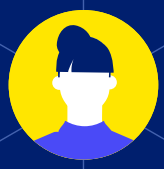
Progress® Data Platform











Closing the Accuracy Gap in GenAI with Progressive Graphs: A Case Study

WHITEPAPER

Table of Contents

Abstract / 3

The Accuracy Challenge / 3

- Graph RAG / 4

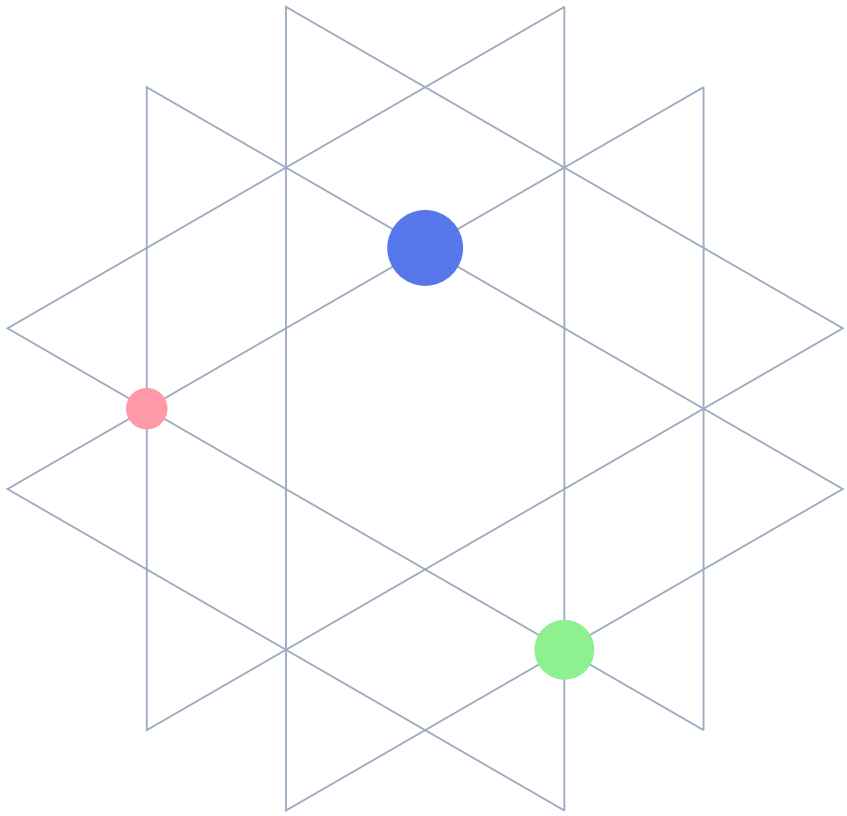
Enhancing Graph RAG with Human-Validated Input / 6

- Progressive Graphs / 6
- The Hybrid Approach / 7
- The Technology Stack / 8

Case Study / 9

- Project Overview / 9
- Methodology / 9
- Results / 10

Conclusion / 14



© 2026 Progress. All Rights Reserved.

Abstract

As firms adopt agentic and generative AI (GenAI) into their products and offerings, maximizing the accuracy of foundational models becomes paramount. Contextual relevance, along with security and scalability for AI implementations, helps reduce business risk by increasing trust, transparency, production readiness and governability.

Organizations have adopted many different approaches to maximizing accuracy, including context engineering, governance validation, model tuning, running queries against multiple large language models (LLMs) and comparing answers to find discrepancies. However, while technology and data quality practices continually improve, AI systems may still provide unsatisfactory answers.

This whitepaper presents a case study on Progressive Graphs, an alternative accuracy-enhancing technique that captures and retains real-time human feedback, enabling AI systems to continually learn and stay true to the industry's or business domain.

Progressive Graphs encapsulate the approach Progress takes on its path toward trusted AI. By combining a semantic retrieval-augmented generation (RAG*) architecture with human-validated input, they incrementally enrich the body of knowledge enterprise AI systems use, enabling a dramatic increase in accuracy from 48% to 90%.

The Accuracy Challenge

LLMs and the vector databases used in RAG* implementations learn from statistical patterns in large text corpora and vector embeddings , rather than from an explicit, verified knowledge base. While this is a powerful approach, it's limited: answers are based on what is most likely given the data and may not be completely accurate.

