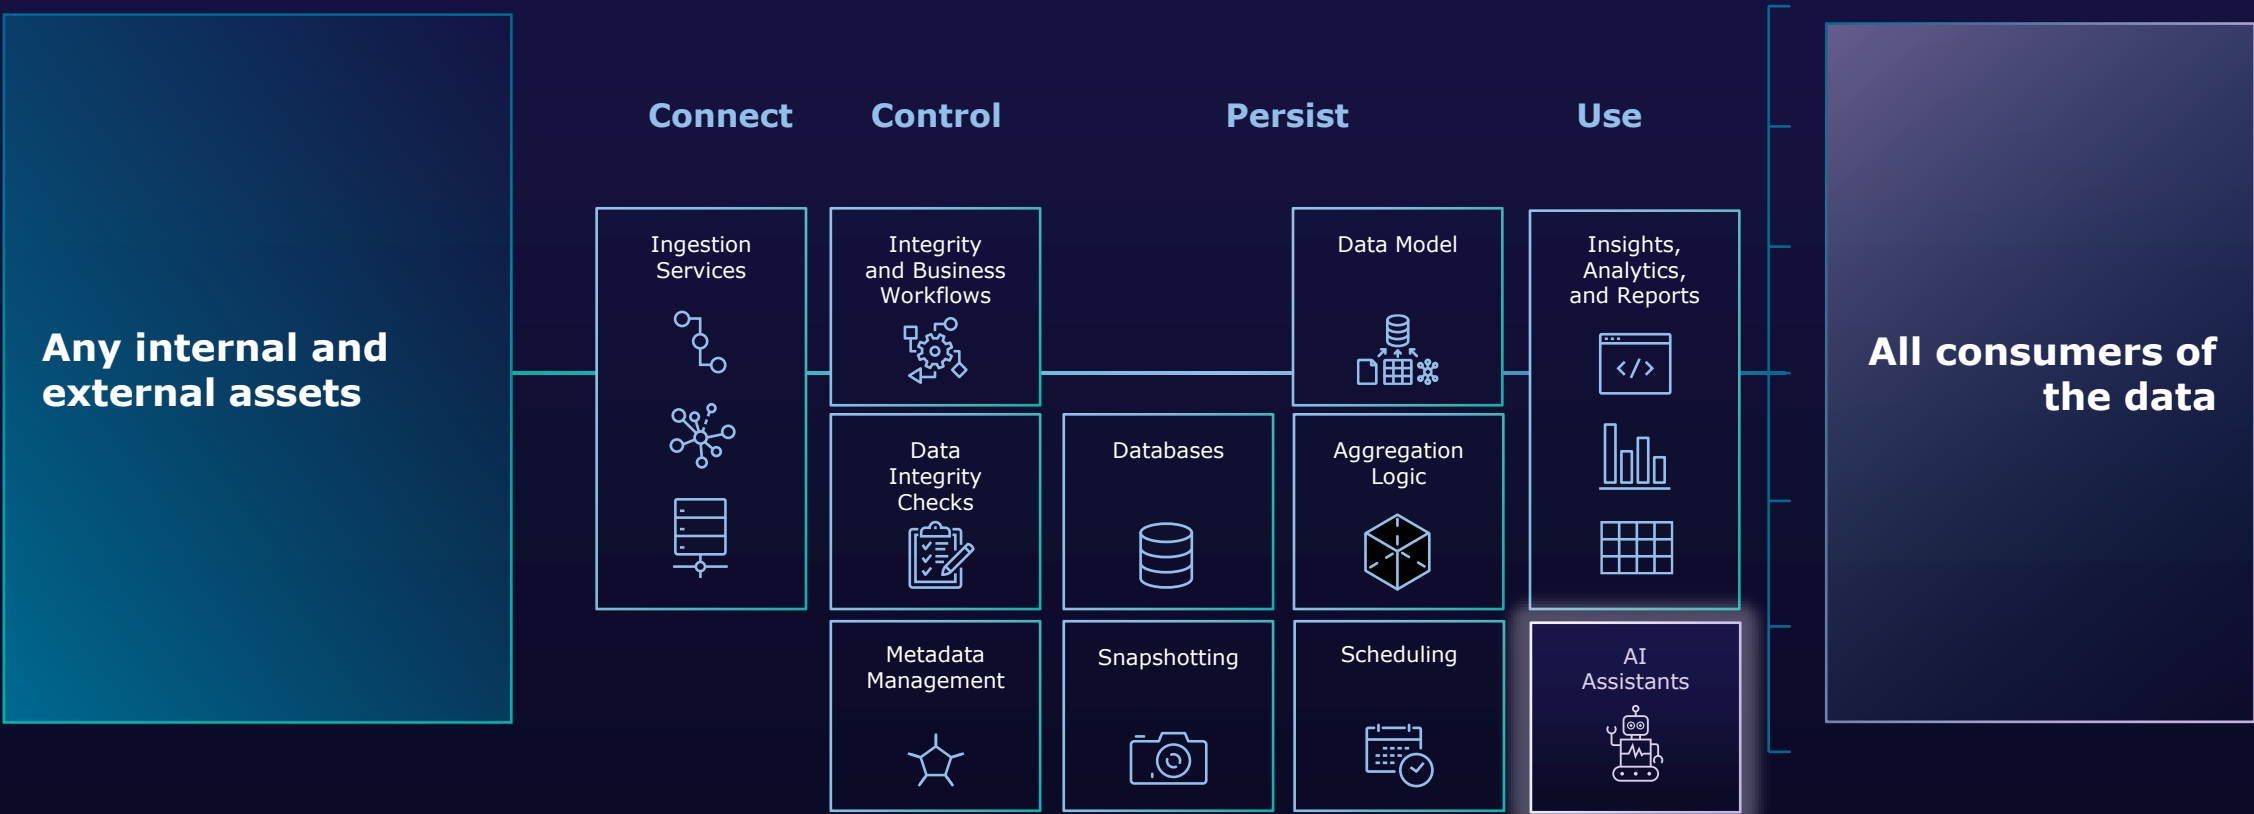


Delivering **Self-Service** Data Management to all Data Consumers

