



esynergy Prepares Proposals 75% Faster with AI Document Retrieval from AWS and DataStax

Learn how DataStax Astra DB helped esynergy enhance document retrieval efficiency with AI and vector search on AWS.

Overview

[esynergy](#) builds data, platform, and artificial intelligence (AI) solutions that deliver business value. In constant pursuit of optimizing its internal processes, esynergy wanted to enhance its document retrieval process with AI to speed up the development of sales proposals. The company chose [DataStax Astra DB](#)—a database that brings real-time vector data to AI apps—for its Sales Copilot solution.

Astra DB integrations with [Amazon Bedrock](#) allowed esynergy to leverage [Claude 2 by Anthropic](#) to generate informed, accurate responses. [LangChain](#) played a crucial role in orchestrating interactions between these components. This combination allowed esynergy to move from proof of concept to production in just 37 days, enabling the company to prepare proposals up to 75 percent faster.

Opportunity | Enhancing Sales Efficiency with Accelerated Document Retrieval

esynergy leverages a community of subject matter experts to design and implement technology solutions that support business objectives. Known for its problem-solving approach to industry challenges, esynergy wanted to build a solution to improve the document retrieval process for customers—and for its own sales team.

"Over the past 10 years, our teams have generated numerous documents stored in various formats within SharePoint, such as PowerPoint proposals, project narratives, visual artifacts, and PDF case studies," said Prasad Prabhakaran, generative AI practice lead at esynergy.



About esynergy

esynergy is a technology consultancy that drives measurable business value for its clients. Working with enterprises and scale-ups in highly regulated industries, the company works with leading cloud providers like AWS to build scalable platforms, products, and services.

AWS services used

- [Amazon Bedrock](#)
- [AWS PrivateLink](#)

Benefits

- 37 days to develop a production-ready solution
- 75% reduction in proposal-preparation time
- 1–2 days to complete a proposal

"When customers ask for information on practices like migrating workloads to the cloud, our sales representatives often must consult multiple sources and navigate through a vast array of documents on SharePoint, which can take days. We knew many of our customers were dealing with similar issues."

Existing AI solutions failed to meet esynergy's requirements for context awareness, security, and scalability. The company needed a cost-effective, scalable database to handle large volumes of data and provide rapid, accurate search results using AI. The solution also had to integrate seamlessly with existing tools, such as SharePoint for document storage and Amazon Web Services (AWS).

Solution | Empowering Sales Teams with AI-Powered Responses

To address the challenges faced by its sales teams, esynergy developed Sales Copilot, an AI-powered conversational assistant that enables rapid access to relevant customer data and generates informed responses. The solution leverages DataStax Astra DB, a scalable and cost-effective vector database, to handle large volumes of data and provide fast, accurate search results using AI.

"We chose Astra DB for its serverless architecture, which significantly reduced our operational costs," said Prabhakaran. "Additionally, its caching capabilities and ability to manage multiple turnarounds to the large language model (LLM), Claude 2, allowed us to deliver fast and accurate responses to our users, even under high concurrency. These features made Astra DB the perfect fit for our Sales Copilot application."

The Sales Copilot application seamlessly integrates with the organization's SharePoint collaboration platform, ingesting customer profiles, communication records, and documents like case studies. LangChain, a framework for building applications with LLMs, plays a crucial role in the data preprocessing stage, enabling efficient splitting of files into manageable chunks for subsequent processing. These chunks are then transformed into dense vector representations using embeddings generated by Astra DB, which are optimized for fast and accurate retrieval. The "retrieve-augment" architecture relies on esynergy's customized component to return only the most relevant data points from Astra DB.

The Claude 2 LLM, integrated through Amazon Bedrock, crafts relevant responses tailored to the user's query, informed by the retrieved customer data and documents. The modular components of LangChain make the implementation of this architecture straightforward, allowing for rapid prototyping and experimentation with different strategies. "Along with Astra DB, Amazon Bedrock provided us with the flexibility to experiment with different AI models and find the one that best suited our needs," said Prabhakaran.



"It allowed us to easily integrate Claude 2 into our solution, providing powerful natural language processing capabilities without the hassle of managing the underlying infrastructure."

