

# How Knowledge Graphs Enhance GenAI Outcomes

Together, Neo4j, Deloitte, and Google Cloud Platform help organizations meet the GenAI challenge in the enterprise

In a [recent Deloitte survey](#) of business and tech leaders, 79% of respondents said they expect GenAI to transform their business in the next three years. We are witnessing not a wave but a sea of change that spans industries and geographies.

Businesses that tap into GenAI's potential *will* have a strategic edge. But this advantage comes with a contingency: For GenAI to deliver enterprise value, it must provide insight into the organization's own data. At the heart of GenAI transformation lies a question that has persisted since the advent of digital data itself: How do organizations effectively centralize a staggering amount of structured and unstructured data and make it useful for decision-making?

Knowledge graphs are the cornerstone of GenAI's ability to provide tangible business value. Knowledge graphs enable organizations to connect data spread across myriad systems and sources, delivering insights that drive business transformation.

#### **Why Knowledge Graphs Are Business-Critical**

If you ask Gopal Srinivasan, Deloitte's Technology Sector AI Lead and Alphabet Google AI Alliance Lead, who should pay attention to knowledge graphs, his answer is simple: everyone.

"The amount of data any enterprise is dealing with has exploded exponentially, and so traditional approaches of how data is encoded, represented, and managed are difficult to scale," says Srinivasan. "And that's where knowledge graphs come in."

#### **What Is a Knowledge Graph?**

A [knowledge graph](#) is a structured data representation of a collection of entities (nodes) and the relationships (edges) connecting them. It's a massive, interconnected web of everything your organization knows, laid out so AI can traverse from one fact to the next. Through inferential reasoning, knowledge graphs provide context-driven insights and a deeper understanding of the topics, filling in the gaps of vector-only databases.

“ Knowledge graphs enable drawing new conclusions and discovering implicit knowledge based on the explicit facts and relationships encoded within the graph...This capability for inference and reasoning makes knowledge graphs especially powerful tools for tasks such as building recommendation systems and complex data analysis, where understanding context and relationships is crucial.”

— Excerpt from “[Responsible Enterprise Decisions with Knowledge-enriched Generative AI](#)” by Deloitte and Vrije Universiteit Amsterdam



Even before the advent of GenAI, knowledge graphs were transforming data analytics in industries ranging from healthcare to [transportation](#) and the [energy sector](#). And now, that benefit is poised to increase exponentially. When knowledge graphs power GenAI, two business-critical outcomes emerge:

- **Enterprise knowledge is accessible and useful.** Organizations need to be able to ask complicated, inference-driven questions about their own data. The ability to query enterprise domain knowledge for insight, a use case called enterprise search, has been a considerable source of value. Previously, asking intricate questions about organization-specific data entailed significant work, requiring a handful of people with the proper permissions and technical acumen to create and interpret tables manually.

The process was labor-intensive, time-consuming, and costly. Knowledge graphs remove this technical burden, democratizing data access and making it easier and faster for people to ask questions and make informed decisions.

- **Returned data is explainable and accurate.** Organizations must be able to trust the data that underpins their choices. Tracing data lineage — the ability to track data back to its exact source used by GenAI in delivering a particular response or insight — is essential, especially in industries with stringent regulatory compliance and safety protocols, such as healthcare.

### When GenAI Meets Knowledge Graph

What do you get when you combine the proven value of knowledge graphs with the limitless potential of GenAI? The [most powerful way](#) to gain smarter insights from enterprise data.

To start, there's a [highly symbiotic relationship](#) between GenAI Large Language Models (LLMs) and knowledge graph construction. LLMs make it faster

and easier to create organization-specific knowledge graphs from disparate sources of information. Increasingly, LLMs can derive context from entirely [different types of information](#) (such as slide decks, presentation videos, PDFs of organization-specific documentation, code, employee data, and so on) and store it all in the same knowledge base.

Regardless of how you build your knowledge graph, [adding GenAI as the layer between the knowledge graph and the user](#) makes all this centralized insight accessible for enterprise search via conversational, humanlike language.

#### Knowledge Graphs and RAG

Knowledge graphs offer an effective means of making information accessible to GenAI for [retrieval-augmented generation](#) (RAG). Retrieval-Augmented Generation (RAG) is a machine learning technique that infuses GenAI with information from external sources, such as enterprise data, enabling users to ask questions and receive answers conversationally through natural language processing.

GraphRAG, a subset of RAG that leverages knowledge graphs, employs an LLM-generated knowledge graph to perform complex queries on private datasets, enabling sophisticated data analysis and insights. Not only does [GraphRAG demonstrate impressive inference and contextual reasoning](#), but it also dramatically increases the reliability and traceability of information returned.