Data Sources

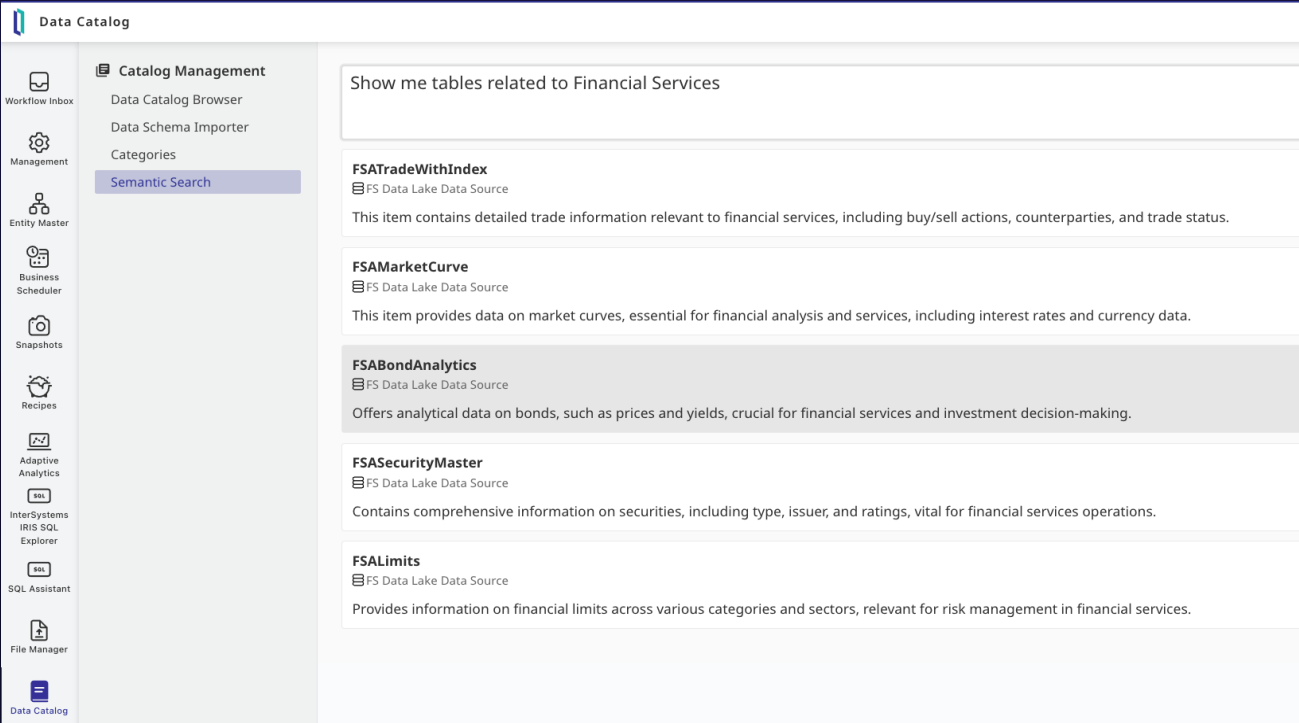
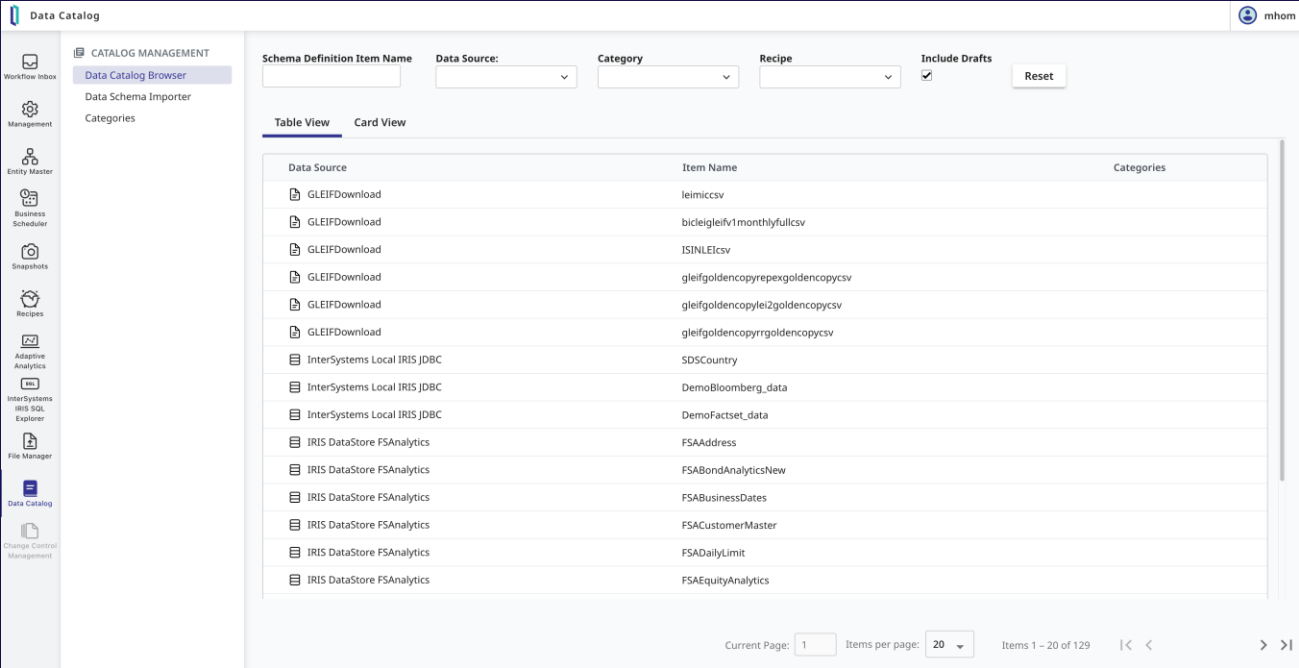
InterSystems Data Studio