The Sales Copilot application integrates seamlessly with esynergy's existing AWS infrastructure, including [AWS PrivateLink](#) for enhanced security and compliance. By leveraging the easy-to-use features of Astra DB, Amazon Bedrock, and LangChain, esynergy was able to develop and deploy the Sales Copilot application within a short timeframe, enabling rapid prototyping and iteration based on user feedback and changing requirements.

Outcome | Generating Proposals Faster, From Weeks to Days

esynergy's Sales Copilot solution has yielded significant benefits for its sales teams. Overall, Sales Copilot has sped up the document retrieval process by 75 percent, and proposals can now be developed in a day or two rather than a week. The application has also enhanced the accuracy of search results, reducing the risk of misinformation and improving the quality of customer interactions. The serverless architecture of Astra DB has led to significant cost savings in data storage and processing. Today, esynergy's sales representatives can quickly access the information they need, leading to better customer service and faster response times. Analytics show regular usage of Sales Copilot at esynergy, with at least 80–90 daily prompts.

"Getting started with the proof of concept was simple with DataStax. We completed development and deployment in just 37 days," said Prabhakaran. "We were able to get up and running with the entire ecosystem without needing extensive support from DataStax. The ease of use allowed us to focus on understanding how we could leverage the platform more effectively to meet our customers' needs."

"Getting started with the proof of concept was simple with DataStax. We completed development and deployment in just 37 days."

Prasad Prabhakaran

Generative AI Practice Lead,
esynergy

About the AWS Partner DataStax

DataStax serves production-level AI applications with Astra DB, its powerful vector database fueled by real-time, real-life customer and streaming data. Built on Apache Cassandra, Astra DB is the data engine for responsive AI applications that consistently deliver accurate results, without artificial hallucinations.



esynergy Prepares Proposals 75% Faster with AI Document Retrieval from AWS and DataStax

Learn how DataStax Astra DB helped esynergy enhance document retrieval efficiency with AI and vector search on AWS.

Overview

[esynergy](#) builds data, platform, and artificial intelligence (AI) solutions that deliver business value. In constant pursuit of optimizing its internal processes, esynergy wanted to enhance its document retrieval process with AI to speed up the development of sales proposals. The company chose [DataStax Astra DB](#)—a database that brings real-time vector data to AI apps—for its Sales Copilot solution.

Astra DB integrations with [Amazon Bedrock](#) allowed esynergy to leverage [Claude 2 by Anthropic](#) to generate informed, accurate responses. [LangChain](#) played a crucial role in orchestrating interactions between these components. This combination allowed esynergy to move from proof of concept to production in just 37 days, enabling the company to prepare proposals up to 75 percent faster.

Opportunity | Enhancing Sales Efficiency with Accelerated Document Retrieval

esynergy leverages a community of subject matter experts to design and implement technology solutions that support business objectives. Known for its problem-solving approach to industry challenges, esynergy wanted to build a solution to improve the document retrieval process for customers—and for its own sales team.

"Over the past 10 years, our teams have generated numerous documents stored in various formats within SharePoint, such as PowerPoint proposals, project narratives, visual artifacts and PDF case studies," said Prasad Prabhakaran, generative AI practice lead at esynergy.



About esynergy

esynergy is a technology consultancy that drives measurable business value for its clients. Working with enterprises and scale-ups in highly regulated industries, the company works with leading cloud providers like AWS to build scalable platforms, products, and services.

AWS services used

- [Amazon Bedrock](#)
- [AWS PrivateLink](#)

Benefits

- 37 days to develop a production-ready solution
- 75% reduction in proposal-preparation time
- 1–2 days to complete a proposal

"When customers ask for information on practices like migrating workloads to the cloud, our sales representatives often must consult multiple sources and navigate through a vast array of documents on SharePoint, which can take days. We knew many of our customers were dealing with similar issues."

Existing AI solutions failed to meet esynergy's requirements for context awareness, security, and scalability. The company needed a cost-effective, scalable database to handle large volumes of data and provide rapid, accurate search results using AI. The solution also had to integrate seamlessly with existing tools, such as SharePoint for document storage and Amazon Web Services (AWS).

Solution | Empowering Sales Teams with AI-Powered Responses

To address the challenges faced by its sales teams, esynergy developed Sales Copilot, an AI-powered conversational assistant that enables rapid access to relevant customer data and generates informed responses. The solution leverages DataStax Astra DB, a scalable and cost-effective vector database, to handle large volumes of data and provide fast, accurate search results using AI.

