GRAPHWISE
AI THRIVES ON WHOLE DATA

Authors: Robert David & Atanas Kiryakov

WHITE PAPER

Building a Semantic Layer and Enabling Generative AI for Enterprises

An introduction to semantic layers, their applications for Generative AI, and the critical technology choices to implement them.

TABLE OF CONTENTS

Semantic Web standards were originally invented to publish machine-readable data with semantics on the Web. In time, they have proven to be sufficiently flexible and expressive to be used in enterprise contexts for solving data integration challenges by building knowledge graphs.

Nowadays, many enterprises increasingly use knowledge graphs as the core of their data strategy, creating semantic layers and data fabrics to best use their information.

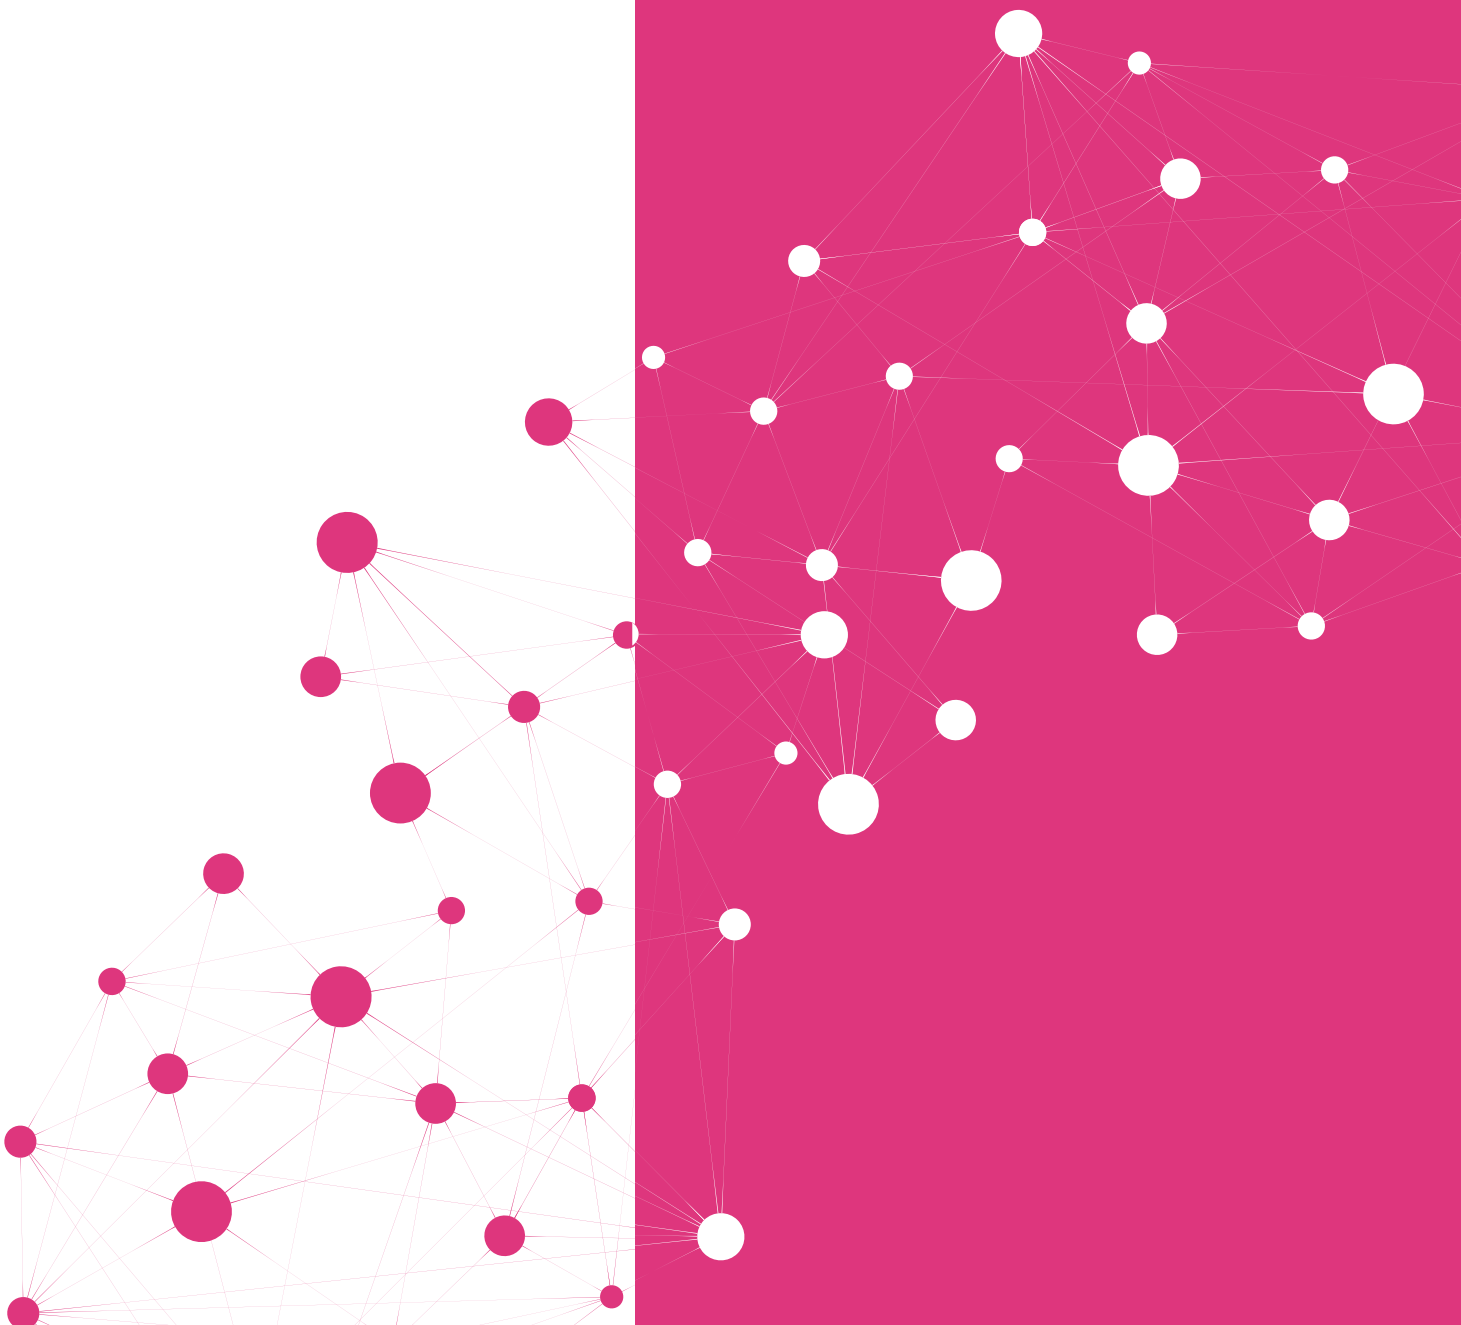

The challenge of data integration

The challenge addressed by semantic layers is the heterogeneity of enterprise data. This includes unstructured, semi-structured, and structured sources that need to be integrated consistently to provide a unified view across all data.

The essential role of the semantics layer is to create a common understanding of the meaning of enterprise data and thereby provide a consistent representation of business facts for humans and applications. This way it drives efficient data governance and supports data quality. It also enables inference - the process of surfacing relationships and deriving insights by connecting the dots across data silos.

Creating a common representation is especially challenging for unstructured data, which needs to be integrated in a way that interweaves it with other kinds of data sources. Such a challenge requires semantic tagging of documents using standards, like the Simple Knowledge Organisation System (SKOS) and Schema.org, and linking them to the structured knowledge in the semantic layer. The semantic layer is extended with metadata, which describes the documents and links them to the structured knowledge in the graph.

