



News Release

Teradata Transforms Tera into an Agentic Coworker; Introduces Tera Context Engine and Tera Harness

2026-09-22

Designed specifically for enterprise data work, Tera brings governed business knowledge and intelligent execution to every data role, at lower cost and greater speed

SAN DIEGO, Sept. 22, 2026 /PRNewswire/ -- **Teradata** (NYSE: TDC) today announced a major evolution of **Tera**, transforming it into an agentic coworker for enterprise data work. Where general-purpose AI assistants generate answers, Tera delivers outcomes. Through natural language and guided execution, everyone from business analysts to platform engineers and database administrators can analyze data, build AI applications, operate infrastructure, and automate complex workflows – all from a single governed environment that reaches enterprise data across platforms, not just within Teradata. Industry expertise and business knowledge are embedded across every interaction, secured by enterprise identity and policy, so no deep Teradata specialization is required.

Tera includes three key capabilities:

- ***NEW*** Tera Context Engine, a vendor-neutral context and orchestration layer that gives AI governed business knowledge;
- ***NEW*** Tera Harness, an intelligent execution layer that routes work across the right skills, tools, data, and models; and
- ***NEW*** Agent Skills, purpose-built for data engineering, data analysis, and data science.

Tera addresses two requirements enterprise AI has lacked: the ability to understand the business and the ability to act on that understanding. Teradata customers gain more reliable AI outcomes, better economics from every model interaction, and faster time to value, supported by **Teradata AI Services** for organizations that want to accelerate deployment.

Benchmark Results

Tera was purpose-built for enterprise data work, optimizing execution across analytics, vectors, and models. In testing on SWE-bench Pro using the same Opus 5 model, Tera consumed 73% fewer tokens than **Claude Code** while achieving higher task completion rates, completed work 42% faster, and incurred 58% lower total cost. On data-eng-bench, a benchmark developed by Snowflake Labs and Bespoke Labs for data pipeline engineering, Tera delivered 53% lower cost per reliably solved task than **Snowflake Cortex Code** using Opus 5, based on their published benchmark data. Across data-eng-bench and ADE-bench, Tera earned benchmark-leading accuracy, achieving the highest Pass³ score on data-eng-bench and tying for the top score on ADE-bench. Full benchmark methodology and results, including materials for independent validation, are available [here](#).

Teradata Executive Quote

"Most enterprises are not starting from scratch with AI. They are dealing with tools that do not work together and a skills gap that makes those tools hard to use at scale. Tera is designed to work across that environment, putting business context, intelligent execution, and pre-built expertise into the hands of every person working with data. And equally important is what enterprises do not give up: control over their models, their data, and where everything runs. The result is AI that actually gets work done, at lower cost, with less overhead."

-Sumeet Arora, CPO, Teradata

Tera Details and Components

Tera is part of the Teradata Autonomous Knowledge Platform and brings together three purpose-built capabilities that can be combined or used independently: Tera Context Engine for governed enterprise knowledge, Tera Harness for intelligent execution, and Tera Agents for specialized expertise. Together, they make Tera an agentic coworker that can understand the business, coordinate the right capabilities, and help drive data work through to an outcome. Organizations retain the ability to choose which models they use, where workloads run, and how enterprise data is accessed, across cloud, on-premises, and sovereign environments.

Tera executes AI natively within the Teradata environment, including the Teradata Console for database admins, running analytic and ML workloads directly on enterprise data without external model calls. This eliminates data movement and latency, and reduces hallucination risk for quantitative tasks like forecasting, regression, and segmentation.

- ***NEW*** Tera Context Engine

Tera Context Engine is an open, neutral context and orchestration layer that operates above and across the data environments enterprises already use. It connects databases, structured and unstructured data platforms, pipeline engines, catalogs, models, and AI agents without moving data, replacing existing technologies, or committing to a single vendor. Importantly, Tera Context Engine is not limited to Teradata-managed data, giving organizations a consistent foundation of trusted knowledge regardless of where their

data lives.

Native Context Graph: By transforming scattered enterprise data into reusable business knowledge, Tera Context Engine connects metadata, lineage, semantics, and business meaning as relationships rather than flattening them into relational structures. Unlike catalogs or ontologies alone, it keeps that knowledge connected to governance, lineage, and access controls as it moves across systems and agents. It also reads from and writes back to existing systems of record, continuously learning from real enterprise use and improving over time with significantly less human effort to maintain. The result is more complete and explainable context, with AI outputs that carry provenance organizations can understand and defend.

Neurosymbolic Models: The new offering combines autonomous semantic mapping and deterministic retrieval with Industry Knowledge Models, the product of Teradata's human-validated expertise and engagement with some of the world's most complex enterprises. Rather than requiring agents to construct business meaning from scratch, Industry Knowledge Models provide a symbolic knowledge foundation that combines explicit enterprise knowledge with statistical AI, capturing the specific terminology, relationships, policies, and operating conditions that define how an industry works. Enterprises get the benefit of fewer agent errors and faster time to value.

Deterministic Query Economics: Governed semantic context and deterministic retrieval shift work away from expensive probabilistic reasoning toward governed execution paths, so models spend fewer calls on established definitions. This delivers higher accuracy, lower inference costs, and more predictable economics as agentic workloads scale. Critical for regulated industries, AI outputs stay auditable and traceable to source data because lineage, access policies, and compliance controls travel with the knowledge as it moves across systems.