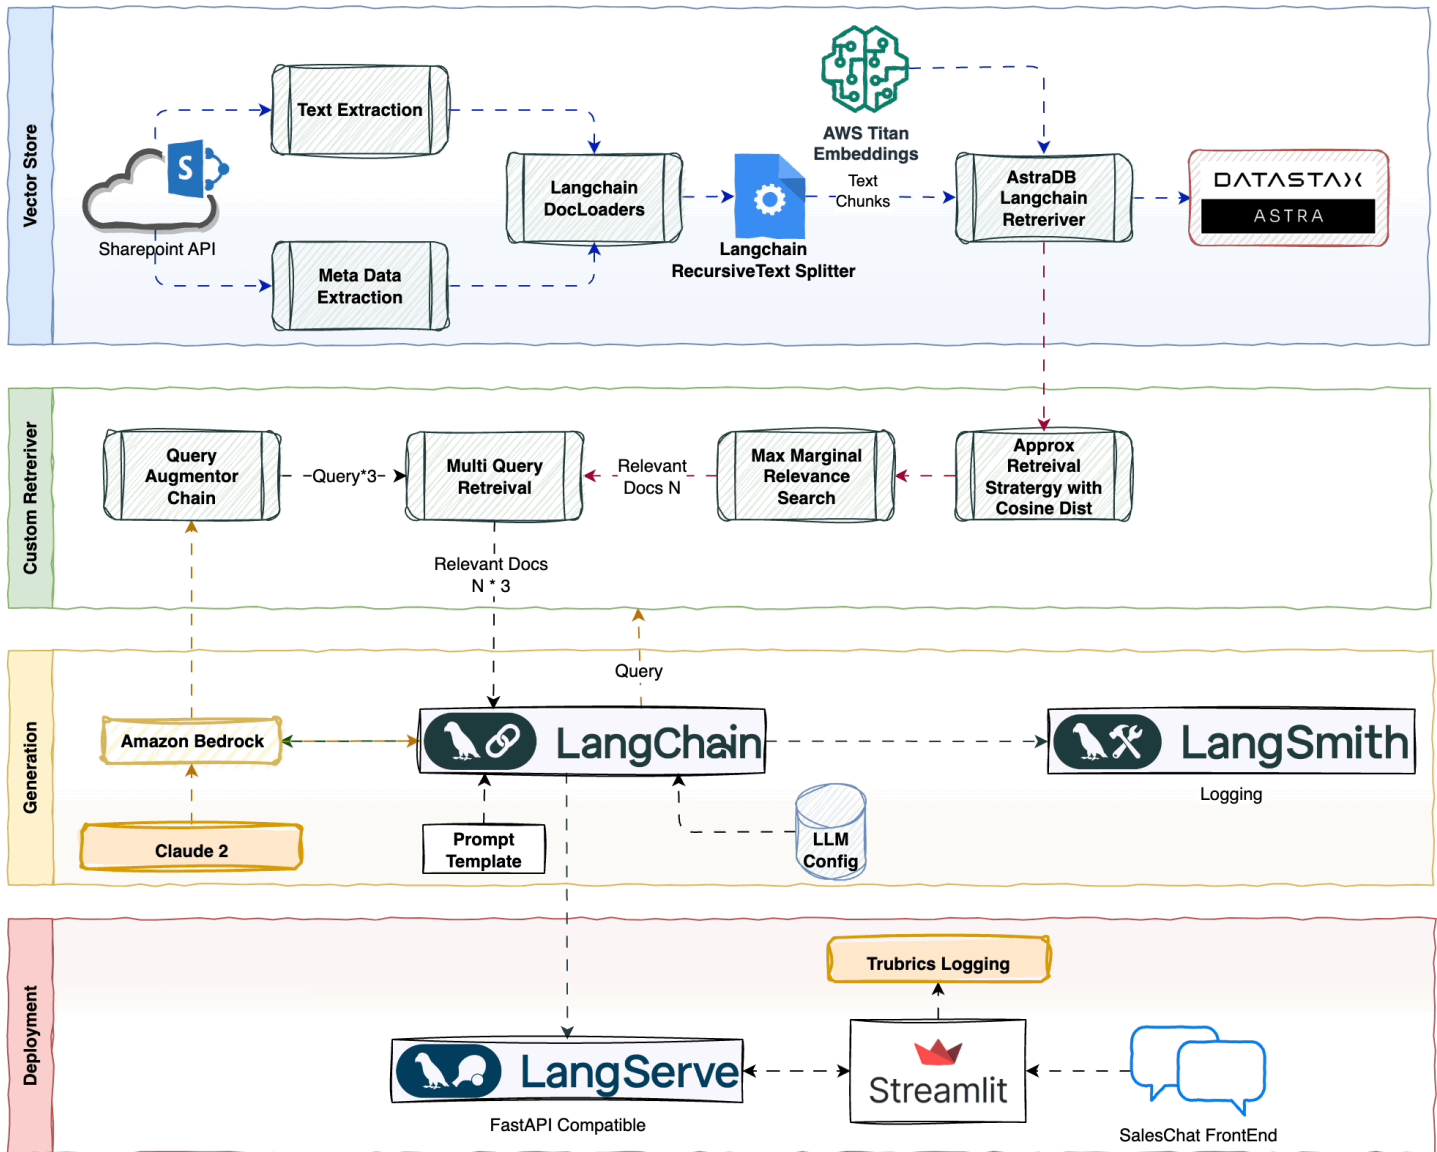
"We chose Astra DB for its serverless architecture, which significantly reduced our operational costs," said Prabhakaran. "Additionally, its caching capabilities and ability to manage multiple turnarounds to the large language model (LLM), Claude 2, allowed us to deliver fast and accurate responses to our users, even under high concurrency. These features made Astra DB the perfect fit for our Sales Copilot application."