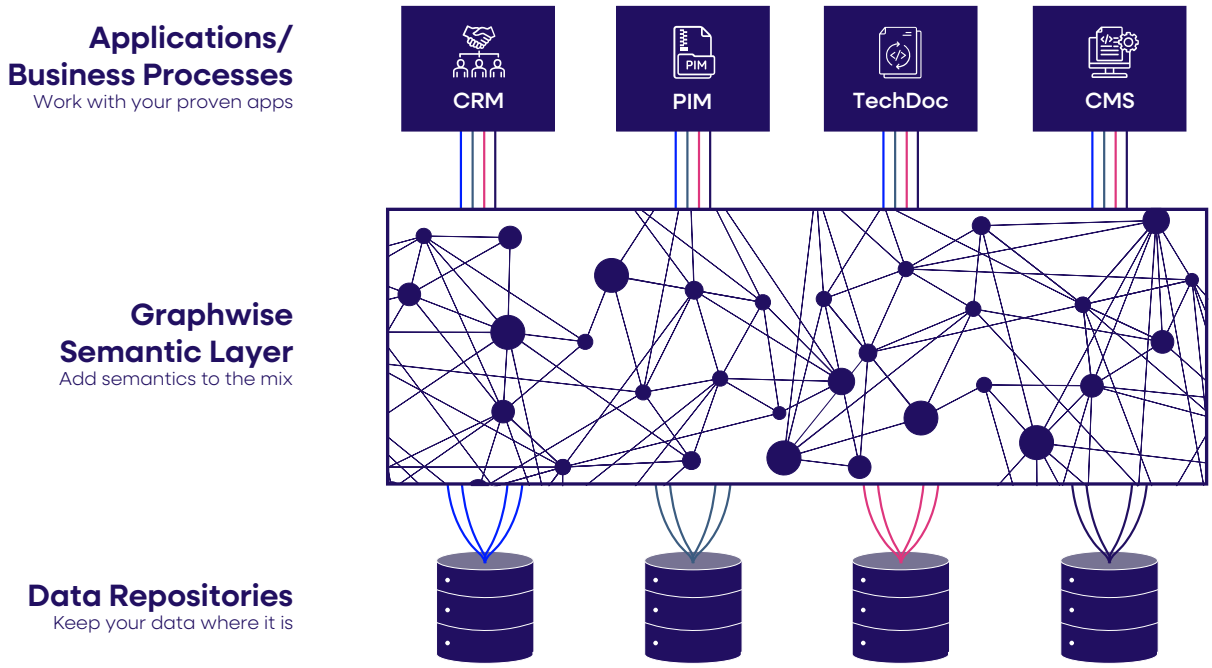

Understanding the meaning of your data

To execute a semantic layer strategy in an enterprise, it is crucial not only to have the data sources linked into one graph. It is also important to provide explicit semantic schemas to describe the information and align the meaning of the data elements. This will enable humans and applications to have a common ground and source of truth.

This way, everyone can understand enterprise data without ambiguities, avoid misinterpretations, and keep it consistent across the enterprises' data sources. Knowledge models with formal semantics provide the contract about the meaning of the data and allow for automatic validation and inference.

Semantic data models for standardizing domain knowledge

Existing ontologies, datasets, and vocabularies can be (re)used to build knowledge graphs without any additional modeling and integration efforts. If needed, they can be easily adapted, combined, and extended to serve specific use cases. The key enabler for reusing knowledge is the W3C's Semantic Web data model, called Resource Description Framework (RDF). It comes with a rich stack of standardized specifications for semantic schemes, query languages, protocols, and exchange formats. This allows interoperability, enables decentralized data architectures, and minimizes vendor lock-in.

This way proprietary data and content can be augmented with domain knowledge and 3rd party data, which allows for deeper interpretation of enterprises' information. End users can get additional insights as well as more complete and accurate results from search and data analytics. The same domain knowledge is precisely what Generative AI needs as context and grounding.

Enterprise-ready technology

Graph databases optimize for fast querying and standardized management of interconnected data. Data models like RDF make it easier to share, transfer, and process data. Reasoning using the Web Ontology Language (OWL) and the Shapes Constraint Language (SHACL) can be made most efficient by running it inside the database, as near to the data as possible.

Graph databases like Graphwise GraphDB provide enterprise-ready storage and processing for semantic layers, including ACID (atomicity, consistency, isolation, durability) compliance for database transactions, high availability, and enterprise-grade security. GraphDB has been designed and has matured for 20 years to deliver predictable performance across a wide variety of workloads. It is the only engine audited for both LDBC Semantic Publishing and Social Network benchmarks, proving that it efficiently handles graph analytics and metadata management.

Integration with content management systems (like Sharepoint and Adobe Experience Manager), SQL-speaking systems, and enterprise buses (like Kafka) make it easy to source data and content. They also help deliver multi-modal retrieval, involving full-text engines and document stores.

