



News Release

Teradata Launches Integrated Enterprise Vector Store to Help Customers Be Ready to Implement Trusted Agentic AI

2025-03-03

Designed for cost-effective, sub-second response times at all data volumes, the new offering can solve multi-dimensional, complex problems by combining structured and unstructured data

Enterprise Vector Store will use NVIDIA NeMo Retriever microservices for accelerated compute, optimized RAG. Teradata (NYSE: **TDC**) today announced **Teradata Enterprise Vector Store**, an in-database solution that brings the speed, power and multi-dimensional scale of Teradata's hybrid cloud platform to vector data management, a crucial element for **Trusted AI**, with future expansion to include integration of **NVIDIA NeMo Retriever microservices, part of the NVIDIA AI Enterprise software platform**. Featuring the ability to process billions of vectors and integrate them into pre-existing enterprise systems, with response times as quick as in the tens of milliseconds, Enterprise Vector Store is designed to cost-effectively deliver the sophistication required for getting real value out of complex, multifaceted business challenges.

The offering creates a single, trusted repository for all data and builds on the strong support Teradata offers today for retrieval-augmented generation (RAG), while working towards dynamic agentic AI use cases, such as "augmented call center" (see example below).

Vector stores are foundational for any organization looking to leverage agentic AI, but most vector stores require trade-offs that make it prohibitively hard or expensive to use in solving the most challenging (and potentially the most lucrative) business problems. They can be fast, but only for small data sets. Or they can manage vector volumes, but not at the speed that agentic AI use cases require. The real magic happens when organizations can apply both lightning-fast speed and massive compute to unstructured datasets that hold real value when combined



with mission-critical structured data.

“Vector stores are at the root of how we bind truth to generative AI models and agentic AI. They are essential to any data management practice, but their impact is limited when they are slow or siloed,” said Louis Landry, Teradata’s CTO. “Teradata’s long-standing expertise in high concurrency and linear scale, as well as the critical ability to harmonize data and support RAG, means Teradata Enterprise Vector Store delivers on the dynamic, trusted foundation large organizations need for agentic AI.”

Teradata’s **Enterprise Vector Store** is designed to be a performant way to enable use cases that require vector capabilities and RAG applications. With cost-efficient scaling and near seamless integration built-in, Enterprise Vector Store is expected to help enterprises maximize value and insight from unstructured data while reducing spend. Given Teradata’s advantage in hybrid, Enterprise Vector Store is a natural choice for organizations that want to scale flexibly across cloud and on-premises environments, building towards an agentic AI future while making the most of current infrastructure.

By managing unstructured data in multi-modal formats — text, video, images, PDFs, and more — Teradata’s Enterprise Vector Store unifies structured and unstructured data for holistic analysis. It also:

- Engages with the full lifecycle of vector data management, from embedding generation and indexing to metadata management and intelligent search
- Processes this work within the existing Teradata system, which thrives in flexible deployment options including cloud, on-premises, or hybrid
- Supports industry-leading frameworks like LangChain and RAG, along with the comprehensive data management and governance practices needed for **Trusted AI**
- Adds planned temporal vector embedding capabilities, which is designed to boost trust and explainability by tracking changes to data over time, improving accuracy and decision making.

A scalable, in-database vector solution built with NVIDIA AI

Teradata Enterprise Vector Store is expected to integrate NVIDIA NeMo Retriever to provide a leading information retrieval solution with high accuracy and data privacy, enabling enterprises to generate business insights in real-time. Developers can fine-tune NeMo Retriever microservices in combination with community or custom models to build scalable document ingestion and RAG applications which can be connected to proprietary data wherever it resides. NVIDIA NeMo Retriever extraction is designed to enable customers to use information and insights from unstructured data sources such as PDFs, enabling developers to build RAG-based applications which leverage real-time knowledge appended with information from across the corporate IT estate.

“Data is essential to accurate inference for AI applications,” said Pat Lee, Vice President of Strategic Enterprise



Partnerships at NVIDIA. “Teradata Enterprise Vector Store, integrated with NVIDIA AI Enterprise and NVIDIA NeMo Retriever, can unlock the institutional knowledge stored in PDFs and other unstructured documents to power intelligent AI agents.”

Use Case: Augmented Call Center

The augmented call center use case demonstrates how the Teradata Enterprise Vector Store uses agentic AI and RAG to transform customer service to be faster, more efficient and tailored to each customer’s needs. AI agents also enable upsell and cross-sell opportunities during customer interactions.

For example, an insurance company stores contracts for its millions of customers in PDF format in an object store. It also uses a hybrid data platform for mission-critical customer 360 data. When a customer calls in, a multi-agent system uses lightning-fast access (as low as tens of milliseconds) to harmonized data to provide precise, context-aware answers to each individual customer.

- “Hello, how can I help you today?”
 - “Customer Interaction” Agent communicates in real time with the customer using a natural language interface which is powered by popular LLMs running as **NVIDIA NIM** on NVIDIA accelerated compute.
- “I’m traveling to Malaysia. Does my insurance cover medical expenses? Should I add anything?”
 - “Contract Analyzer” Agent quickly retrieves coverage details from the PDF copy of the contract using RAG with Enterprise Vector Store, which has extracted the information from PDFs and stored as embeddings using NVIDIA NeMo Retriever in Teradata Enterprise Vector Store.
 - “Insurance Advisor” Agent uses reasoning and decision making to recommend adding dental coverage for the duration of the trip, using a propensity-to-buy model and Teradata’s trusted predictive and explainable AI capabilities.
- “Ok, let’s add dental please.”
 - “Actions” Agent uses operational analytics and customer 360 (structured) data in Teradata’s hybrid environment to create a contract for customer signature.

Availability

Teradata Enterprise Vector Store is now available in private preview, with general availability expected in July.

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**.

Media Contact

Jennifer Donahue

Teradata

Jennifer.Donahue@Teradata.com