For example, if a term or concept rarely appears in the training data, a vector-only RAG system may fail to retrieve relevant information even when it exists. On the other hand, it may surface irrelevant content simply because it contains words that are commonly associated with the user's query.

© 2026 Progress. All Rights Reserved.

When an LLM's training has adequately covered the area relating to the question, the LLM is more likely to provide a correct result. This is one important reason that new LLM releases are more accurate than the versions they replace. Because when an LLM lacks sufficient knowledge about a topic, it relies more heavily on statistical patterns, increasing the risk of plausible—but potentially incorrect—answers.

That's why domain-specific knowledge is essential when using AI systems in highly regulated industries such as life sciences, financial services, healthcare and public safety or government, where risk tolerance is low.

***Retrieval-Augmented Generation (RAG)**

RAG connects GenAI models with an organization's private, trusted data to provide accurate insights that are grounded in context. Using internal data sources to reduce hallucinations, RAG allows for reliable AI responses that can adapt to data changes without the need for retraining the model.

Graph RAG

RAG implementations that use domain-specific knowledge captured in ontologies, stored in knowledge graphs and retrieved via semantic search* are often referred to as Graph RAG or ontology-driven AI. Rather than relying solely on statistical relationships, these approaches model the actual semantic relationships between concepts and data. This helps ground AI responses in validated knowledge, improving accuracy and relevance.

Adding graphs to RAG implementations has been widely recognized in the industry, but adoption has been slow. Here's why:

© 2026 Progress. All Rights Reserved.

- Time and Cost Efficiency: The generation of knowledge graphs can be time- and labor-intensive—and there is a concern that the effort involved is not cost-effective. Specifically, there is a belief that manual graph creation requires substantial amounts of skilled labor, which often lack needed business domain knowledge. The second concern is that knowledge graphs sometimes quickly go out of date and can be difficult to maintain.

- Accuracy and Risk: Luckily, GenAI is being adopted to help generate ontologies as a strategy to compensate for the lack of skilled experts or time inefficiencies this causes. When the creation of the graph is outsourced purely to GenAI, however, enterprises run the risk that the resulting graphs will include invalid and irrelevant terms; in turn, undermining its value.

***Note on Semantic Search**

In AI and data management, the term semantic search can refer to two different approaches that are often conflated: querying vector-based data (vector embeddings) and querying graph-based data. Both aim to understand meaning, but they represent and retrieve information differently.

Vector embeddings represent words, images and other content as numerical points in a high-dimensional space, enabling similarity-based searches. Graph data (RDF triples), including metadata and semantic triples used in knowledge graphs, captures explicit relationships between entities, i.e.,—who or what is connected to whom and how.

Both approaches can improve search relevance, but they work differently. Vector search finds conceptually similar content based on proximity in the vector space, while graph-based search follows defined relationships between entities using graph indexes and query languages like SPARQL. Together, they provide complementary ways to retrieve and reason over information.

© 2026 Progress. All Rights Reserved.

Enhancing Graph RAG with Human-Validated Input

Progressive Graphs

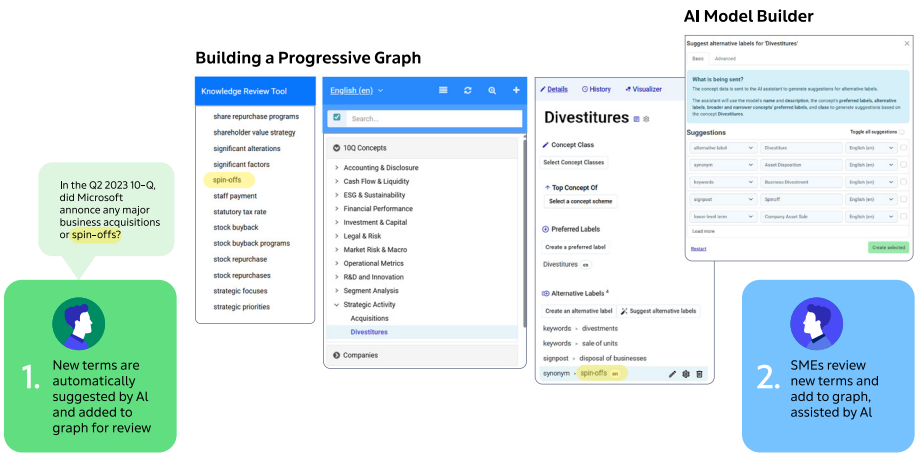
Progressive Graphs is an approach designed by Progress® to help teams quickly generate and incrementally enrich accurate knowledge graphs that are tailored to a specific GenAI deployment. Embedding them into AI systems and workflows significantly improves the accuracy, relevance and reliability of AI-generated responses.