Intelligent applications providing valuable insights

Semantic layers easily consolidate enterprise information sources for a unified view on data. On top of that, semantic applications drive business processes and effectively reduce time and costs. The most basic applications are semantic search and recommendations for knowledge discovery. Enterprises can specifically model context-aware assistants for different personas and business tasks, thereby providing highly precise information for decision-making.

The versatility of knowledge graphs makes it easy to source data for downstream systems focused on analytic tasks like neo4j, Databricks, AI, and BI tools. Those benefit from the unified and enriched data and can thereby improve the quality of results.

Graph RAG - Generative AI you can trust

Reliable facts with precise semantics are specifically important when introducing Generative AI. Although it is a powerful new technology applicable to a variety of use cases, the main problem is the completeness and reliability of the information used to train it or to augment generation. Generated false information, commonly known as “hallucinations”, is a show-stopper in an enterprise environment when business decisions require precise and explainable facts.

A knowledge graph provides the necessary grounding to Generative AI to mitigate hallucinations and to feed proprietary data. This makes it an essential component and a reliable information source. With Graph RAG, we use the best of both worlds in hybrid architectures. We leverage the commonly used retrieval augmented generation (RAG) pattern in combination with semantic layers, to retrieve reliable and explainable information for decision-making.

The Semantic Web came a long way from the original idea to finally providing semantic layers for Graph RAG in enterprises. This establishes Graph AI as the next step for Generative AI. “Without including unstructured and semistructured data management into data fabric processes, the GenAI experience in the enterprise will continue to have major hallucination problems.” (Gartner®, “Emerging Tech: Data Fabrics With Multimodal Data Focus for Generative AI-Enabled Applications”, Radu Miclaus, Sharat Menon, Ehtisham Zaidi & Ramke Ramakrishnan, published 17 September 2024. GARTNER is a registered trademark and service mark of Gartner, Inc. and/or its affiliates in the U.S. and internationally and is used herein with permission. All rights reserved.). The same report points to “Semantic Web as a Central Methodology for Multimodal Data Fabric.”

Building on reliable standards

There are two major graph technology stacks building on top of two different data representation models: RDF and labeled property graphs (LPG). The story in the previous section elaborates on how semantic layers and multimodal data fabrics can be developed using the RDF stack.

On the other hand, the LPG model is designed with graph analytics in mind. Several LPG engines push the limits of scalable graph analytics of the type: find a path between two nodes in a graph or identify clusters of (more densely) connected nodes. The entire technology stack is optimized to minimize the learning curve and the cost for one-off graph analytics projects.

As good as it is for this purpose, this DNA also makes LPG inappropriate for semantic layers and other applications where data governance, integration, and quality are of the essence. The following capabilities not supported by LPG are a requirement for semantic layers:

- Standardized data schema language and serialization formats
- Formal semantics, reasoning, and validation capabilities
- Publicly available ontologies and datasets
- A flexible data model that can gracefully accommodate the structure and the semantics of diverse types of data, metadata, and knowledge