Data Users



InterSystems Data Studio

Data Catalog Powered By Semantic Search



InterSystems Data Studio

Research Assistant Generating Visualizations with AI

Visualize

Research Assist

SQL Assist

Catalog

Schedule

Create

Snapshot

Stream

Remediate

Manage - System

Manage - Calendars

Manage - Files

Manage - Models

Train

Research Assistant

mhom

Graph the results

Chef 8:45 AM, 18 Aug 2025

The graph visualizing the revenue results of Microsoft, Apple, Nvidia, and Tesla over the last five years has been created. Here are the key observations based on the data:

- Microsoft:**
 - Revenue increased steadily from approximately \$143 billion in 2020 to \$245.1 billion in 2024.
- Apple:**
 - Revenue showed a consistent upward trend, starting at \$260.2 billion in 2019 and reaching \$383.3 billion in 2023.
- Nvidia:**
 - Experienced significant growth, with revenue rising from \$10.92 billion in 2020 to \$60.92 billion in 2024.
- Tesla:**
 - Revenue grew from \$24.58 billion in 2019 to \$96.77 billion in 2023, showing a strong upward trajectory.

Overall, Apple maintained the highest revenue among the four companies, while Nvidia showed the most rapid growth rate. Tesla also demonstrated substantial growth, although it started from a lower base compared to the others.

If you need further details or specific visualizations, please let me know!

Chef 8:45 AM, 18 Aug 2025

Revenue Results of Microsoft, Apple, Nvidia, and Tesla (Last 5 Years)

Year	Microsoft	Apple	Nvidia	Tesla
2019	143.0	260.2	10.92	24.58
2020	145.0	275.0	12.0	30.0
2021	168.0	366.0	18.0	54.0
2022	198.0	394.0	28.0	83.0
2023	212.0	383.3	28.0	96.77
2024	245.1	383.3	60.92	96.77

Type a message...

Business User Experience: Build in the **Solution Space**



Workflow Inbox

Management

Entity Master

Business Scheduler

Snapshots

Recipes

Adaptive Analytics

InterSystems IRIS SQL Explorer

File Manager

Data Catalog

Change Control Management

Entity Details

Entity List

Add Entity

Search entities

> EU - London

> US

US - NYSE

US - SIFMA

Entity Details

Add Child Entity

☒ Enable Entity

Entity Description*

US - SIFMA

☒ Use Parent Time Zone?

☒ Use Parent Fiscal Calendar?

☒ Use Parent Business Week?