Governed Agentic Automation: Beyond supplying context for AI to consume, Tera Context Engine enables agents to act on governed context, automating data product creation, pipeline specifications, and validation controls while applying enterprise definitions and policies consistently. The result is faster delivery of trusted data products with less manual preparation and governance effort, at enterprise scale.

- ***NEW* Tera Harness**

Unlike general-purpose agents that generate responses and leave users to act on them, Tera Harness turns intent into outcomes. It maintains context across workflows and coordinates the right skills, tools, data, and models at each step without manual routing.

Cost-Optimized Loop: Tera Harness limits unproductive iterations before they translate into cost, batching independent work, removing model and tool interactions that do not improve the outcome, and

bounding execution based on task progress. Rather than relying on repeated LLM reasoning at every step, Tera applies 84 proven execution patterns pre-inference, so agents start with an execution plan. The result is governed, multi-step work that completes reliably at enterprise scale, without runaway model iterations.

Loop-Embedded Guardrails: Before execution begins, Tera Harness applies pre-inference guardrails, loads relevant memory, and plans the work ahead. Safety, policy controls, and human approvals are embedded directly inside the agent loop, blocking high-risk actions before they execute. Unlike standard agents that rely on self-policing or external filters, Tera enforces governance at the point of execution, protecting against unauthorized actions and destructive operations.

Go-Native Performance Core: Tera Harness is built on a Go-native engine and gRPC to deliver high-concurrency agent execution and low-latency communication. Where standard agent frameworks run each agent as a separate process, driving up infrastructure costs and latency at scale, Tera supported 512 concurrent agents on a single 8-vCPU VM, serving 279 tool calls per minute. That means enterprises can scale agentic workloads without the infrastructure costs and latency spiral that other frameworks introduce with every agent added.

Durable Agentic Execution: As execution unfolds, Teradata Harness monitors progress, recovers from failures rather than abandoning tasks, and provides full observability into how work was completed. State management and checkpointing are built directly into the runtime layer, so agents can run for days or weeks, pause for human approval at zero compute cost, and resume exactly where they left off after infrastructure failures. Unlike most agent frameworks, which require a separate orchestration layer to survive failures while continuing to consume cloud resources during long-running waits, Tera handles durability natively.

For enterprises managing AI at scale, this translates directly into fewer model calls per outcome, lower token costs, better accuracy, and faster completion.

- ***NEW* Agent Skills**

Agent Skills are reusable, task-specific capabilities that package decades of Teradata knowledge into AI-callable functions. They can be invoked by any agent or harness, and accessed directly by users through natural language, making advanced data and platform work accessible without requiring deep technical specialization. Rather than reasoning from scratch, models receive clear guidance for common data and platform tasks, improving consistency and reducing the risk of going off-track. Skills load automatically based on the task at hand, and Tera routes each to the correct models and tools. MCP connectivity allows organizations to extend Tera further with their own tools and integrations.

Tera includes two categories of purpose-built Tera agents that draw on these skills. Platform Agents handle the operational work of managing the Teradata environment, including workload tuning, compute sizing, telemetry, and FinOps. Analytics Agents tackle data work directly, from natural language to SQL and Python to query optimization.

Teradata AI Services

For organizations looking to accelerate deployment, Teradata AI Services is designed to help customers avoid the experimentation trap and move directly toward production outcomes. AI Services helps identify the use cases where governed context will create measurable value, configure Industry Knowledge Models, and get enterprise knowledge into production faster powered by Tera. Teradata's AI Value Engineering methodology ensures AI programs are designed and developed for production-readiness, not experimentation, and pre-developed agents and tooling built for every stage of the AI development lifecycle accelerate time-to-customer-value.

Availability Details

Tera Context Engine, Tera Harness and Agent Skills will be available in Q4 2026.

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 foundation, governance layer, and performance backbone that companies need now. The Teradata Autonomous Knowledge Platform puts AI into production across cloud, on-premises, and hybrid environments.

Forward-Looking Statements

This press release contains "forward-looking statements" about the expectations, beliefs, plans, and intentions relating to Tera, including Tera Context Engine, Tera Harness, and Agent Skills. Such statements include statements regarding future product availability, capabilities, and offerings, and expected benefits to Teradata customers. Forward-looking statements are subject to known and unknown risks and uncertainties and are based on potentially inaccurate assumptions that could cause actual results to differ materially from those expected or implied by the forward-looking statements. If any such risks or uncertainties materialize or if any of the assumptions prove incorrect, Teradata's results could differ materially from the results expressed or implied by the forward-looking statements made. Teradata undertakes no obligation, and does not intend to update the forward-looking statements. Factors that may cause actual results to differ materially from those in any forward-looking statements include: (i) delays and unexpected difficulties and expenses in executing the planned product capabilities and offerings, (ii) changes in the regulatory landscape related to AI and (iii) uncertainty as to whether customer adoption will justify the investments in the product capabilities and offerings. Further information on factors that could affect Teradata's financial and other results is included in the filings Teradata makes with the

Securities and Exchange Commission from time to time.

The Teradata logo is a trademark, and Teradata is a registered trademark of Teradata Corporation and/or its affiliates in the U.S. and worldwide.

MEDIA CONTACT

Jennifer Donahue

Jennifer.Donahue@Teradata.com

View original content to download multimedia:<https://www.prnewswire.com/news-releases/teradata-transforms-tera-into-an-agentic-coworker-introduces-tera-context-engine-and-tera-harness-302885832.html>

SOURCE Teradata Corporation