



Enterprise-grade Agentic RAG: From Data Unification to Trustworthy Insights

A blueprint for scalable, auditable AI knowledge systems

Viktor Botev
CTO & Co-Founder

IRIS.AI

World
Summit 
08-09 October 2025 | Amsterdam

The Pilot Trap

95%
AI Pilots Fail



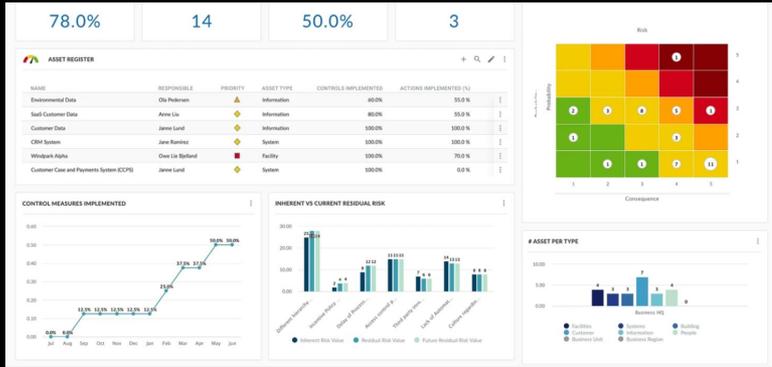
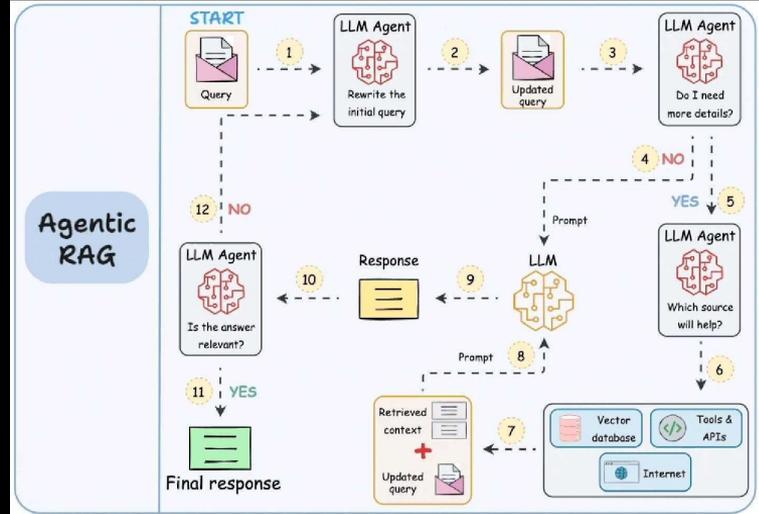
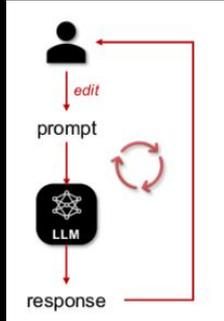
The Real Enterprise Challenge

- Disconnected silos across departments.
- No unified context for reasoning.
- High cost of hallucinations & non-compliance.



From Retrieval to Agentic RAG

SEARCH



Enterprise Agentic RAG Architecture

Unified Data

Fabric

Connect

to any structured or unstructured data source and delivers agent-ready data without rebuilding pipelines, even PDF!
(Heavy Indexing)

Agentic

RAG

Develop

domain-specific LLM agents to identify the best data sources align to every questions & answer with the right reasoning

LLM

Evaluation

Optimize

the quality of entire agentic workflow and ability to evaluate LLMs continuously.