☐ Use Parent Holiday Calendar?

Holiday Calendar

Year

Search holidays

Toggle	Holiday ↓↑	Date ↓↑
<input checked="" type="checkbox"/>	Martin Luther King Jr. Day	2022-01-17
<input checked="" type="checkbox"/>	Washington's Birthday	2022-02-21
<input checked="" type="checkbox"/>	Good Friday	2022-04-15
<input checked="" type="checkbox"/>	Memorial Day	2022-05-30
<input checked="" type="checkbox"/>	Juneteenth	2022-06-20
<input checked="" type="checkbox"/>	Independence Day	2022-07-04
<input checked="" type="checkbox"/>	Labor Day	2022-09-05
<input checked="" type="checkbox"/>	Thanksgiving Day	2022-11-24
<input checked="" type="checkbox"/>	Christmas Day	2022-12-26
<input checked="" type="checkbox"/>	New Years Day	2023-01-02
<input checked="" type="checkbox"/>	Martin Luther King, Jr. Day	2023-01-16
<input checked="" type="checkbox"/>	Washington's Birthday	2023-02-20
<input checked="" type="checkbox"/>	Good Friday	2023-04-07
<input checked="" type="checkbox"/>	Memorial Day	2023-05-29
<input checked="" type="checkbox"/>	Juneteenth	2023-06-19
<input checked="" type="checkbox"/>	Independence Day	2023-07-04
<input checked="" type="checkbox"/>	Labor Day	2023-09-04
<input checked="" type="checkbox"/>	Thanksgiving Day	2023-11-23
<input checked="" type="checkbox"/>	Christmas Day	2023-12-25
<input checked="" type="checkbox"/>	New Years Day	2024-01-01

Developer Experience: Build with the **Tools You Know** & ❤️



The screenshot shows a VS Code workspace named 'cinema2025 (Workspace)'. The Explorer sidebar on the left shows the project structure:

- EXPLORER
 - OPEN EDITORS (1 unsaved)
 - Film.cls cls/Cinema
 - cinema2025-em... 1
 - CINEMA2025 (WORKSPACE) (1K+)
 - cinema2025
 - .vscode
 - bin
 - cls/Cinema
 - Film.cls
 - HugeFilm.cls
 - Review.cls
 - Util.cls
 - data
 - include
 - lib
 - screenshots
 - sql
 - cinema.sql
 - cinema2025.sql M
 - www

The main editor shows the file 'cinema2025-embeddings-compute-max.py' with the following code:

```
32
33 # Updating our Cinema embeddings at max velocity
34 import iris
35 import torch
36 from sentence_transformers import SentenceTransformer
37 import ray
38
39 @ray.remote # Engage max distributed power with Ray!
40 def update_embeddings(object_ids):
41     # Push Apple Silicon to the max with Metal GPU support!
42     device = torch.device('mps')
43     # Load the sentence transformer model
44     model = SentenceTransformer('all-MiniLM-L6-v2', device=device)
45     update_count = 0
46     for id in object_ids:
47         # Update the embeddings for a single IRIS Film object
48         film = iris.cls('Cinema.Film')._OpenId(id)
49         embedding = model.encode(film.overview)
50         film.load_embedding(embedding)
51         sc = film._Save()
52         if sc == 1:
53             update_count = update_count + 1
```

Developer Experience: Build with the **Tools You Know** & ❤️



iris-mcp main



SQLTool.cls

Class User.XData

ExampleService.cls

cls\AI\MCP\Service\SQLTool.cls

```
1  /// This service provides access to SQL tools within an InterSystems IRIS instance.
2  /// It is meant to be used by developers to manage and query data, not for use within
3  /// an application.
4  Class %AI.MCP.Service.SQLTool Extends %AI.MCP.Service
5  {
6  /// Service name for MCP protocol
7  Parameter SERVICENAME = "sql-tools";
8
9  /// Service version
10 Parameter VERSION = "1.0.0";
11
12 /// Instructions for this service
13 XData INSTRUCTIONS [ MimeType = "text/markdown" ]
14 {
15 # %AI.MCP.Service.SQLTool
16
17 An MCP service exposing SQL tooling over InterSystems IRIS.
```

Agent Experience: InterSystems AI Hub



AI Development on IRIS



Accelerate **native Agent development** through seamless, governed low-code access to LLMs, MCP tools, RAG, tool-calling templates & more

