

teradata.

Autonomous Intelligence **World Tour**

 Experience

Karachi, Pakistan | September 29, 2026



The Autonomous Enterprise Starts Here: Lakehouse, Semantics & Agentic AI at Scale

Paul Ibberson, Enterprise Architect, EMEA, Teradata

Your data can do more than report what happened

Evolve from systems of record to systems of intelligence, and turn every data point—from claims and contracts to customer interactions—into action.



Enterprise Challenges



Can your AI platform deliver enterprise outcomes with velocity?

AI not equal to ROI!



Does your AI run where your data and knowledge lives?

AI is data starved!



Are you locked into single vendor usage pricing, limiting control over cost and risk?

Agents never sleep!

AI breaks traditional data sovereignty models

The gap between AI ambition and the ability to govern it is widening.



This isn't just about where data lives. It's about how AI, analytics, and governed data operate together, under your control.

teradata.

Autonomous Intelligence

World Tour 2026

Why AI sovereignty and governance fail at scale

Traditional sovereignty model

AI, analytics, and data live in separate systems—you need to move data to AI

Governance applied after the fact or patched together

Sovereignty weakened by movement and duplication

New sovereignty model

Unified analytics and AI on shared data

Open, flexible data foundation with governance built in by design

Maintain full data and AI sovereignty with one control plane for governance

Teradata Vision



Autonomous

Deliver autonomous experience across platform layers and end-to-end data and AI processes



AI

Build, activate and manage AI outcomes through Agents, AI/ML, Apps, and Marketplaces



Knowledge

Enabled through building systems of intelligence informed by knowledge that adapt, self-evolve and are agentic

Autonomous AI + Knowledge

 Humans

 AI

 Workspaces

 Agents

 Agent Execution

 AI Tools

 Knowledge Context

 Knowledge Foundation


Active & Elastic
Compute
CPUs & GPUs


Object & Block
Storage


On-Prem


Multi-Cloud


Sovereign Cloud


In-Database AI
Capabilities

 Gemini


ANTHROPIC

 Mistral AI


LLaMA
by Meta

Open source

 OpenAI



Teradata Principles




Colocation

AI lives where data lives.



Performance

Data, vectors, models everywhere.



Context

Knowledge informed by enterprise data



Choice

Compute, storage, models, engines, location. You choose.



Workspace

Build, deploy, govern AI together.

Why Teradata



Versatile Engine

Agents don't sleep. We provide the best engine for agents, SLM, SQL, Vector, AI execution



Best Context

Autonomous data products, Industry-grounded, knowledge informed by enterprise data, closes a critical gap in agentic AI adoption



Tooled for Agentic AI

Build, activate and manage AI outcomes through Agents, AI/ML, Apps, and Marketplaces

AI
Services

Multi-Cloud | Hybrid Cloud | On-Prem | Sovereign AI + Data

Trust | Resilience | Scale | Fit for Purpose | Velocity

Introducing Teradata Factory

**Sovereign Data and AI.
Open architecture.
One integrated system.**

- Modern NVIDIA + Dell powered rack
- Specialized AI compute Infrastructure
- Integrated analytics + AI + Lakehouse
- Governance and workload protection by design
- Runs AI where enterprise data lives

teradata.



DELL
Technologies



The Teradata Lakehouse vision

Be the Lakehouse Platform of choice for Humans and Agents — standards-based and fully governed

1

Open Lakehouse — bring compute to data

Position Teradata compute to run anywhere data lives — without owning storage and catalog, with zero ETL

2

Apache Iceberg and Delta as the standard formats

Open Iceberg or Delta standards become the connective tissue across all analytics environments

3

Agentic AI is the real endgame

Be the AI Lakehouse of choice — the platform that powers knowledge and AI at Agentic scale, ready for enterprise SLAs

Teradata's vision is to be the platform of choice for Knowledge and AI — open, standards-based data architecture delivering high-performance, enterprise-grade SLAs for data stored in an open lakehouse

Teradata Lakehouse strategy



Open and
standards
driven




Leader in
price-
performance



Centralized
security and
governance




Open table format: Unparalleled choice for data access and analytics



Reduced cost and complexity

OTFs provide a flexible, cost-efficient, and easy-to-use storage abstraction layer, enabling customers to manage enterprise data with minimal data movement, reduced complexity, and cost



Open and connected

Teradata's open and connected approach enables **cross-read**, **cross-write**, and **cross-query** of data stored in a variety of OTFs and open catalogs in multi-cloud and multi-data-lake environments



World-class analytics

Explore AI use cases with a future-ready, truly modern data strategy with unmatched agility and flexibility to deliver Trusted AI at scale—all without the need to move, replicate, or transform data

teradata.

Autonomous Intelligence

World Tour 2026

AI use cases built on a unified lakehouse



Intelligent operations

- Real-time risk and fraud detection
- Predictive maintenance and anomaly detection
- SLA-protected tactical decisioning with in-place inference



Customer & revenue growth


- Customer 360
- AI-driven personalization across channels
- Consistent metrics and insights



Risk & compliance

- Workload isolation ensuring reliable, compliant performance
- AI models running on governed, auditable datasets
- Automated regulatory and financial reporting

Many data sources, multiple analytic techniques, and a single query



```
SELECT * FROM
( ON (
  SELECT * FROM
  (
    ON LOCAL.Customer AS InputTable
    ON ( SELECT * FROM
      ( ... ) ) AS FitTable
  ) AS a
) ) AS predict

JOIN ( SELECT * FROM nPath (
  ON ICEBERG.ClickstreamJSON
  PARTITION BY cust_id ORDER BY date
) ) AS np

JOIN SPARK.SocialMediaSentiment AS sm
WHERE predict.value = 1
AND np.path LIKE '%cancel%'
AND sm.sentiment_score < 0.4;
```

What this demonstrates:
Query data across systems | Apply multiple analytic techniques | Leverage specialized processing engines | Deliver results through a single query experience

How Teradata Fabric Extends Your Analytics Ecosystem

Four core capabilities enable seamless analytics across distributed data environments.