The Sales Copilot application seamlessly integrates with the organization's SharePoint collaboration platform, ingesting customer profiles, communication records, and documents like case studies. LangChain, a framework for building applications with LLMs, plays a crucial role in the data preprocessing stage, enabling efficient splitting of files into manageable chunks for subsequent processing. These chunks are then transformed into dense vector representations using embeddings generated by Astra DB, which are optimized for fast and accurate retrieval. The "retrieve-augment" architecture relies on esynergy's customized component to return only the most relevant data points from Astra DB.

The Claude 2 LLM, integrated through Amazon Bedrock, crafts relevant responses tailored to the user's query, informed by the retrieved customer data and documents. The modular components of LangChain make the implementation of this architecture straightforward, allowing for rapid prototyping and experimentation with different strategies. "Along with Astra DB, Amazon Bedrock provided us with the flexibility to experiment with different AI models and find the one that best suited our needs," said Prabhakaran.

"It allowed us to easily integrate Claude 2 into our solution, providing powerful natural language processing capabilities without the hassle of managing the underlying infrastructure."

The Sales Copilot application integrates seamlessly with esynergy's existing AWS infrastructure, including [AWS PrivateLink](#) for enhanced security and compliance. By leveraging the easy-to-use features of Astra DB, Amazon Bedrock, and LangChain, esynergy was able to develop and deploy the Sales Copilot application within a short timeframe, enabling rapid prototyping and iteration based on user feedback and changing requirements.





Outcome | Generating Proposals Faster, From Weeks to Days

esynergy's Sales Copilot solution has yielded significant benefits for its sales teams. Overall, Sales Copilot has sped up the document retrieval process by 75 percent, and proposals can now be developed in a day or two rather than a week. The application has also enhanced the accuracy of search results, reducing the risk of misinformation and improving the quality of customer interactions. The serverless architecture of Astra DB has led to significant cost savings in data storage and processing. Today, esynergy's sales representatives can quickly access the information they need, leading to better customer service and faster response times. Analytics show regular usage of Sales Copilot at esynergy, with at least 80–90 daily prompts.

"Getting started with the proof of concept was simple with DataStax. We completed development and deployment in just 37 days," said Prabhakaran. "We were able to get up and running with the entire ecosystem without needing extensive support from DataStax. The ease of use allowed us to focus on understanding how we could leverage the platform more effectively to meet our customers' needs."

"Getting started with the proof of concept was simple with DataStax. We completed development and deployment in just 37 days."

Prasad Prabhakaran

Generative AI Practice Lead,
esynergy

About the AWS Partner DataStax

DataStax serves production-level AI applications with Astra DB, its powerful vector database fueled by real-time, real-life customer and streaming data. Built on Apache Cassandra, Astra DB is the data engine for responsive AI applications that consistently deliver accurate results, without artificial hallucinations.