AI Enablement of IRIS



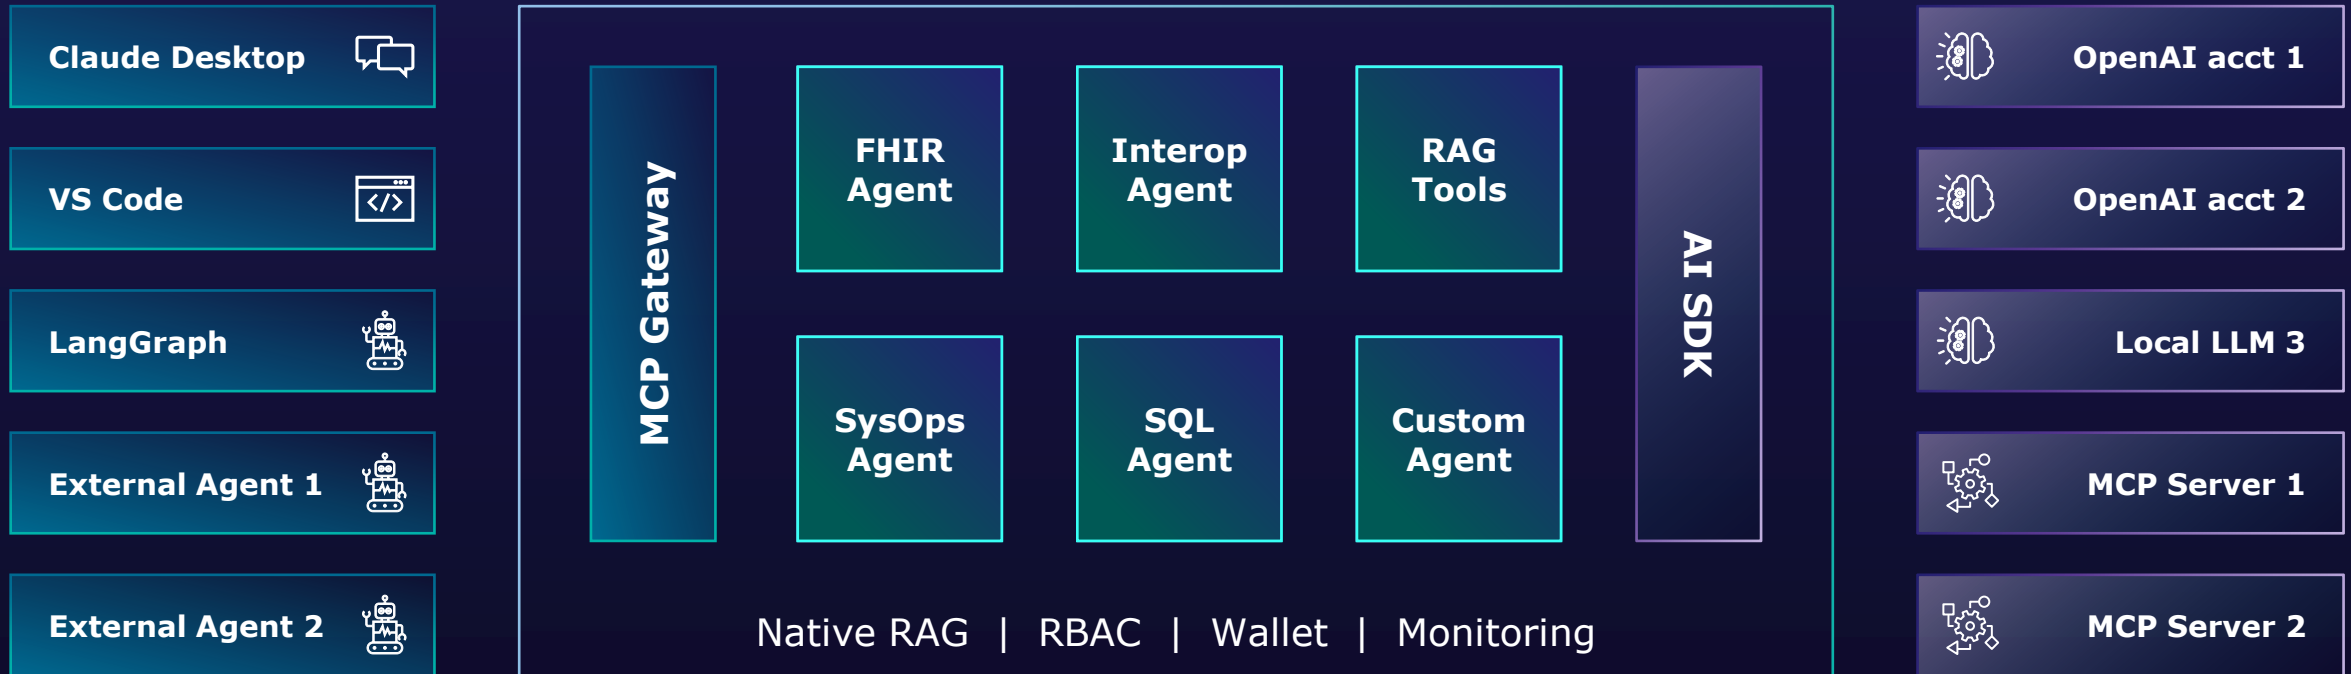
Accelerate **external Agent development** through robust, secure no-code projection of IRIS-based capabilities to MCP

