Harness structured and unstructured data for Trusted AI with Teradata and NVIDIA

As generative and agentic AI evolve, enterprises seek value from unstructured data, which will account for over 80% of global data in 2025. Teradata's Enterprise Vector Store offers a cost-effective solution for managing this data at scale, integrating structured and unstructured data capabilities. This enables seamless storage, search, and retrieval of billions of vectors, driving actionable insights and enhanced customer experiences. The platform supports advanced use cases like augmented call centers and delivers Trusted AI at scale, accelerating customer experience innovation and meeting the demand for faster ROI.

Get the freedom to unlock unstructured data

Unstructured data, growing three times faster than other data types, presents enterprises with significant challenges in management and analysis. Traditional vector databases often fall short due to data movement requirements, lack of integration with non-vector data, and high costs.

Teradata's Enterprise Vector Store uses NVIDIA NeMo Retriever for data extraction, reranking and embedding, optimizing RAG performance and transforming customer service by unlocking insights from unstructured data

Teradata's Enterprise Vector Store addresses these issues with a scalable, in-database solution that supports the full lifecycle of vector data management. It integrates structured and unstructured data, enabling enterprises to gain holistic insights while reducing data movement, enhancing privacy, and minimizing costs. The platform supports AI development frameworks like LangChain, and leading data Retrieval Augmented Generation (RAG) solutions powered by NVIDIA NeMo Retriever, part of the NVIDIA AI Enterprise software platform. It provides a leading information retrieval solution with high accuracy and data privacy, which enables enterprises to extract information from multi-modal PDFs, simplify complex AI workflows, and generate business insights in real-time.

Integration with NVIDIA NeMo Retriever

NVIDIA NeMo Retriever is a collection of generative AI solutions that provide leading information retrieval with high accuracy and data privacy — built with NVIDIA NIM™. They can be fine tuned in combination with community or custom models to build scalable data extraction pipelines, document ingestion and Retrieval Augmented Generation (RAG) applications.

Through the Teradata Enterprise Vector Store, customers can connect their data to RAG-backed generative AI applications for sophisticated use cases like augmented call centers. This unified platform enables organizations to drive operational efficiency, enhance customer experiences, and achieve breakthrough business outcomes. Additional benefits of the NVIDIA partnership include:

- Accelerated unstructured data processing: NVIDIA NeMo Retriever provides GPU-accelerated PDF parsing, vector embeddings, object detection, table extraction, OCR, and chart interpretation, transforming unstructured documents into actionable insights at scale.

- Optimized embedding and retrieval: NVIDIA NeMo Retriever and Teradata Vantage enable fast, accurate, and scalable embedding generation and search and retrieval with intelligent search using algorithms such as in-database Vector-Distance, K-Means, and hierarchical navigable small world (HNSW).
- Tracked changes with temporal embeddings: Trust and explainability are enhanced with Teradata Vantage's temporal embeddings that create a timestamp of embeddings, so changes to PDF documents can be tracked.

Build to an agentic AI future for dynamic CX and more

Agentic AI autonomously solves complex challenges with advanced reasoning and planning. Teradata's Enterprise Vector Store supports this by providing a scalable, trusted data foundation for agentic AI. It integrates vast datasets with end-to-end processing capabilities, enhancing problem-solving and operationalizing agentic AI. Known for enterprise-scale analytics, Teradata offers cost-effective scalability for inference economics, driven by generative AI applications. Agentic AI can revolutionize productivity, efficiency, and customer experience across various use cases, including augmented call centers.

Augmented call center use case

Generative AI agents enhance customer service by handling complex tasks, personalizing interactions, and assisting human agents. Enterprises can use Retrieval Augmented Generation (RAG) to provide precise, context-aware answers from multiple data sources, improving resolution times and customer satisfaction. Teradata's Enterprise Vector Store transforms customer service by unlocking insights from unstructured data. It simplifies data ingestion, vectorization, and intelligent search, combining structured and unstructured data into actionable insights. Enhanced by NVIDIA NeMo Retriever microservices, it delivers high scalability, cost efficiency, and improved agent productivity.

Maximize value with less data movement

Vector databases are crucial for generative AI but scaling them is complex. Teradata's Enterprise Vector Store overcomes these challenges with a scalable, in-database solution that processes billions of vectors seamlessly. It integrates with existing systems, minimizing data movement and enhancing security. The platform combines data embedding, vector indexing, and intelligent search for structured and unstructured data. With hybrid deployment options, it supports flexible scaling across cloud and on-premises environments. This comprehensive approach enables cost-efficient scaling, optimized workflows, and maximized value from unstructured data, empowering enterprises to unlock their data's full potential.
Test-drive Teradata's free hands-on demo environment.

About Teradata

At Teradata, we believe that people thrive when empowered with trusted information. We offer the most complete cloud analytics and data platform for AI. By delivering harmonized data and Trusted AI, we enable more confident decision-making, unlock faster innovation, and drive the impactful business results organizations need most.
See how at Teradata.com.

17095 Via Del Campo, San Diego, CA 92127 Teradata.com

teradata.