Progressive Graphs expand and update their underlying knowledge model over time as users search, ask questions and interact with your RAG-based AI systems. They leverage LLMs to suggest new concepts coming from user queries that are not part of your knowledge graph. Subject matter experts (SMEs) then review and validate these additions, helping keep the production knowledge model accurate and current.

This flow is designed to be easy to implement and can lay on top of existing GenAI or agentic AI applications, so SMEs spend minimal time building or updating the graphs per use case.

A major benefit of Progressive Graphs is that the knowledge graph is dynamic and stays current. As users ask new questions and use different terminology, the graph evolves to reflect changing information needs and business context.

Figure 1. Progressive Graph Workflow Illustration



© 2026 Progress. All Rights Reserved.

Figure 1 shows the Progressive Graph workflow for creating an application-specific graph from an initial skeleton. The process is as follows:

1. A user asks about “spin-offs”
2. The observing GenAI determines this term is not in the graph and suggests it
3. An SME decides whether the term should be included as a synonym of “divestiture”
4. The system recommends related terms and concepts that may also be worth adding to the knowledge graph

The result is a knowledge graph that evolves alongside user needs while remaining grounded in expert validation and governance.

The Hybrid Approach

Progressive Graphs can be incorporated into a hybrid approach to GenAI that leverages multi-model data and retrieval strategies for maximum accuracy within your RAG implementation.

You can apply a tailored mix of data and indexes for your specific use case. For example, you can start with vector search, then add keyword and metadata search capabilities. As your needs evolve, you can incorporate geospatial data and ontology-based graph relationships.

Most industries have standards groups that author graphs for particular verticals or use cases, and you can incorporate these too as governed starting graph layers. By combining vector search with industry knowledge graphs, organizations can achieve acceptable levels of accuracy while they build a more advanced solution.

To achieve this hybrid approach in practice, we built the Progressive Graphs search flow to combine the top results from each retrieval technique using Progress technology and its built-in Reciprocal Rank Fusion (RRF) algorithm.

© 2026 Progress. All Rights Reserved.

Reciprocal Rank Fusion (RRF) is a popular algorithm used in search and AI systems to combine and re-order relevance scores of search results from multiple independent retrieval methods (like keyword search and vector search) in a single list.

The Technology Stack

The Progressive Graphs workflow is backed by the Progress® Data Platform capabilities. Our semantic AI and modeling suite, the Progress® Semaphore™ platform, allows us to generate industry-specific ontologies and handles the progressive enrichment of the knowledge graph. Our multi-model data management platform, Progress® MarkLogic®, handles structured, semi-structured and unstructured content together with the knowledge graphs in a single platform.

By combining documents, graphs and search indexes with 100% transactional consistency, you benefit from using multiple data models to achieve the richest context possible, while also using a hybrid retrieval approach to improve the quality and accuracy of RAG-based applications.

In the following case study, we leverage the Reciprocal Rank Fusion algorithm in the MarkLogic platform to unify relevance ranking by combining results from vector, keyword, graph-based and geospatial searches. Initially, you can rely heavily on vector search, then adjust ranking weights as the knowledge graph becomes more complete. This enables a smooth transition to a hybrid retrieval model while maintaining high-quality search results throughout the process.

© 2026 Progress. All Rights Reserved.

Case Study

Project Overview

We've compared the AI response accuracy between a vector-only approach to a RAG system and one using the Progressive Graphs solution. We loaded publicly available 10-Q documents into the MarkLogic multi-model database, then built an application interface to query them using our own UI development toolkit. A user interface enabled us to compare outputs from both approaches by displaying their responses to the same questions side by side.

We borrowed questions and a base set of generated answers from an open-source project, along with some additional questions to test the system.

The interface provided visibility into the entire retrieval process, including the query sent to the LLM, its supporting documents and the graph-derived terms used to enhance the search. It also displayed what documents the LLM used in its response, along with links to the documents.