The InterSystems AI Hub **accelerates time-to-value** by letting you focus on business functions rather than AI infrastructure, ensuring **coherent security, monitoring & governance** for your AI apps & agents

Agent Experience: Trusted **Data**, Governed **Action**



InterSystems IRIS



Data & AI Platform **Innovation**



Data Platform for AI-Ready Data

- PB-Size Transactional/Analytical Database
- Security Database
- Security Wallet
- Foreign Tables / Data Virtualization Enhancements
- BI/Analytics Experiences
- Monitoring and Guardrails for Data Management
- Native Open Table Format Support

Differentiated AI Capabilities

- Agentic Framework with Multi-Path Reasoning
- Customizable Agents w/ Knowledge/Memory
- AI Assistants for Complex Transformations
- MCP Gateway
- AI Gateway
- RAG Recipes

Experiences for Ease-of-Use

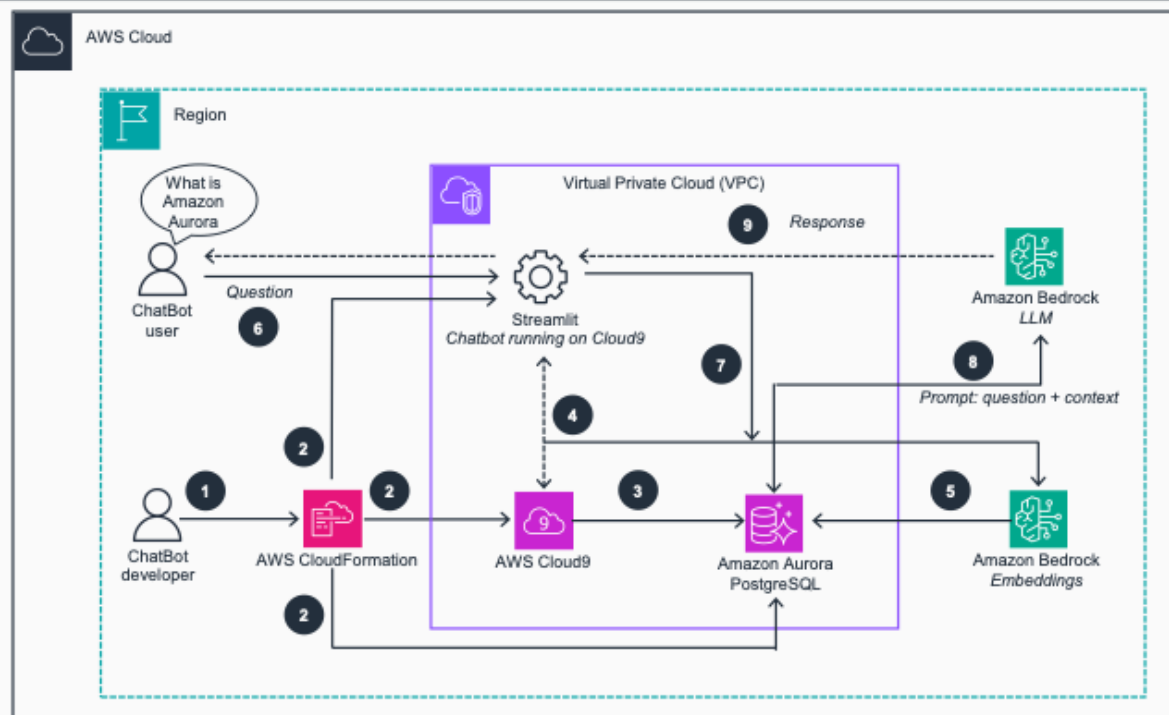
- Data Studio Health Module
- Advanced Interoperability Experiences
- Enhanced Python Experiences
- Data Governance, Compute Management Experiences in Data Studio
- New API/UI experiences for System Administration

The Data Experience: 2025



Guidance for High-Speed RAG Chatbots on AWS

This architecture diagram shows how to build an artificial intelligence (AI)-powered chatbot that lets you ask questions based on content in your PDF files in natural language. Once you upload PDF files to the application, type in questions in simple English. The AI-powered application will process the questions and use the Retrieval-Augmented Generation (RAG) technique to generate a response based on the relevant content from the PDFs.