Orchestrate query execution across Teradata as well as other databases and platforms without migrating data



Launch queries in place and execute pushdown processing across hybrid and multi-cloud environments



Join and transfer data across tables and systems from multiple sources



Reduce data movement and minimize egress charges

Knowledge Context

Why enterprises need trusted knowledge and context



Agents are the new data consumers

AI agents operate 24/7, generating sustained high-concurrency workloads

They need trusted context for every decision: semantics, lineage, policy, and KPIs



Most agents fail in production

Early deployments fail when scaled due to a missing governed knowledge foundation

Data lives in silos — structured, unstructured, open-format — with no unified context



The cost of fragmentation

Fragmented stacks create data movement, duplicated governance, and semantic drift

Cross-capability joins at decision time are impossible without colocation

teradata.

Autonomous Intelligence

World Tour 2026

Autonomous knowledge makes agentic AI enterprise-ready

Knowledge gives AI agents and people shared, trusted business context to reason and act in real time

-  Unifies data, metadata, policies, and decision traces into a living context graph
-  Eliminates schema-first modeling and per-use-case rebuilds as AI evolves
-  Enables agents to discover, understand, and act on enterprise knowledge without human mediation



Introducing Teradata Context Engine

Context Engine is a new context layer across the Teradata Autonomous Knowledge Platform that transforms enterprise and industry knowledge into reusable AI-ready data products — giving AI agents the business meaning, relationships, usage patterns, and governance context they need to act with precision, not guesswork.

Teradata Context Engine



Unify Data

Structured + unstructured data
across all environments



Add Meaning

Business definitions, lineage,
and governance context



Build Trust

Governed, policy-aware AI
outputs you can act on



Scale AI

Reusable context across
production AI workloads

Data Access



Data Understanding

AI Experimentation



Precise AI Outcomes

Raw Data



AI-Ready Data Products

Manual Prep



Reusable Semantic Context

The Context Gap in Action

Without Context Engine

Structured + unstructured data in silos
Metadata is fragmented across systems
Business definitions remain inconsistent
Teams manually prepare and reconcile data
AI lacks trusted business context

Result

Generic, inconsistent answer
not trusted or actionable



With Context Engine

Unified structured + unstructured data
Shared business definitions and relationships
Usage-aware contextual metadata
Lineage and policy-aware governance
Minimal manual data preparation

Result

Precise, explainable answer
consistent, trusted, ready to act on

Tera

Autonomous Intelligence
World Tour 2026




Tera

teradata.

© 2026 Teradata. All rights reserved.

Introducing Tera

The agentic coworker for
enterprise data work



What Tera can do

Tera provides specialized experiences for analytics, development, and platform operations within a single governed environment.

01



Analyze

Explore data, uncover insights, create visualizations, and turn questions into analysis through natural language.

02



Build

Create AI models, agents, applications, and code faster with Teradata skills and guided execution.

03



Operate

Manage workloads, monitor usage, tune queries, investigate incidents, and optimize platform operations.


04




Automate

Coordinate multi-step data and AI workflows using specialized skills, agents, MCP, and intelligent orchestration.

Teradata AI Studio



Tera Delivers Better Outcomes at Three Levels



Level 3: Tera maximizes value from every token

Tera's intelligent harness turns token efficiency into business value by applying the right expertise, maintaining context, and delivering better outcomes from every AI dollar spent.

Level 2: Tera accelerates critical jobs-to-be-done across every data role

From analysts to engineers, Tera delivers dramatically faster outcomes through specialized skills and context. Turning weeks into days and days into hours.

Level 1: Tera makes it easy to use Teradata. No expertise required.


Tera unlocks the full potential of your Teradata platform by making advanced data and AI capabilities accessible through natural language and guided execution.

FROM HUMAN-LED TO AGENT-LED ENTERPRISE

EXAMPLE: CREDIT RISK MEMO PROCESS


Before agents

RM performs each task manually
2–4 days per memo



After agents

AI agent performs all steps;
human reviews completed tasks



Impact
20–60%
productivity gain
30%
faster decisioning speed

Modern AI Outcomes

Modern AI outcomes harness multiple AI capabilities to drive measurable ROI.



Predictive AI



Generative AI



Agentic AI



Multi-modal AI



Vision AI



Conversational AI



Hybrid RAG



Real-time Intelligence



Decision Optimization




Sovereign AI

Examples of Modern AI Outcomes

Defense Intelligence

Visual Intelligence driven protection of \$15 million defense assets



Agentic Logistic

Agentic AI for port congestion solution with potential cost saving of \$120 billion per year



R&D AI-Assistant

Knowledge Intelligence unifies vehicle telemetry with vectorized vehicle specifications



AI Agent for every customer

A personal financial-advice agent to every banking customer



Agentic hyper-Personalized Healthcare

An agent for every patient, monitoring health and fitness data predicting likely medical issues




Maritime AI Surveillance

Geo-spatial analysis of high-frequency vessel movement data



The autonomous enterprise does not start with an agent.

It starts with open data, shared meaning and governed context.



Lakehouse → Semantics → Trust → Action