Built-in

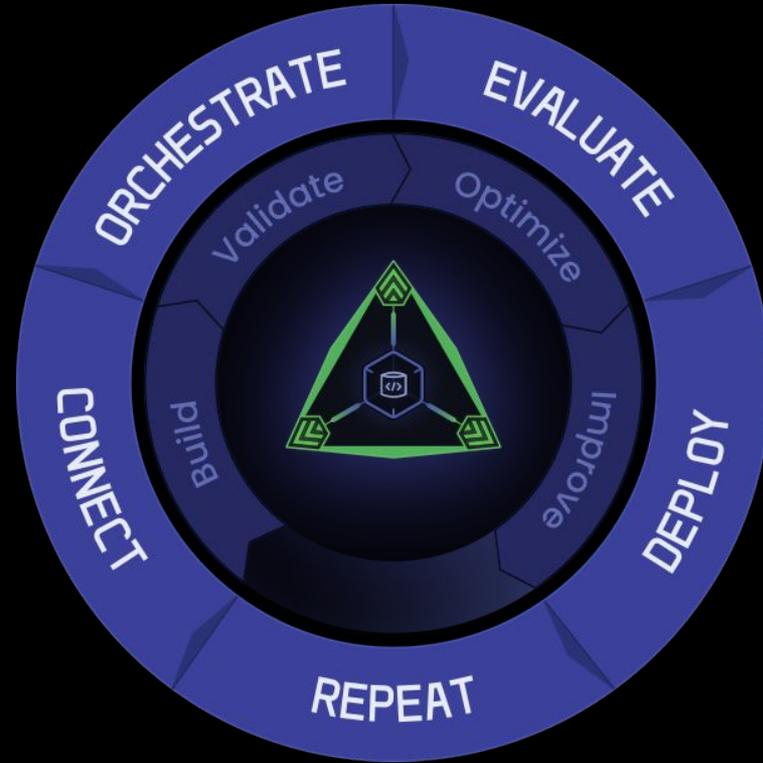
Governance

Monitor

with fine-grained access policies, real-time quality & efficiency metrics tracking for continuous improvement

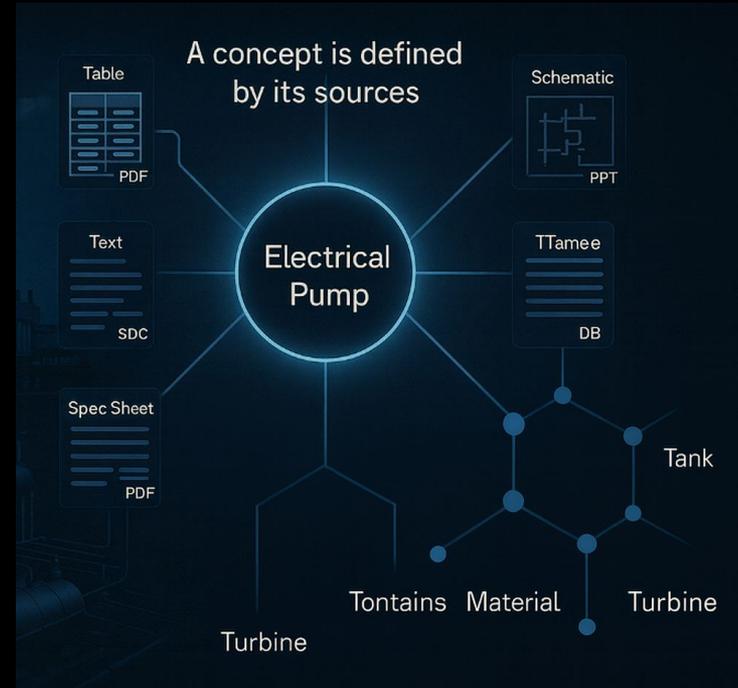
Enterprise Agentic RAG Architecture

- **CONNECT** – Your Data
(Extract, OCR, Index, Ontology)
- **BUILD** – RAG Core
(Retriever, Assembler, Generator)
- **ORCHESTRATE** – Specialized Agents
- **EVALUATE** (ConSens, WISDM)
- **DEPLOY** (Monitoring, Cost Control)