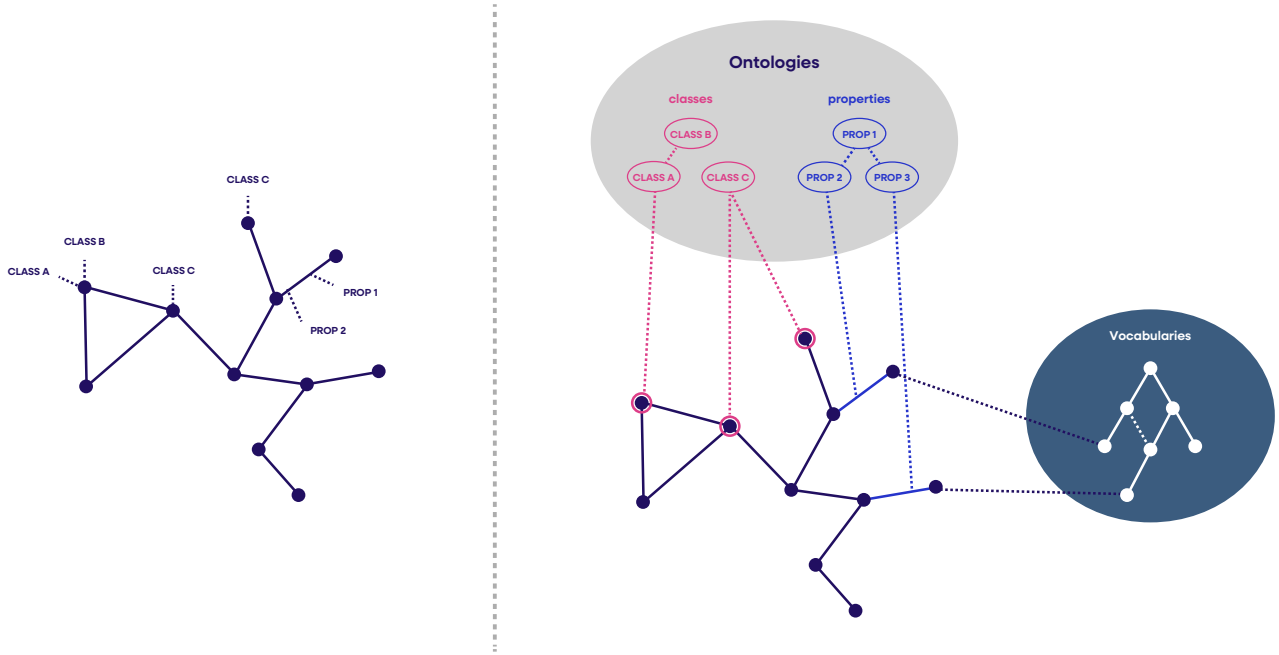

The semantic layer as the foundation of Enterprise AI

To wrap it up, RDF is the most appropriate technology stack for semantic layers because it enables efficient continuous data integration, unification, governance, reuse, and publishing. LPG technology fits well for one-off projects and downstream analytics. Graph data can easily be sourced from the semantic layer to LPG for these purposes and, if necessary, existing LPG data can also be imported into it.

Generative AI applications need access to unified data with clear semantics and as much context as possible. The task of searching for a path in a graph is irrelevant to Graph RAG applications, which strive to get holistic descriptions of entities, concepts, and relationships. Such applications need access to an RDF-based semantic layer.

What makes Graphwise GraphDB special is that it is the leading RDF graph database engine, which has been also proven (and audited!) to deal efficiently with graph analytics. To reduce integration costs and burdens, it integrates a long list of database and search engines as well as many data management platforms. It is also available both on-premise and from major cloud platforms.

To serve content management, knowledge management, and Generative AI applications well, on top of GraphDB, Graphwise offers the most comprehensive knowledge graph management platform. This includes PoolParty, which covers the whole linked data life cycle, and a list of accompanying tools and AI models.

Graphwise philosophy is to put knowledge engineers, data engineers, and scientists in control. So, it offers a toolset that allows architects to choose the most appropriate combination of tools and AI models to deliver an optimal solution in terms of both cost and performance.

GRAPHWISE

AI THRIVES ON WHOLE DATA

Graphwise enables organizations to unlock ROI for enterprise AI by delivering the most comprehensive and trusted industry solution in the field of knowledge graphs and semantic AI technologies. As enterprises pour millions into AI investment, Graphwise delivers the critical knowledge graph infrastructure to ensure enterprises are ready to realize the technology's full potential, is trusted, and can be implemented at scale.

Graphwise is the result of the merger between tech visionaries Ontotext and Semantic Web Company who are vendors of GraphDB and PoolParty. The Graphwise Product Suite comprises GraphDB's RDF database and PoolParty's taxonomy management and NLP services to help consolidate data from various formats and identify hidden connections between data objects.

These capabilities serve as the foundation for enterprises to scale their systems effortlessly, while ensuring data and content remains well-structured and classified. Graphwise has over 200 employees worldwide, with offices located across North America, Europe and APA

graphwise.ai

Graphwise New York
Ontotext
Select Office Suites Flatiron
1115 Broadway, 16 Madison Square West
Suite 1200, New York, NY 10010,
United States
+1 929 239 0659

Graphwise Sofia
Ontotext
Twins Centre
fl.3, 79 Nikola Gabrovski str.
1700 Sofia
Bulgaria
+359 2 974 61 60

Graphwise Vienna
Semantic Web Company GmbH
Neubaugasse 1, Top 8
1070 Vienna
Austria
+43-1-4021235