Methodology

The environment was obtained and the GenAI application and infrastructure was set up in less than three days. The pilot began with a small set of PDF documents and a representative set of 200 sample user questions. Limiting the initial dataset made it easier to validate that the LLM was retrieving and using correct information.

After testing, we spit the 10-Q filings into smaller chunks of about 1 KB each to reduce costs and improve retrieval performance. In addition, we removed the display formatting from the PDFs, leaving only the underlying content, including text-based tables.

After running the questions against the application in bulk, it suggested a large number of terms. An SME evaluated the output and either rejected or accepted the terms (although most were rejected and replaced with more valuable terms). The document chunks were then classified against the knowledge graph to improve retrieval. For example, documents were tagged with terms like "most recent" to help the system identify the latest published content when responding to user queries.

Next, we tested the updated knowledge graph using the original questions, along with some further questions designed to evaluate a range of scenarios.

© 2026 Progress. All Rights Reserved.

Results

We evaluated the answers against a definition of a correct answer, which required them to have:

- No errors
- No important facts left out
- Correct links to the data used for the answers
- Verified data within the database needed to answer the question correctly

Figure 2. Correct Results

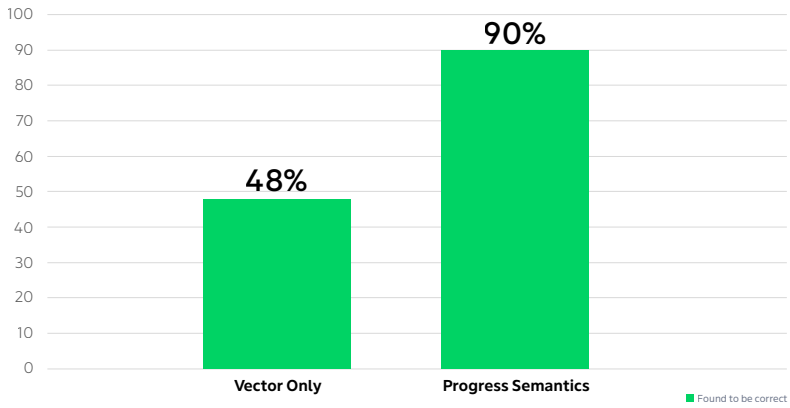


Figure 2 shows a comparison between the vector-only and Progress Graph approaches. While the vector-only approach answered 48% of questions correctly according to the criteria, the Progressive Graph approach achieved a 90% accuracy rate.

Figure 3. 10-Q Results Found to Be Inaccurate

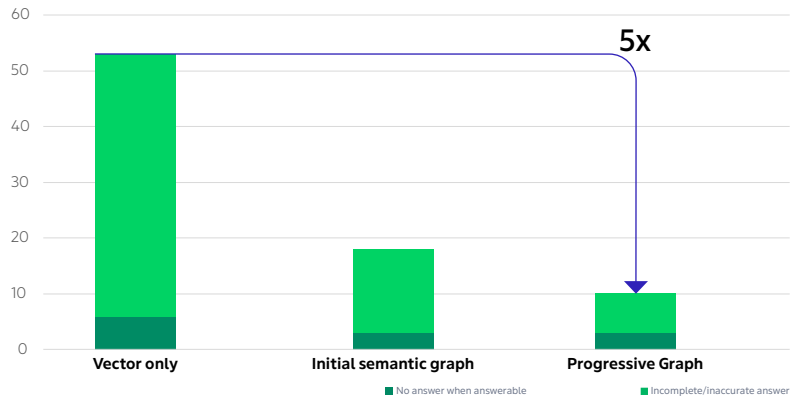


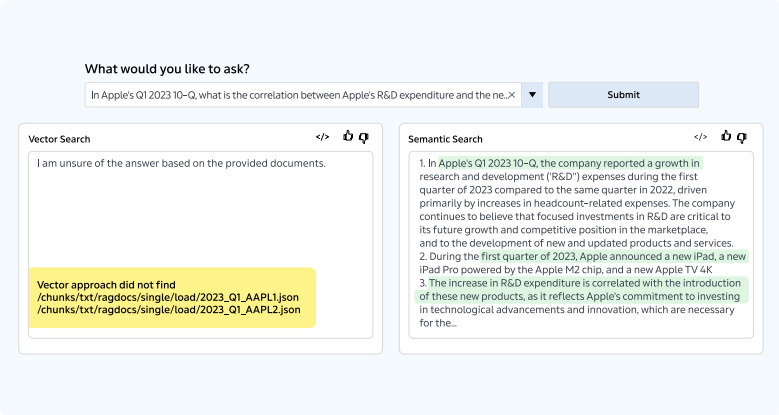
Figure 3 shows the most common causes of incorrect responses. Although some questions received no answer, most errors involved missing information or inaccurate facts.