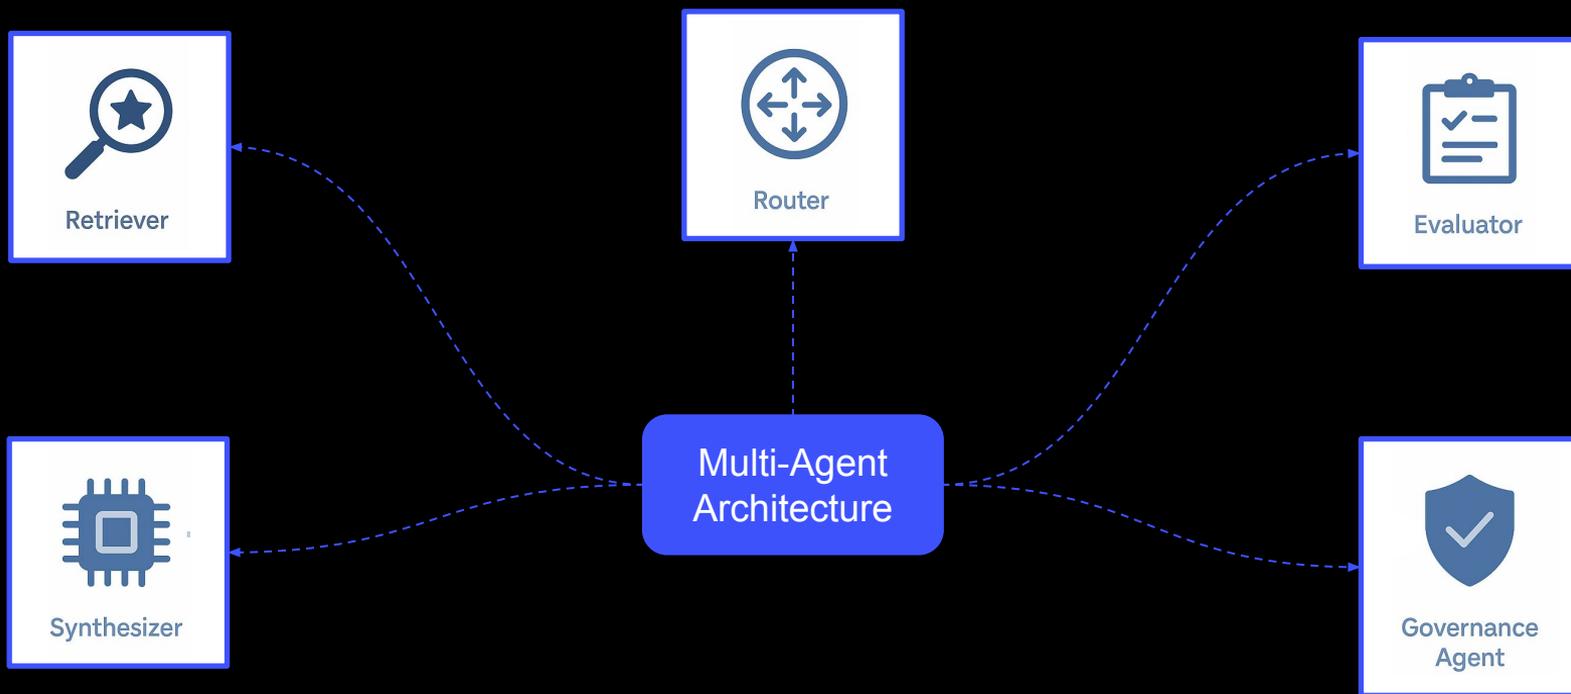


Data Unification

- Multimodal **extraction** (text, tables, images, schematics).
- Linking entities and context into reusable **knowledge assets**.
- **Asset** – an Entity connected to specs, drawings, and databases.

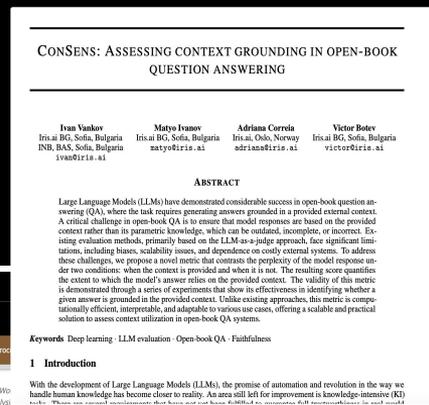


Orchestrating Specialized Agents



Evaluation & Governance – The Trust Layer

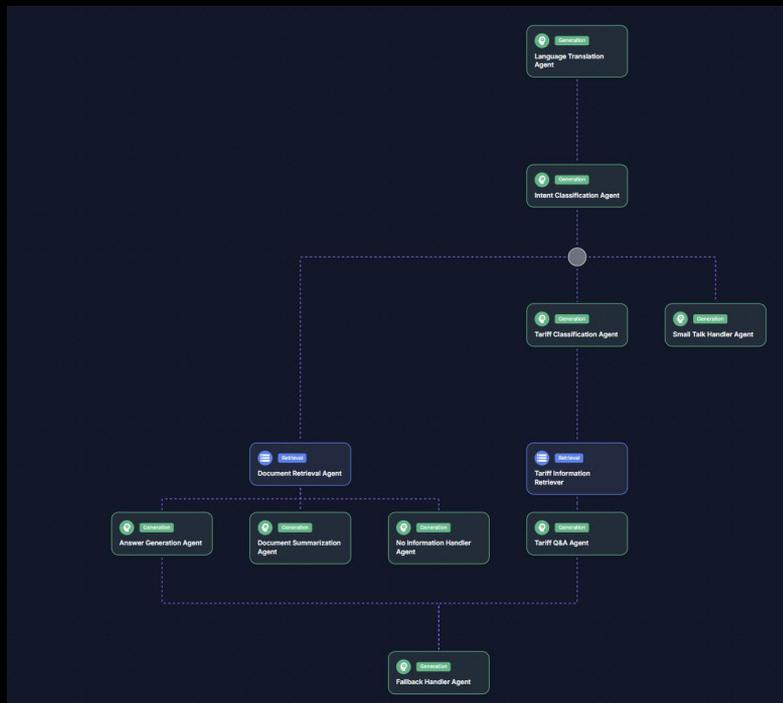
- Metrics: ConSens, WISDM, groundness, factual consistency.
- Continuous monitoring of outputs for quality and compliance.
- Alignment with EU AI Act and ISO 42001 frameworks.



