News Release

Teradata Enables AI Agents to Autonomously Process Text, Images, and Audio at Enterprise Scale

2026-03-09

Teradata Enterprise Vector Store unifies structured and unstructured data with agentic capabilities across hybrid environments, enabling rapid deployment of production-ready AI systems

SAN DIEGO, March 9, 2026 /PRNewswire/ -- Teradata (NYSE: TDC) today announced new agentic and multi-modal data capabilities for Teradata Enterprise Vector Store, a unified solution that enables organizations increasingly to harness the full potential of generative AI and autonomous agents across hybrid, cloud, and on-premises environments. Integrated with Unstructured, this release marks a significant evolution in Teradata's enterprise AI infrastructure, combining multi-modal data integration, agentic capabilities, and advanced hybrid search to unlock new levels of intelligence and efficiency.

New Features

Teradata Enterprise Vector Store delivers a complete pipeline—from embedding generation to indexing, metadata management, and AI framework integration—with the following advanced features:

- Unstructured Integration: Automated ingestion and processing of documents, PDFs, images, and audio with upcoming video support
- Hybrid Search: Combines semantic and lexical search with metadata-driven techniques for more accurate, context-aware retrieval
- Multi-Modal Embeddings: Support for text, image, and audio embeddings with richer semantic representations
- Higher Embedding Dimensions: Up to 8K dimensions for enhanced accuracy and nuance
- LangChain Integration: Direct integration enabling enterprise-scale RAG pipelines, rapid prototyping to production, and agentic execution that extends beyond search—allowing AI agents to retrieve context and operationalize outcomes through governed actions and autonomous workflow orchestration

Why Now: The Enterprise AI Challenge

With the explosive growth in unstructured data—which Gartner estimates is growing at three times the rate of structured data—traditional vector databases have proven insufficient for enterprise-scale AI deployments, particularly as AI models become increasingly multimodal, processing text, images, audio, and video simultaneously.

As adoption surges – **nearly 80% of companies are already deploying AI agents with most projecting 100%+ ROI from agentic AI initiatives** – external research shows enterprises face significant barriers to scaling: fragmented data silos, limited scalability, and lack of unified access to structured and unstructured content alike. These constraints prevent organizations from realizing the full potential of agentic AI at enterprise scale. Closing that gap requires an enterprise vector store built for the scale, performance, and governance that modern AI demands.

Why Teradata: Industry-Leading Scale and Performance

Teradata was built for exactly this moment. **Forrester research** notes that "high-end scale and performance still require considerable effort, especially when supporting tens of billions of data points (vectors)." Most vector solutions hit practical limits at a few hundred million embeddings. Teradata Enterprise Vector Store was engineered for enterprise-scale AI, with the ability to ingest millions of documents, thousands of files per hour, and multi-modal data streams with appropriate configuration and data characteristics.

In combination with Vantage's proven enterprise architecture, Teradata Enterprise Vector Store has been shown to deliver: linear scalability across billions of vectors and high-dimensional embeddings; 1,000+ concurrent queries without performance degradation; optimized cost structures that eliminate duplicated infrastructure; and enterprise-grade governance across cloud, on-premises, and hybrid environments.

How Enterprises Will Use It

Teradata's integrated approach, in partnership with Unstructured, eliminates the complexity of point solutions by automatically parsing and transforming unstructured data into high-quality embeddings and unifying structured and unstructured data within a single governed platform. This enables AI agents to autonomously access comprehensive enterprise context and execute complex workflows without manual intervention.

Process Diverse Data Types at Scale:

Through the Unstructured partnership, organizations can automatically parse and transform documents, PDFs, images, and audio into high-quality embeddings at enterprise scale. This enables AI systems to reason across vastly different data sources with shared semantic understanding.

Real-World Example: Healthcare Visual Q&A: Medical institutions combine structured patient records with clinical

notes, medical images, and audio dictations to support faster diagnosis and treatment planning. Teradata-LangChain agents orchestrate a governed workflow that applies vision models, runs multi-modal vector search, and grounds responses with trusted documentation—delivering explainable, source-traceable results.

