

# Unlock Scalable, Multimodal Intelligence With Teradata Enterprise Vector Store

## Enterprises are rapidly adopting generative and agentic AI, but most struggle to operationalize unstructured and multimodal data at scale.

Text, documents, images, audio, video, and other unstructured data now account for most enterprise information yet remain siloed in disconnected systems. Standalone vector databases introduce complexity, data movement, and rising costs, while limiting access to structured enterprise data and governance controls.

Teradata Enterprise Vector Store solves these challenges by embedding vector storage, search, and retrieval directly inside Teradata Vantage®. By unifying structured and unstructured data in a single, governed platform, organizations can build accurate, context-aware AI applications without duplicating infrastructure or pipelines. The result is faster time to value, lower operational cost, and a trusted data foundation for enterprise scale RAG and agentic AI.

## Capabilities designed for enterprise AI

Teradata Enterprise Vector Store is built on core capabilities that operationalize enterprise-scale RAG and agentic AI. Together, these capabilities unify multimodal data, deliver efficient scale and price performance, and provide an open developer experience that accelerates AI innovation without sacrificing governance or trust.

### Unified multimodal data across cloud and on-premises deployments

Vectors and structured data live in one database, creating a unified semantic layer that combines relational data, metadata, and multimodal embeddings. By eliminating isolated vector systems and external pipelines, Teradata enables richer context, hybrid and fusion search, and more accurate RAG and AI agents grounded in complete, governed enterprise data.

### Efficient scale and price performance

Built on Teradata Vantage® proven MPP architecture, Enterprise Vector Store scales linearly to billions of vectors while supporting high throughput, highly concurrent workloads. Organizations can operationalize vector workloads alongside transactional and analytical use cases, delivering consistent performance and predictable cost efficiency as unstructured and multimodal data grows.

### Superior developer experience

Open integrations with LangChain, Python, SQL, and SDKs enable developers, data scientists, and analysts to build, test, and deploy RAG pipelines and AI agents quickly. By removing unnecessary infrastructure and orchestration, Teradata accelerates experimentation and time to value while maintaining enterprise grade governance, security, and control.

## Multimodal enterprise data foundation

Teradata Enterprise Vector Store enables enterprises to generate, store, and query vector embeddings across text, documents, images, audio, and video while keeping vectors and structured data together in one governed system. This unified, multimodal foundation expands the semantic context available to AI models and agents, enabling richer understanding, more accurate RAG, and enterprise-scale AI without moving data to external vector databases across cloud, on-premises, or hybrid environments.

To support diverse enterprise needs, Teradata offers flexible ingestion and embedding options. Unstructured provides ETL pipelines with a drag-and-drop UI for extracting, enriching, chunking, and embedding content across 70+ file types. NVIDIA (NIM™/NVAIE) delivers high-performance embeddings, reranking, and guardrails through GPU-based microservices for performance-intensive workloads. Cloud service providers (CSPs) enable access to cloud-hosted model endpoints through native consoles and LangChain workflows, offering flexible model choice and elastic scaling.

## Hybrid search for accurate, context aware retrieval

Enterprise AI requires more than similarity search alone. Teradata Enterprise Vector Store delivers hybrid search by combining vector similarity, lexical search, metadata filtering, and structured SQL queries into a single retrieval experience. This approach dramatically improves accuracy for real-world enterprise queries by blending semantic meaning with exact business context. Hybrid and fusion search allow applications and AI agents to retrieve insights across tables, documents, and embeddings simultaneously without stitching together multiple systems. By grounding AI responses in both meaning and enterprise facts, Teradata reduces hallucinations and increases trust, making it ideal for mission critical use cases such as customer experience, compliance, fraud detection, and operational intelligence.

## LangChain integration for RAG and AI agents

Teradata Enterprise Vector Store integrates natively with LangChain, enabling developers to build and operationalize RAG pipelines and AI agents directly on enterprise data. Agents can retrieve relevant context from vectors and structured tables, reason over the results, and execute governed actions without manual orchestration.

This integration allows AI systems to move beyond retrieval into agentic execution, supporting multistep workflows that combine search, reasoning, and action. Developers can prototype quickly using familiar Python and SQL interfaces, then scale seamlessly to production using Teradata's proven MPP architecture. By unifying data, vectors, and agents in one platform, Teradata accelerates AI innovation while maintaining enterprise grade governance, security, and control.

## Maximize value with reduced data movement

Teradata Enterprise Vector Store is purpose built for enterprise scale, not experimentation. Unlike standalone vector databases, it is fully integrated into Teradata Vantage®, eliminating data duplication, reducing cost, and enabling AI workloads to run alongside operational and analytical workloads. Plus, Teradata Enterprise Vector Store offers:

- **An integrated—not point—solution:** Vectors, structured data, metadata, and SQL in one platform
- **Enterprise scale performance:** Linear scalability to billions of vectors and high dimensional embeddings
- **Hybrid and fusion search:** Unified retrieval across structured and unstructured data
- **Open developer experience:** Native LangChain integration, Python and SQL interfaces
- **Deployment flexibility:** Cloud, on-premises, or hybrid environments with consistent capabilities
- **A trusted foundation:** Governance, security, and control built into the data layer.

Teradata Enterprise Vector Store provides the trusted data foundation enterprises need to operationalize RAG and agentic AI at scale, combining structured and unstructured multi-modal data into real business outcomes.

## About Teradata

Teradata Autonomous AI and Knowledge Platform activates enterprise intelligence by unifying data, knowledge and business context to achieve tangible outcomes. With Teradata, organizations can provide agents with full context for impact when it matters. Our solution lets businesses connect and scale on premises, in the cloud, or through a hybrid approach. Teradata delivers real business value with AI. [Learn more at Teradata.com.](https://www.teradata.com)