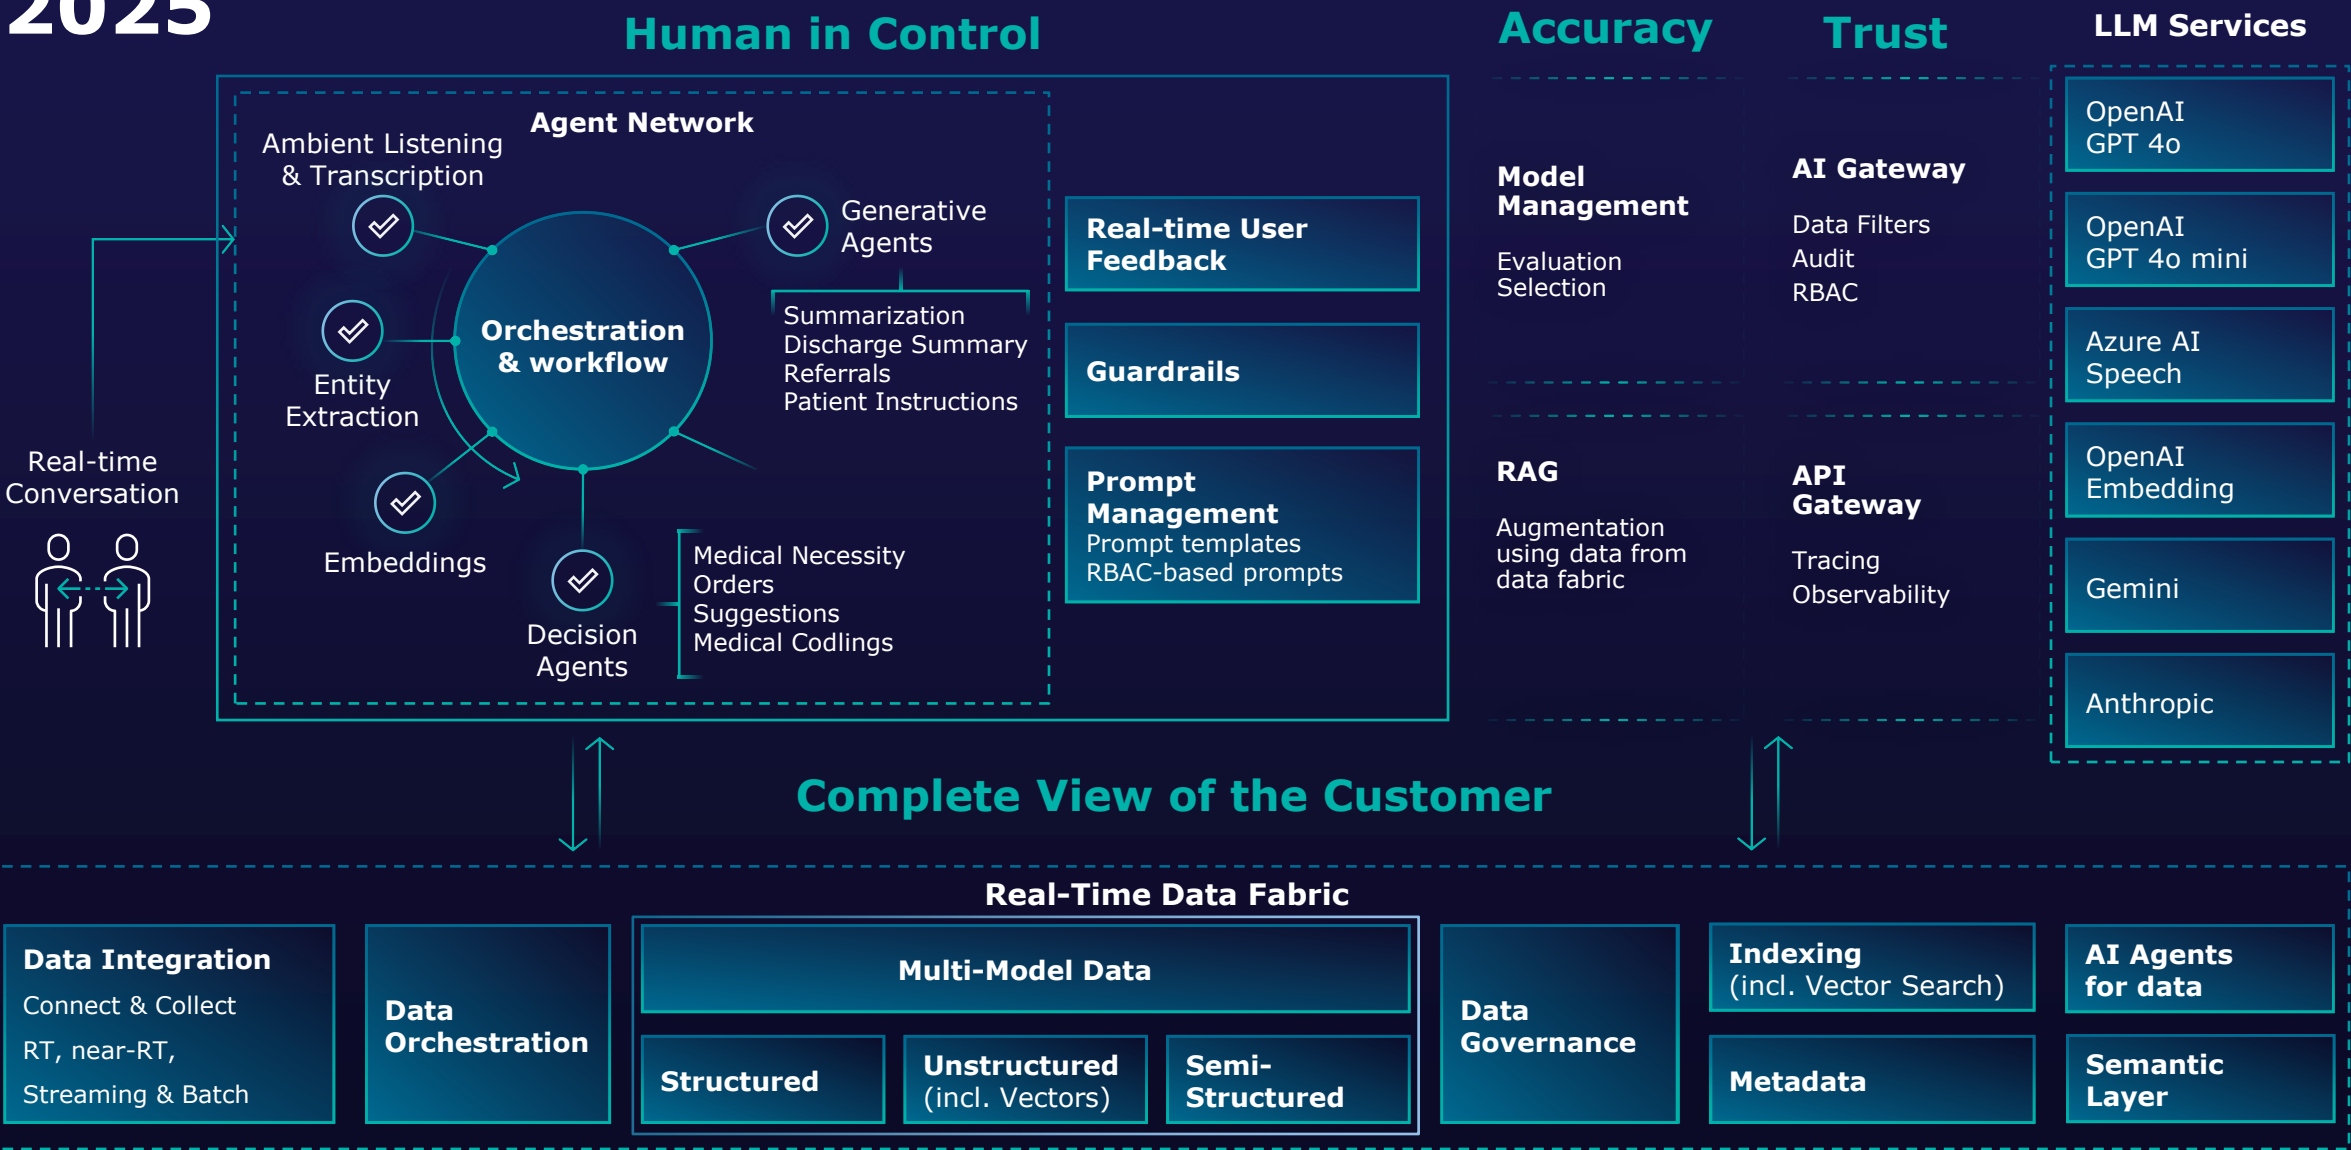


Reviewed for technical accuracy May 30, 2024
© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

AWS Reference Architecture

- 1 Download the **AWS CloudFormation** template from the GitHub repository and deploy the **CloudFormation** stack.
- 2 The **CloudFormation** stack deploys an **AWS Cloud9** instance, an **Amazon Aurora PostgreSQL** cluster, a **Streamlit** custom chatbot application, and other pre-requisites required for this Guidance.
- 3 Set up the environment variables to connect to the **Aurora PostgreSQL** instance, create the **pgvector** extension, and start the **Streamlit** application.
- 4 Once the **Streamlit** application starts, upload the PDF document for processing. This will segment the document into chunks and convert them into vectors using an **Amazon Titan** model from **Amazon Bedrock**. **Amazon Bedrock** is a fully managed service that offers a choice of high-performing foundation models (FMs).
- 5 Load the vector embeddings into an **Aurora PostgreSQL** cluster.
- 6 The user asks a question in natural language in the chatbot application.
- 7 The question from the **Streamlit** application is converted into embeddings using the **Amazon Titan** model. The vectors are then compared with the **Aurora PostgreSQL** vector store to identify the most semantically similar vectors.
- 8 Pass the user question and the context from the vector database to the large language model (LLM). In this example, the **Claude 3** model from **Anthropic** on **Amazon Bedrock** is used.
- 9 The LLM generates a response based on the relevant content and displays the response in the chatbot application.

The IRIS Data Platform Experience: 2025





Thank You