Enable Autonomous Workflows: AI agents can independently retrieve context, take action, and orchestrate complex workflows through seamless LangChain integration, transforming AI from simple chatbots into fully autonomous, production-grade systems capable of sophisticated decision-making.

Real-World Example: Insurance Claims Automation: Claims adjudication agents process damage photos and policy PDFs alongside structured claims data, extracting information from images and documents while cross-referencing coverage rules and claim history—delivering faster, explainable decisions with full audit compliance.

Deliver Context-Aware Intelligence: Hybrid search combines semantic vector search with lexical and metadata-driven techniques while fusion search enables unified retrieval across structured and unstructured data. This multi-layered approach can help dramatically improve reliability and reduce AI hallucinations by incorporating comprehensive context into every query.

Real-World Example: Defense Intelligence: Military organizations transform static camouflage doctrine into adaptive, intelligence-driven protection by having troops capture images of camouflaged assets via secure apps. These images are processed in the Enterprise Vector Store alongside terrain patterns and threat signatures, with LangGraph-orchestrated agents delivering real-time survivability guidance at the speed of the battlefield.

Eliminate Data Silos: Unlike traditional vector databases that operate in isolation, Teradata's agentic enterprise vector store enables AI agents to simultaneously pull insights from tables, logs, documents, images, and metadata within a single governed environment—eliminating data duplication and pipeline complexity.

Real-World Example: Business Loyalty Agents: Financial services firms build governed agents that combine unstructured policy definitions with structured business data to answer complex questions like loyalty discount eligibility—bridging the gap between documents and databases that SQL alone cannot address.

Accelerate Development and Deployment: Open integrations with SQL, Python, and LangChain enable developers to design and orchestrate autonomous agent workflows that seamlessly access both structured and multi-modal unstructured data using familiar tools and skills—from rapid prototyping to production deployment across cloud, on-premises, or hybrid environments without architectural constraints.

Real-World Example: From Prototype to Battlefield: Defense organizations rapidly deploy secure mobile apps that enable troops to capture field imagery, which is instantly processed through the Enterprise Vector Store with

LangGraph-orchestrated agents delivering real-time tactical guidance—demonstrating how familiar development tools enable fast deployment of mission-critical AI systems in demanding environments.

Executive Commentary

"We're entering an era where AI agents will become the primary interface for enterprise intelligence—autonomously orchestrating workflows, making decisions within defined governance frameworks, and uncovering insights across every data type," said Sumeet Arora, Chief Product Officer at Teradata. "Stand-alone vector databases can't deliver on this vision. Teradata Enterprise Vector Store fundamentally reimagines how enterprises operationalize AI by unifying structured and multi-modal unstructured data with autonomous agent capabilities within a single governed platform. Organizations can now move from prototype to production-grade agentic systems in some cases within hours, not months—while maintaining the governance, security, and sovereignty that mission-critical AI demands."

"Enterprises shouldn't have to choose between data security and AI readiness. By embedding Unstructured natively inside Teradata Enterprise Vector Store, Teradata customers get production-quality, AI-ready data at scale, with no external tools, no data leaving the platform, and no compromise on governance," said Brian Raymond, Founder and CEO of Unstructured.

Availability

New agentic and multi-modal capabilities for Teradata Enterprise Vector Store are generally available to Teradata customers starting April 2026.

For more information, visit: **https://www.teradata.com/platform/clearscape-analytics/enterprise-vector-store**

About Teradata

Teradata empowers enterprises to turn intelligence into autonomous action, grounding AI agents in deep business context and trusted data. As AI agents multiply, Teradata is the context engine, governance layer, and performance backbone that companies need now. The Teradata Autonomous AI and Knowledge platform puts AI into production across cloud, on-premises, and hybrid environments.

MEDIA CONTACT

January Machold

january.machold@teradata.com

View original content:**https://www.prnewswire.com/news-releases/teradata-enables-ai-agents-to-autonomously-process-text-images-and-audio-at-enterprise-scale-302707423.html**

SOURCE Teradata Corporation