### Minimize Risk, Maximize Benefits

We know this by now: Like all burgeoning and powerful tech, as GenAI continues to iterate and scale, so too do its risks. In fact, only 25% of directors and C-suite leaders believe their organizations are “highly” or “very highly” prepared to address governance and risk issues related to GenAI adoption, according to a [2024 Deloitte survey](#). And according to an [analyst report from Enterprise Strategy Group](#), many organizations have faced challenges in transitioning from experimental to enterprise-ready GenAI, with nearly 40% of organizations citing “difficulty validating and evaluating results, employee hesitancy to trust recommendations, and ethical considerations and biases in generated content.”

Deploying GenAI with a knowledge graph mitigates many of the challenges that can arise from using GenAI alone. Together, they allow organizations to:

- **Reduce hallucinations and improve effectiveness.** Knowledge graphs power more accurate, transparent, and explainable GenAI outcomes by dramatically reducing inaccurate responses and ensuring that all material is traceable. They also improve overall GenAI system quality and effectiveness.
- **Gain better insights.** A connected view of domain knowledge allows organizations to find contextually related information, uncover hidden relationships, and gain richer insights. When answers can be delivered nearly as quickly as questions are asked, organizations can adapt and pivot use cases to ensure the best business outcome.
- **Achieve greater predictive accuracy.** Knowledge graphs combined with GenAI enable downstream models and analysis. Organizations can then understand patterns, quality, and potential biases in grounding data so they can adapt for continuous improvement.

### Understanding the Stakes

To understand just how valuable it is to extract the right insights from enterprise data, Sudhir Hasbe, Chief Product Officer at Neo4j, says to consider the question from the perspective of a global pharmaceutical company with centers and suppliers all over the world. If a hurricane causes a major distribution center to go down for three days, that organization needs to know all possible downstream impacts. Which drugs would be affected? What are the implications for suppliers? How can they ensure critical medications get to pharmacies and hospitals on time?

Ask a typical chatbot these kinds of supply chain analysis questions, and you'll get answers based on the original data that the tool was trained on. This could be the Internet, books, or public data. With the context provided by a knowledge graph, answers can come from your own enterprise data.

“Now, there is knowledge. There are facts. There is context,” Hasbe explains. “And when you combine all this together, you get magic. Because now, you can get the true answers to your questions.”

“*Knowledge graphs provide the perfect complement to LLM-based solutions where high thresholds of accuracy and correctness need to be attained.*”

— Excerpt from “Technical Insights: Why Knowledge Graphs Are the Missing Link Between Data and AI Engineering”, a presentation at the [Gartner Data & Analytics Summit](#)

For GenAI to deliver business value that guides decision-making, it must be powered by an organization's specific enterprise knowledge. With knowledge graphs in play, as countless [real-world use cases](#) demonstrate, organizations can move forward in their GenAI transformation for a strategic edge.

#### Four Steps, Endless Opportunity

GenAI is a singular generational technology because, as Hasbe puts it, "It impacts everything." That stunning reality encompasses both opportunity and overwhelm: If the scale of change is everything, where do you start?

Regardless of your sector, there are four steps toward a successful GenAI project:

1. Identify a business problem.
2. Validate that it's the right problem for the technology.
3. Start small. It's tempting to tackle complicated problems right away, but instead, determine exactly what you want to deploy in the initial use case and prove it.
4. Find the right team to work with — one that will help you identify the right problem, understand the technology, and implement the most impactful solution.

The newly announced partnership between Neo4j, Google Cloud Platform (GCP), and Deloitte is bringing industry-leading capabilities together to help organizations across sectors solve the GenAI challenge for the enterprise. "This is about making sure we can put the best technology in the hands of customers and deliver impactful business value," says Advait Bopardikar, Outbound Product Manager for Cloud AI at GCP.

As the primary cloud provider [leading the AI wave](#), GCP provides integration with Vertex AI, an ML development platform, enabling enterprises to build and deploy improved graph-based machine learning models. GCP's relationship runs deep with Neo4j, the industry leader in knowledge graph technology and the default graph database tool for [Google Vertex](#) and [BigQuery](#) engagements. Together, GCP and

Neo4j make it seamless for organizations to [implement knowledge graph technology at scale](#).

Coming alongside GCP and Neo4j, Deloitte brings deep business acumen and the ability to identify complex organizational problems across industries [that can be solved by deploying GenAI](#) with knowledge graphs. The organization is already working with over 150 clients globally to achieve value-drive AI deployments, and a significant number of those use cases have applied a GraphRAG approach.

"The tri-alliance ecosystem allows us to get to solving the problem faster," says Srinivasan. "It's really as simple as that."

The enterprise GenAI advantage is real, and it's powered by knowledge graphs. For more information on GenAI and Neo4j, please contact:

**Manoj Suvarna, Deloitte** ([manoj-suvarna@deloitte.com](mailto:manoj-suvarna@deloitte.com))

**Azzy Maizer, Deloitte** ([amaizer@deloitte.com](mailto:amaizer@deloitte.com))

**Rupal Jain, Google** ([jainrupal@google.com](mailto:jainrupal@google.com))

**Nathan Barney, Neo4j** ([nbarney@neo4j.com](mailto:nbarney@neo4j.com))