

# Humanitarian Foundation

—

## SemanticRAG POC



### The Challenge

A humanitarian foundation needed to demonstrate the ability of its Graph Retrieval Augmented Generation (GRAG) system to answer complex, cross-source questions. In particular, the task was to evaluate the impact of foundation investments on strategic goals by synthesizing information from publicly available domain data, internal investment documents, and internal investment data. The challenge lay in connecting diverse and unstructured information and also ensuring that the insights generated were precise, explainable, and actionable for executive stakeholders.



### The Solution

To address these challenges, the project team developed a proof-of-concept (POC) that leveraged advanced graph technology and a semantic RAG (Retrieval Augmented Generation) agentic workflow.

The solution was built around several core workstreams:

#### *Defining System Functionality*

The initial phase focused on establishing a clear use case: enabling the foundation to query its data ecosystem with natural language questions and receive accurate, explainable answers. This involved mapping out a comprehensive taxonomy and ontology that could encapsulate the knowledge domain of investments, thereby standardizing how investment documents and data were interpreted and interrelated.

#### *Processing Existing Data*

With functionality defined, the next step was to ingest and transform various data types. Structured data from internal systems and unstructured investment documents were processed and aligned with the newly defined ontology. Advanced techniques, including semantic extraction and graph mapping, were employed to ensure that all data—regardless of source—was accessible within a unified graph database.

#### *Building the Chatbot Model*

Central to the solution was the development of an investment chatbot that could leverage the graph's interconnected data. This was approached as a cross-document question-answering challenge. The model was designed to predict answers by linking query nodes with relevant data nodes across the graph, thereby addressing competency questions that a naive retrieval model would miss. An explainable AI component was integrated to transparently show which data points drove each answer, instilling confidence in the results.

# Humanitarian Foundation

—

## SemanticRAG POC

### *Deploying the Whole System in a Containerized Web Application Stack*

To ensure immediate usability, the POC was deployed, along with all of its dependencies, in a user-friendly, portable web application stack. This involved creating a dedicated API layer to interface between the chatbot and the graph database containers, alongside a custom front end that allowed executive users to interact with the system and view detailed explanations of the generated answers and the source documents upon which they were based. Early feedback highlighted the system's ability to connect structured and unstructured content seamlessly, paving the way for broader adoption.

### *Providing a Roadmap for Further Development*

Beyond the initial POC, the project laid out clear next steps. Recommendations included refining the chatbot's response logic, optimizing performance (notably in embedding and document chunking), and enhancing user experience through additional ontology-driven query refinements. These steps are critical for evolving the system from a demonstrative tool to a fully integrated component of the foundation's data management and access stack.



## The EK Difference

A key differentiator of this project was its adoption of standards-based semantic graph technology and its highly generalizable technical architecture.

The architecture comprises:

### *Investment Ontology and Data Mapping:*

A rigorously defined ontology underpins the entire system, ensuring that all investment-related data—from structured datasets to narrative reports—is harmonized under a common language. This semantic backbone supports both precise data integration and flexible query interpretation.

### *Graph Instantiation Pipeline:*

Investment data is transformed into RDF triples and instantiated within a robust graph database. This pipeline supports current data volumes and is scalable for future expansion. It includes custom tools to convert CSV files and other structured datasets into RDF and mechanisms to continually map new data into the graph.

### *Semantic RAG Agentic Workflow and API:*

The solution utilizes a semantic RAG approach to navigate the complexities of cross-document query answering. This agentic workflow is designed to minimize unhelpful hallucinations, ensuring that each answer is traceable back to the underlying data. The integrated API provides a seamless bridge between the front-end chatbot and the back-end graph, enabling real-time, explainable responses.

# Humanitarian Foundation

—

## SemanticRAG POC

Investment Chatbot Deployment:

Built as a central interface, the chatbot exemplifies how graph technology can be operationalized to address executive-level investment queries. It is fine-tuned to reflect the foundation's language and domain knowledge, ensuring that every answer is accurate and contextually relevant.



### The Results

---

The POC successfully demonstrated that GRAG could answer complex questions by:

- Delivering coherent and explainable recommendations that bridged structured and unstructured investment data.
- Significantly reducing query response time through a tightly integrated semantic RAG workflow.
- Providing a transparent AI mapping that allowed stakeholders to see exactly how each answer was derived.
- Establishing a scalable architecture that can be extended to support a broader range of use cases across the foundation's data ecosystem.

This project underscores the transformative potential of graph technology in revolutionizing how investment health is assessed and how strategic decisions are informed. With a clear roadmap for future enhancements, the foundation now has a powerful, next-generation tool for deep, context-driven analysis of its investments.

Enterprise Knowledge (EK) is a services firm that integrates Knowledge Management, Information Management, Information Technology, and Agile Approaches to deliver comprehensive solutions. Our mission is to form true partnerships with our clients, listening and collaborating to create tailored, practical, and results-oriented solutions that enable them to thrive and adapt to changing needs.

Our core services include strategy, design, and development of Knowledge and Information Management systems, with proven approaches for Data and Information Management, Knowledge Graph Implementation in support of NLP, ML, and AI initiatives, Taxonomy Design, Project Strategy and Road Mapping, Brand and Content Strategy, Change Management and Communication, and Agile Transformation and Facilitation. At the heart of these services, we always focus on working alongside our clients to understand their needs, ensuring we can provide practical and achievable solutions on an iterative, ongoing basis.