SUCCESS STORY

EY applies search with Elastic to pioneering generative AI experience for finance

EY puts Elastic at the heart of GenAI experience for financial sector

EY uses Elastic's ESRE features to unlock valuable insights from unstructured data for banks, enabling financial institutions to take advantage of generative AI to streamline regulatory compliance, boost innovation, and mitigate risk.

Region
US

Industry
Education / Non-Profit

Solution
Elastic Search

Ensures information accuracy

- EY built a generative AI experience that helps clients unlock insights from unstructured data with Elastic's advanced search technology.

Seamless integration with generative AI stack

- With Elastic, EY can easily integrate cutting-edge AI search features with other components of its generative AI stack.

Delivers generative AI innovation at scale

- With Elastic, EY's generative AI experience can handle large volumes of diverse documents and seamlessly scale across many organizational units.

Elastic's AI search toolkit enables EY finance clients to unlock valuable insights for compliance and innovation

For financial institutions, rapid access to trusted information is more important than ever, especially considering increased pressure to comply with evolving regulations like capital adequacy and ESG (Environmental, Social and Governance) standards. For this reason, many banks and their partners are exploring the potential of generative AI to unlock value from vast quantities of unstructured data including PDFs and other documents.

Ernst & Young (EY), one of the largest professional services networks in the world and a leading provider of cutting-edge solutions to the financial sector, is working hard to support banks through these challenges. Vishaal Venkatesh, GenAI Manager, EY, says, "Our primary focus is to champion responsible AI and deliver these solutions effectively to our clients. We ensure they not only understand the technology but also possess the proper framework to implement it successfully."

To achieve this aim, EY has put Elastic Search, built on the Elasticsearch platform, at the heart of its generative AI solution. "Elastic's cutting-edge work in search and retrieval attracted us," says Venkatesh. "We integrated these capabilities into our AI stack for Retrieval Augmented Generation (RAG), which improves accuracy and accelerates retrieval of insights from unstructured data."

elastic

Driving innovation, mitigating risk

Central to EY's solution is ESRE (Elasticsearch Relevance Engine), a suite of tools that combines machine learning models, data transformation, and a vector database and data search and retrieval.

Venkatesh says, "With ESRE, we can use machine learning models to generate vector embeddings stored at large scale. We used the super-efficient embedding model to build out the RAG component, which significantly reduced our processing time." EY also took advantage of enhanced indexing and chunking available through the Elastic product suite.

This collaboration enables EY to target two strategically valuable areas of unstructured data in financial services: ESG and financial reports. "Elastic allows us to help banks streamline reporting on their ESG commitments, including internal metrics and details from the full value chain," he says.

For financial reports, EY's AI solution, powered by Elastic, can extract data and insights from complex documents with numerous tables. "Imagine extracting 14 key variables from a 40-page PDF or comparing information across multiple reports from different years. This is where we save clients significant time and resources," says Venkatesh.

In addition to delivering cutting-edge technology, Elastic offers seamless integration with other components in EY's generative AI stack, including LamaIndex, a data framework for building LLM applications, LangChain, and open-source vector databases such as LanceDB.

> With Elastic, we can simultaneously promote responsible AI and innovation. It means our clients can adhere to regulations, act faster, and mitigate internal risks.

Vishaal Venkatesh
GenAI Manager

elastic

EY observed significant improvements in three key areas following its deployment of Elastic:

- Approximately 10-15% gains in accuracy across various document and extraction methods;
- Elastic's RAG setup was able to produce results 3 times faster when compared to Native RAG approaches;
- Improved scalability, demonstrating the ability to handle diverse document formats and seamlessly scale across various organizational units.

Speed to value was another crucial advantage. "Traditionally, integrating vendor solutions presents a significant challenge due to lengthy ramp-up periods," says Venkatesh. "However, Elasticsearch's unique infrastructure significantly accelerated our development process, enabling us to realize the benefits much faster."

At the heart of the generative AI revolution

Venkatesh adds that Elasticsearch's core strength lies in its ability to perform large-scale information retrieval while simultaneously enabling the integration of diverse AI technologies throughout the entire development lifecycle. This integrated approach extends to offering vital observability and security functionalities, ensuring optimal performance across EY's custom-built AI pipeline.

"By leveraging a robust and secure solution like Elastic, we eliminate the need to invest significant resources in building and maintaining these critical components from scratch," Venkatesh says. "Furthermore, relying on a proven and trusted solution like Elastic mitigates the risk of unforeseen issues with homegrown components."

Throughout the project, EY was able to call on the support of Elastic account managers and technical experts whenever necessary. "Our collaboration with Elastic proved to be exceptionally smooth due to the strong foundation established from the outset. The combined efforts of Elastic solution architects and account managers ensured we could swiftly navigate the complexities of generative AI and achieve rapid prototyping," says Venkatesh.

Above all, it helps keep EY at the vanguard of the generative AI revolution. "With Elastic's powerful search and retrieval functionalities, combined with machine learning for vector embedding, we're ideally positioned to equip our clients with responsible generative AI tools that further innovation, unlock value, and mitigate risk."

> Working with the Elastic team has been fantastic. They help ensure that we're constantly pushing the boundaries of what is possible and responsible with generative AI.

Vishaal Venkatesh
GenAI Manager