© 2026 Progress. All Rights Reserved.

Further investigation showed that in many cases, the RAG implementation using Progressive Graphs sent the correct data to the LLM, but did not use it. **However, the RAG implementation was better at determining what data should be used to construct an answer than the LLM, even when the LLM was only given a small subset of the overall data.**

Examples:

Figure 4. No Answer vs. Expanded Answer



This example illustrates a common limitation of vector-only retrieval: The information needed to provide an accurate answer to the question was distributed across multiple document chunks. By enriching those chunks with knowledge graph classifications, Progressive Graphs could be used by the RAG system to retrieve all relevant data to provide the LLM with the context it needed to generate complete and accurate answers.

Figure 5. Wrong vs. Correct Answer

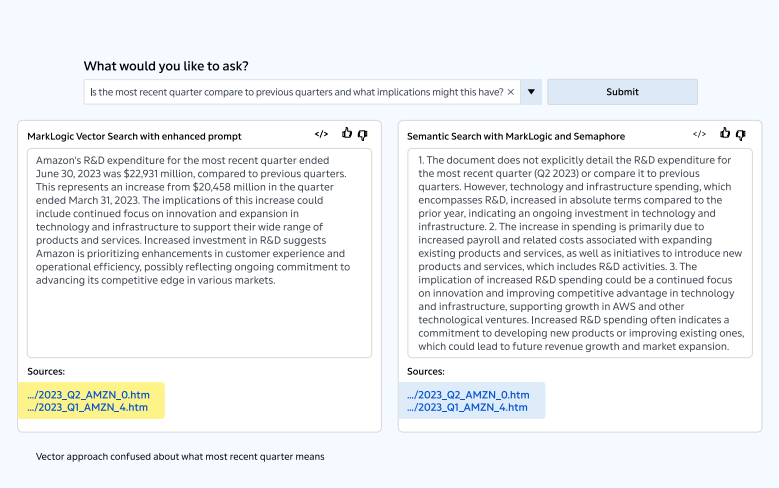


Figure 5 demonstrates that to answer the question correctly, the LLM must recognize that the user is asking about the most recent quarter and retrieve information from the appropriate reporting period.

© 2026 Progress. All Rights Reserved.

LLMs often struggle when answers depend on retrieving specific data rather than recognizing language patterns. By enriching documents and document chunks with relevant classifications, Progressive Graphs help facilitate that the LLM retrieves the right information.

Figure 6. No Answer vs. Correct Answer

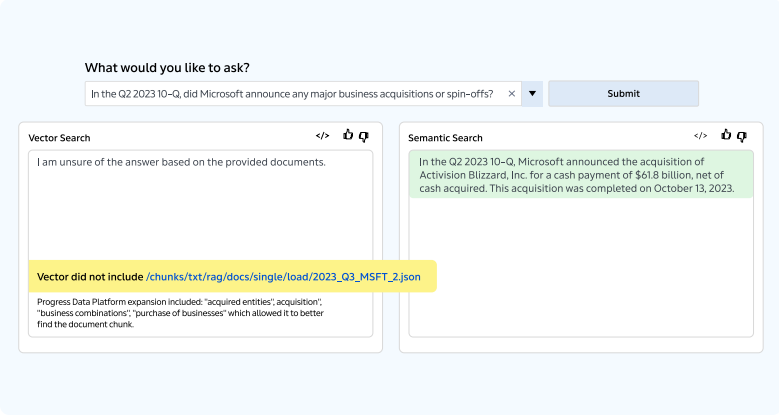


Figure 6 demonstrates the core benefit of Progressive Graphs. The user asked about “spin-offs,” a synonym for “divestiture.” While the vector-first approach would have provided the correct answer for “divestiture,” in this case, it did not retrieve the appropriate document since “spin-offs” is a less common term. Running a semantic search encountered the same issue on the first try, because “spin-offs” was not in the original knowledge graph.