European
Innovation
Council

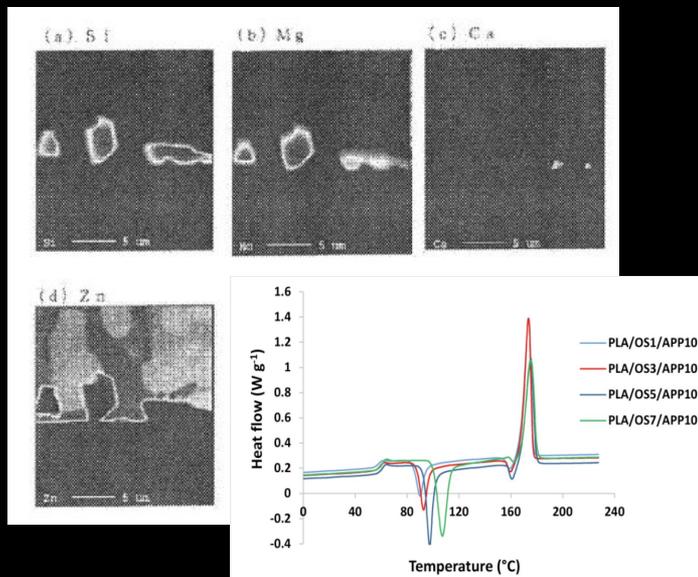
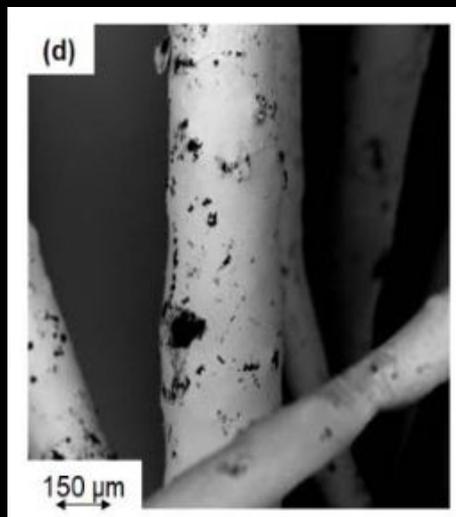


Telecom AI Contact Center



- Easter Europe Telecom
- Deep-knowledge AI customer contact center with **96+%** accuracy
- Escalation rate reduction to **15%**
- Built-in local language **68+**
- Repeatable use of assets & technologies for other use cases
- Optimise cost and efficiency with high quality output
- Maximize trade off between speed and accuracy

Process Manufacturing Scenario



- Collect product intelligence from research patents (IP) to enhance product capabilities
- Ability to contextualize over **68** different languages & custom images (i.e., Chinese patterns)
- Increase accuracy with thousands of documents
- Reduce time to knowledge by **20x**

“At our company, we are constantly looking for ways to optimise our R&D processes and this is what brought us to Iris.ai. Integrating the their tools into our process has made the ingestion and processing of external data significantly easier. It has cut weeks months out of our research and development timelines and gives us the capacity to review more patents.”

Iris.ai Vision

**Iris.ai – 10 years of
Experience serving
across industries,
R&D, and scientific
communities**

- Making Knowledge Machine-Readable.
- 10 years of R&D on scientific and technical text understanding.
- Foundation for explainable, grounded insights.

Let's build enterprise AI systems that are grounded, auditable, and scalable.

FIND US AT BOOTH G62

viktor@iris.ai

IRIS.AI