In contrast, the RAG implementation with the underlying Progressive Graph suggested that the term “spin-offs” be added to the knowledge graph as a synonym for “divestiture.” The next time we ran the test, the application provided the correct answer.

© 2026 Progress. All Rights Reserved.

Going Beyond 90% Accuracy

Although increasing accuracy from 48% to 90% delivers meaningful business value, certain use cases demand nearly 100% accuracy. The Progress Data Platform helps organizations improve retrieval accuracy further with the following capabilities:

1. **Ranking Quality:** We used high-quality data in the 10-Q project. However, not all projects start with a clean, AI-ready data set. Results ranking should take this into account, consider the general quality of the source and its recency.
2. **Traceability:** Understanding the lineage of a document or record is essential in determining its reliability.
3. **Vocabulary:** This whitepaper has discussed Progressive Graphs, which are meant to be deployment-specific; however, there are many preexisting vocabularies and knowledge graphs that you can use immediately to enhance responses.
4. **De-duplication:** Many duplicate documents exist in a RAG system's data sets. If the LLM uses duplicate or near-duplicate data to respond to a query, accuracy can deteriorate. In other words, de-duplicating data sets improves accuracy.
5. **Context:** A fact may be highly relevant in one context and irrelevant in another; determining that context often requires more sophisticated knowledge graph relationships than basic retrieval approaches can provide. Context helps identify which information is relevant and how it should be interpreted and used in responses. Simple word matching and basic graph models often struggle to make these distinctions.
6. **Enhanced Relevancy Rankings:** A critical finding of the 10-Q project is that LLMs will not always use records containing data that corresponds to a question. So it's important to *limit the body of information to include only relevant data* and use sophisticated relevancy rankings. This will force the LLM to base its response on the correct data.

In addition to Reciprocal Rank Fusion discussed in this white paper as one means to achieve this, the Progress Data Platform offers advanced techniques for enhancing and tuning results relevance:

- boosting the relevancy of specific terms based on the user's interests,
- boosting results for specific sections of a document (its metadata, title, etc.),
- boosting based on data quality and adding terms to queries based on user profiles.

In summary, the Progress Data Platform architecture helps maximize accuracy for any LLM use case.

© 2026 Progress. All Rights Reserved.

Conclusion

This case study reinforces the notion that high levels of accuracy are a function of model sophistication and system efficiency when incorporating domain knowledge, contextual relevance and continuous feedback.

While vector-based RAG approaches lay a powerful foundation for AI systems, they are predictive, not deterministic, which is insufficient for high-stakes applications.

The Progressive Graphs approach offers an effective and scalable way to enhance accuracy. As demonstrated in this whitepaper, by combining semantic graph structures with real-world user interactions and human validation, you can continuously refine the knowledge layer and improve response quality by up to 40%.

Importantly, this approach does not require a complete re-architecture of existing GenAI systems. Instead, it complements and enhances established RAG frameworks through a hybrid model that balances statistical inference with governed knowledge. Over time, Progressive Graphs evolve alongside users' needs so AI systems remain aligned with business context, terminology and expectations.

As enterprises pursue increasingly critical AI use cases, especially in regulated industries, the ability to deliver trustworthy, transparent and contextually accurate outputs will become a defining requirement. In this vein, Progressive Graphs provide a viable foundation for achieving this goal, offering a repeatable method to improve accuracy while maintaining scalability and governance.

Talk to a Progress expert to see how you can achieve comparable results with your data and the Progress Data Platform.

About Progress Software

Progress Software (Nasdaq: PRGS) empowers organizations to achieve transformational success in the face of disruptive change. Our software enables our customers to develop, deploy and manage responsible AI-powered applications and personalized digital experiences with agility and ease. Businesses of all sizes get a trusted provider in Progress, with the products, expertise and vision they need to turn AI disruption into a competitive advantage. Millions of developers and technologists at hundreds of thousands of organizations depend on Progress every day. Learn more at www.progress.com

/progresssw
/progresssw
/progresssw
/progress-software
/progress_sw_

© 2026 Progress Software Corporation and/or its subsidiaries or affiliates.
All rights reserved. Rev 2026/06 | RITM0373106

Progress